Essays on Media Dynamics

Sara Shahanaghi

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy
under the Executive Committee
of the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2022

Abstract

Essays on Media Dynamics

Sara Shahanaghi

In this dissertation, I consider how the dynamic nature of news provision may give rise to unique strategic behavior on the part of news firms. Specifically, I consider three separate dynamic environments, corresponding to the three different chapters of this dissertation.

In the first chapter, I present a continuous-time model to understand the strategic forces behind reporting errors. News firms are rewarded for reporting before their competitors, but also for making reports that are credible in the eyes of consumers. Errors occur when firms *fake*, reporting a story despite lacking evidence. I establish existence and uniqueness of an equilibrium, which is characterized by a system of ordinary differential equations. Errors are driven by both a lack of commitment and by competition. A lack of commitment power gives rise to errors even in the absence of competition: firms are tempted to fake after their credibility has been established, capitalizing on the inability of consumers to detect fake reports. Competition exacerbates faking by engendering a preemptive motive. In addition, competition introduces observational learning, which causes errors to propagate through the market. The equilibrium features rich dynamics. Firms become gradually more credible over time whenever there is a preemptive motive. The increase in credibility rewards firms for taking their time, and thus endogenously mitigates the haste-inducing effects of preemption. A firm's behavior will also change in response to a rival report. This can take the form of a *copycat effect*, in which one firm's report triggers an immediate surge in faking by others.

In the second chapter, we study news firms' reporting behavior, including their propensity to misreport, when they are reputation-driven. In our model, a news firm (sender) dynamically learns about a state and reports to a consumer (receiver). Senders are concerned with their reputation at the end of the game, and must choose when to time their report. We find that in equilibrium, the sender *fakes*, i.e., report despite being ignorant of the state, with positive probability in every period. This faking in turn leads to a higher level of misreporting than if the sender were instead truthful. We further find the sender's reputations is endogenously rewarded for both speed and accuracy, and thus we provide a microfoundation to the speed-accuracy tradeoff in the media setting. Finally, we consider dynamics in the sender's strategy, finding that the sender becomes more truthful, and thus less prone to misreporting, as time passes.

In the third chapter, I present a dynamic model of reputation-driven media bias. I extend Gentzkow and Shapiro's (2006) model to a dynamic setting and characterize the equilibrium under the assumption that firms are myopic. Equilibrium bias is driven by two factors: an appeal to the consumer's prior bias and an effort to appear consistent. I then consider a forward-looking firm, and show that it will behave as if it is myopic in equilibrium. Finally, I consider an extension of the baseline model in which firms face a lying costs. I show that assuming the firm is myopic is with loss in this case, and that this reversal is driven by the equilibrium separation that lying costs induce.

Table of Contents

List of Figures

Acknowledgements

I would like to thank my advisors Yeon-Koo Che and Navin Kartik for their invaluable guidance throughout my PhD. I am also grateful to my committee members Andrea Prat, Qingmin Liu, and Elliot Lipnowski for their many insights. Finally, I would like to thank the participants of the Columbia Microeconomic Theory Colloquium and Seminar for their helpful comments and suggestions regarding the three chapters presented herein.

Chapter 1: Competition and Errors in Breaking News

1.1 Introduction

> What a newspaper needs in its news, in its headlines, and on its editorial page is terseness, humor, descriptive power, satire, originality, good literary style, clever condensation, and accuracy, accuracy, accuracy!
>
> — Joseph Pulitzer

Accuracy is often considered to be the core tenet of news media. This belief is widely held by consumers of news: when asked in a 2018 Pew survey, the majority of respondents listed accuracy as a primary function of news, valuing it over thorough coverage, unbiasedness, and relevance.

Despite this, public perceptions of news accuracy are not favorable. In a 2020 Survey, 38% of respondents stated that they go into a news story thinking it will be largely inaccurate. While many factors may contribute to this skepticism, consumers express particular concern about hasty reporting: 53% of respondents believe that news breaking too quickly is a major source of errors.

These concerns are supported by a multitude of instances in which news media have made major factual errors. In the immediate aftermath of the 9/11 attacks, cable news stations made several statements that were false: NBC reported an explosion outside the pentagon, CNN reported a fire outside the national mall, and CBS claimed the existence of a car bomb outside the state department. Erroneous reporting has been endemic to terrorist attacks in general, with news media misidentifying perpetrators or other key details of the Boston bombings, Sandy Hook massacre, London bombings, and Oklahoma City bombings. Furthermore, such errors are not limited to terrorist attacks. Notoriously, in 2004 CBS news, under the direction of Dan Rather, published the Killian Documents, a collection of memos which called into question George W. Bush's military record. These documents could never be authenticated and were widely believed to be forged.

More recent media blunders are ever present: in 2017, ABC news falsely reported that Michael Flynn would testify that Donald Trump had directed him "to make contact with the Russians." In 2019, ABC News headlined its nightly news broadcast with what it claimed to be exclusive footage of the ongoing air strikes on Syria. It was later revealed that this footage was from a machine gun convention in Oklahoma.

While such errors are commonplace, they are also costly to news firms. For one, exposure of errors can be reputationally damaging. This was acutely true of the *Rolling Stone* scandal, in which the magazine falsely accused a group of University of Virginia students of sexual assault. Not only was the journalistic failure widely reported by other firms, the error resulted in several publicized lawsuits against the magazine. Furthermore, major errors often lead firms to part with valuable journalists in an effort to protect their reputations. This was evident in the terminations of Dan Rather and Brian Ross —both lead journalists at major news stations—following their respective reporting blunders.

The objective of this paper is twofold. First, I seek to understand why reporting errors are pervasive despite their costliness to firms. In particular, I explore how strategic forces can induce firms to commit errors that are completely avoidable. My second objective is to understand *when* reporting errors are most probable, and relatedly, when firms are less trustworthy. That is, I seek to understand both the dynamics of reporting errors and the environmental factors that can make them more prevalent.

Model To answer these questions, I present a dynamic model of breaking news. I consider a continuous-time setting where multiple firms dynamically and privately learn about a story and must choose if and when to report it. Firms learn by seeking confirmation that the story is true. Reporting errors occur when firms *fake*, i.e., report the story despite lacking confirmation. Because reports are public, firms also learn by observing the reports of their competitors. I thus account for an important feature of the newsroom setting: firms learn privately but also observationally.

Firms in this model seek viewership. Error-prone reporting conflicts with this objective, and is thus costly to the firm, in two ways. First, errors harm firms ex post (after they have been exposed). This ex-post cost captures the detrimental effect of errors on a firm's future livelihood. Importantly, error-prone reporting is also costly ex ante (before errors can be unearthed). This is due to the fact that a firm's viewership hinges on its *credibility*, i.e., the consumer's belief that the report is not fake. This belief is formed rationally with knowledge of the firm's reporting strategy: firms who fake more achieve lower credibility in equilibrium. By making this assumption, I take the stance that a story is valued to the extent that there is trust in a firm's journalistic standards, a notion that is informed by consumers' demonstrated preference for accurate news.

Finally, this model accounts for one of the most salient qualities of the breaking news problem: competition. All else equal, a firm who preempts its rivals (e.g., by being the first to report) is rewarded with greater viewership. This allows us to understand the impact of competition on the propensity of firms to err. Doing so is especially pertinent given the rise of digital news. Since the ascent of the internet, there has been a documented shift from print to digital news.[1] This shift has arguably contributed to a news industry where firms feel greater pressure to get stories out quickly in order to beat out competitors. This is due to the fact that, while print news is limited to daily publication at most, digital news faces no such constraints.[2] By considering a continuous-time setting, one can better understand 24-hour news environment, where preemptive concerns are not only present but ceaseless.

Analysis I analyze this model, establishing both existence and uniqueness of an equilibrium. Under this equilibrium, fake reports do not occur at set times, but are rather distributed continuously over time. This mixing implies an *indifference condition*: at any time in which the firm must fake, it must be indifferent between faking immediately and after some short wait. Formally, this condition implies an ordinary differential equation (ODE) on the firm's reporting behavior. I thus show that the equilibrium is characterized by a system of ODEs, a result which is central to our analysis and guides many of the economic implications that follow.

[1] While 16% of 2018 survey respondents often receive news from print newspapers, 33% do so from news websites.
[2] This is also true of TV news, which remains the most popular news medium in the United States.

Economic Implications I find that errors are strategic responses to two features of the news environment: a lack of commitment by firms, and competition.

To this end, I begin by showing that competition alone is not responsible for reporting errors. In particular, if the ex-post cost of error is relatively small —because consumers are less aware or critical of them —even a monopolist will fake. Such errors are driven by a firm's *inability to commit* to a reporting strategy: a firm is tempted fake after its credibility has been assessed. This is due thte fact that firms cannot observe whether a firm is faking, and thus the firm is not directly punished for doing so. I substaniate the notion that a lack of commitment causes errors by proving that a firm who can commit will always report truthfully, and thus never err.

I then show that competition exacerbates errors, and does so through two separate channels. First, competition can give rise to a preemptive motive in equilibrium: firms have an incentive to speed up their reporting in order to beat out competitors. This incentive for speed induces firms to fake and thus err. Second, competition causes errors through another, less obvious channel: observational learning. When one firm reports a story, other firms become more confident that the story is true. This increased confidence in turn yields firms more likely to fake. I thus find that observational learning exacerbates errors not by giving rise to them in the first place, but by causing existing errors to propagate through the market.

This paper also sheds light on the dynamics of reporting behavior and credibility. These dynamics take two different forms in equilibrium: gradual changes that happen in the absence of new reports and discrete changes that occur in response to a new report.

I first show that firms become gradually more truthful —i.e., less inclined to fake—as time passes. Furthermore, firms become more credible over time whenever preemptive concerns are present. In other words, consumers are less trusting of reports that are made quickly. This model thus justifies consumers' expressed concerns about hasty reporting. The reason for this gradual improvement in credibility lies in the firms incentives. The risk of being preempted introduces an endogenous cost to delay. That is, the firm must somehow be compensated for this cost to ensure that its indifference condition is satisfied. This is achieved by means of increasing credibility. That is, increasing credibility mitigates the haste-inducing effects of preemption.

In addition to this gradual increase in credibility, dynamics can take a second form: discrete changes in a firm's reporting behavior and credibility in response to a rival report. This can entail a *copycat effect*, in which one firm's report causes an instantaneous boost in faking by others. The copycat effect implies that when one firm's report is quickly repeated by other firms, such follow-up reports will often lack credibility because they are not independently verified. It illustrates that firms can herd on both the reports themselves and the *timing* of their reports. This provides an explanation for the "clustering" of reporting errors that can occur in breaking news.[3]

In addition to these core results, I consider comparative statics and an extension of the model. I find that, unsurprisingly,credibility is improved by both a higher ex-post cost of error and a higher learning ability. I also further explore the role of competition by considering the marginal effect of an additional firm in the market. Whenever preemptive concerns are present, adding a competitor will make each individual firm more likely to fake early on by increasing the preemptive threat they face. However, this is mitigated later on by the effects of observational learning: existing firms are able to learn that the story is false more quickly by observing the silence of an additional competitor, which will yield them less willing to fake. Finally, I extend the model to allow for heterogeneity in firms' ability to learn. This extended model gives rise to an intuitive result: firms with greater ability to learn are also more credible in equilibrium. Though there are many potential reasons why ability and accuracy can correlate in the market for news, this model provides a novel explanation: firms with lower ability face a greater preemptive threat, and are thus more willing to fake.

Related Literature The preemption literature has modeled a variety of scenarios, including R&D races ([1]), technology adoption ([2]), the strategic exercise of options ([3]), and financial bubbles ([4]). This paper contributes to this literature in two key ways. The first is in the endogeneity of the payoff function. In the existing literature, a player's decision to preempt does not affect its underlying payoff function. That is, the benefit of preempting may be stochastic (e.g., [3]), but

[3]Examples of this include the reporting surrounding the Boston bombings and the 2000 US presidential election.

it is exogenous. In my setting, however, a firm's payoff from reporting hinges on the consumer's beliefs about its reporting behavior. Such beliefs are important in the market for news because consumers may not be able to immediately observe the quality of a news report, e.g., whether it was verified before being reported. This assumption has implications for the nature of the firm's incentives. While in the existing literature, players earn some exogenous benefit from delaying their actions which counteracts the incentive to preempt, this is not true in our setting. Rather, I find that even if no such benefit exists exogenously, it will arise endogenously.

This paper is not the first to consider observational learning in a preemption setting. In [5], firms can only observe their own payoffs, and thus draw inferences about the payoffs of their competitors by observing when and whether they act. Meanwhile, in [6], players receive breakthroughs which are privately observed, and thus at every moment are uncertain about how much competition they face. In contrast, I assume that firms learn observationally about their own payoffs. This form of observational learning is also present in [7], in which players privately observe bad-news Poisson signals in a winner-takes-all setting. Like this paper, the lack of action by one's opponent implies that acting is unprofitable. However, in contrast to [7], I consider a payoff environment that is not necessarily winner-takes-all (i.e., nth movers may enjoy positive payoffs). It is precisely because of this more general payoff structure that observational learning gives rise to herding on both actions themselves and the timing of these actions. In this sense, this paper also connects to the literature on herding with endogenously-timed decisions ([8], [9], [10]). In particular, the notion that an action by one individual can trigger others to quickly follow suit arises in [8]. While such behavior is efficient in their setting, that is not the case in ours, where it can cause errors to propagate through the market.

This paper contributes to a recent literature on preemption in news ([11], [12], [13]). As in this paper, [11] considers a setting where firms dynamically learn about a single story and must decide whether and when to report it. Meanwhile, [13] suppose that firms report on a sequence of potential political scandals, finding that competition improves the quality of information but deteriorates consumer welfare. I contribute to this literature by modeling two vital features of the breaking news environment, i.e., the role of credibility and observational learning. These two features are what drive many of the economic implications of this paper, including the importance of commitment, herding, and dynamics in firm behavior.

This paper also contributes more broadly to a literature on competition in news. This literature is surveyed by [14], with more recent contributions by [15], [16], [17], and [18]. [17] and [16] specifically consider the effects of competition on news accuracy. In both papers, firms compete for the attention of consumers and face constraints or costs to accuracy. Meanwhile, in my setting, accuracy is not intrinsically costly. Rather, accurate reporting entails an indirect cost, namely that of being preempted. I contribute more generally to this literature in two ways. First, I consider the effects of competition on a different notion of accuracy, namely the prevalence of factual errors. Second, this paper also sheds light on the dynamics of firm behavior. This allows one to understand the effects of competition on not only on news quality as a whole, but also on its time path.

Finally, this paper connects broadly to the literature on the strategic provision of information. Unlike frameworks where a sender seeks to induce a particular action from receivers ([19], [20]), firms in my model treat information as a good, aiming to maximize its appeal to consumers. This notion underlies the literature on demand-driven media bias. In [21], firms bias their reports in an appeal to consumers' preferences for having their beliefs confirmed. Meanwhile, in [22] bias arises purely in response to reputational concerns, and is thus driven by an aim for long-term profitability. My framework accounts for both the short-term and long-term objectives of a news firm. This sheds light on an intertemporal tradeoff faced by news media: low-quality reporting may benefit a firm in the short run, but can cause damage in the long run. Separately, I note that the kind of deception firms engage in shares common threads with other work. The notion of faking is also studied in [23] in a competition-free setting that encorporates discounting. Furthermore, the endogenous Poisson arrival of inaccurate information, a feature our equilibrium exhibits, also arises in [24], and takes the form of "spamming" by recommender systems.

Outline The remainder of the paper is organized as follows. Section 2 presents the model. Section 3 is dedicated to characterizing its equilibrium, first considering the monopoly benchmark and then encorporating competition. In Section 4, I present the core economic implications of this equilibrium, which pertain to the effects of competition and equilibrium dynamics. In Section 5, I present comparative statics. Section 6 considers an extension of the model in which firms have heterogenous learning abilities. Finally, Section 7 concludes. All formal proofs are relegated to the Appendix.

1.2 A model of breaking news

There are $N \geq 1$ firms, indexed by i, and one consumer. Time, which is continuous and has an infinite horizon, is denote by $t \in [0, \infty)$. There is a time-invariant state $\theta \in \{0, 1\}$, which denotes whether a particular story is true ($\theta = 1$) or false ($\theta = 0$). All players are endowed with a common prior $p_0 \equiv Pr(\theta = 1) \in (0, 1)$.

Each firm privately learns about the state by means of a one-sided Poisson signal: if $\theta = 1$, a signal revealing that $\theta = 1$ arrives to each firm at a Poisson rate $\lambda > 0$. To formalize this learning process, let $s_i \in [0, \infty]$ denote the time at which such a conclusive signal arrives to firm i, with $s_i = \infty$ denoting that a signal never arrives. I assume that $s_i \sim (1 - e^{-\lambda s_i})$ if $\theta = 1$, and $s_i = \infty$ if $\theta = 0$. I further assume that conditional on $\theta = 1$, s_i is i.i.d. across firms. Under this learning process, firms validate a story by receiving breakthroughs. I assume this learning process because it reasonably approximates the learning that takes place in a breaking news setting. One can imagine that given the short lifecycle of a breaking news story, firms do not seek piecemeal evidence but rather pursue reliable sources who can confirm the story. For instance, in the case a terrorist attack, this would entail reaching out to contacts in the police department.

In addition to learning about the story, firms also report about it. Each firm has a single opportunity to make a report over the course of the game. I assume that firm does not choose what to report, but instead whether and when to do so. As the payoff function will soon illustrate, the content of this report can be interpreted as an assertion that the story is true, i.e., that $\theta = 1$. A report history H is a partially ordered set of pairs (i, t_i), pairing each firm i who has reported with a report time t_i, with elements ordered according to the order in which the respective reports were made.[4] Report histories are public information: at every time t, all players observe the current report history. Thus, firms not only learn about θ via their private signal, but also observationally by means of their rival firms' reports.

A firm who never reports earns a payoff of 0. A firm who does report earns

$$k_n \alpha - \beta \mathbb{I}(\theta = 0). \tag{1.1}$$

The first term of (1.1), $k_n \alpha$, denotes the immediate market share (i.e., viewership or readership) that the firm enjoys from reporting a story. k_n captures the role of the firm's order n in its payoff, while α denotes the credibility of the firm's report. Formally, an index n denotes that the firm was the nth to report. The k_n are constants, where $k_1 \geq k_2 \geq \ldots \geq k_N \geq 0$. This assumption accounts for competition. All else equal, firms who report early compared to their competitors enjoy greater market share. The firm's payoff is also increasing in the credibility of its report, α. A report's credibility is the consumer's belief, at the time the report was made, that the firm has received evidence that $\theta = 1$. Formally, this is the belief that $s_i \in [0, t]$, where t is the time of the firm's report. While the k_n are exogenous, α is endogenous. In assuming a product form for the market share, I take the stance that a report is profitable only insofar that consumers believe it was informed. This captures the notion that consumers value accuracy in journalism, and thus only consume news to the extent that they find it credible.

The second term of (1.1), $-\beta \mathbb{I}(\theta = 0)$, captures the ex-post penalty of error: a firm who reports when $\theta = 0$ incurs a penalty, given by a constant $\beta > 0$. This penalty captures the reputational harm a firm suffers from making a report that is later uncovered to be false.

[4] Formally, elements are ordered according to relation $\succsim$, where $(i, t_i) \succ (j, t_j)$ if $t_i > t_j$ or $t_i = t_j$ but j reported first, and $(i, t_i) \sim (j, t_j)$ indicates that the reports were made simultaneously.

1.2.1 Equilibrium

A *Markov strategy* F is a set of distributions $F_{p,n}$ over future report times for each belief $p \equiv Pr(\theta = 1)$ and order $n \in \{1, ..., N\}$ of the next firm to report.[5] Formally, let t denote the span of time the firm waits before reporting conditional on not receiving a conclusive signal. Then, t is distributed according to $F_{p,n} \in \Delta[0, \infty]$, where $t = \infty$ denotes a lack of report altogether.[6] I restrict attention to symmetric equilibria, and thus will omit the firm's index from the $F_{p,n}$ in much of the analysis below.

I place some restrictions on F. First, I assume that for all (p, n), $F_{p,n}$ must be piecewise twice differentiable and right-differentiable everywhere on $[0, \infty)$. This restriction grants analytical convenience and ensures that all equilibrium objects are well-defined.[7]

Second, I impose a selection criterion (SC): a firm immediately reports once it has learned the story is true. This criterion is stated formally as follows:

Definition 1. F satisfies (SC) if

$$F_{1,n}(t) = 1 \text{ for all } t \geq 0, n \in \{1, ..., N\}.$$

(SC) imposes that firms do not abstain from reporting a story they know to be true. It serves the purpose of ruling out unintuitive equilibria with periods of silence, which can only be supported by off-path beliefs that reports made during these gaps entail little or no credibility. An implication of this assumption is that fixing any starting belief p, all players who have not yet reported will share

[5]Formally, $n = |H| + 1$, where H denotes the current history. I assume that if m firms report at the same history H, one firm will be assigned order n, another $n + 1$, etc., with their identities randomly determined.

[6]By defining strategies in this way, I assume that firms react can instantly to a competitor's report. To illustrate this point, suppose $F_{p,2}(t) = 1$ for all t and p. Then if some firm makes the first report at t, all other firms will also report at t.

[7]Note that F satisfies the above restrictions if and only if there exist two functions on p, q_n and b_n, where for all (p, n) and $t \geq 0$,

$$F_{p,n}(t) = \sum_{s \leq t | q_n(p(s)) > 0} q_n(p(s)) + \int_0^t b_n(p(s)) ds$$

such that b_n are piecewise differentiable and $q_n(p) = 0$ at all but a countable number of p. Namely, q_n denotes the *point mass* of reports, while b_n denotes the right *hazard rate* of reports.

the same *common belief* about the state after t time has passed. I denote this common belief by $p(t)$.

While defining strategies in this way, i.e. with a separate distribution for each (p, n), is convenient, it introduces redudancy. Thus, I must impose a consistency condition to ensure that the $F_{p,n}$ are consistent with each other whenever on-path.[8] This condition stipulates that $F_{p,n}$ and $F_{p(t),n}$ are related via the following formula:

$$F_{p(t),n}(s) = \frac{F_{p,n}(s+t) - F_{p,n}(t_-)}{1 - F_{p,n}(t_-)} \text{ for all } s \geq 0 \text{ whenever } F_{p,n}(t) < 1, \tag{1.2}$$

where $F_{p,n}(t_-) \equiv \lim_{\tau \uparrow t} F_{p,n}(\tau)$. This formula is an immediate result of Bayes Rule.

Before proceeding, I define two intuitive terms to describe reporting behavior: *faking* and *truth telling*. A report is *fake* if it is made by a firm despite lacking independent confirmation, i.e., a signal $s^i \neq \emptyset$. Meanwhile, a report that is made in response to such a signal is *truthful*. I use these terms to not only describe a firm's report, but also its behavior: a firm is faking if it is sending a fake reports, while it is truth telling if its reports are exclusively truthful. Given the above selection assumption, strategies only differ in the distributions they place over fake reports.

I seek a symmetric perfect Bayesian equilibrium of this game. This is defined as a Markov strategy F paired with beliefs α and p at each history such that F satisfies sequential rationality and both α and p are consistent with Bayes Rule.

The consistency of α with Bayes Rule implies that it must be given by the following formula at all (p, n) on-path: [9]

$$\alpha_n(p) = \begin{cases} \frac{\lambda p}{\lambda p + b_n(p)} & \text{if } F_{p,n}(0) = 0 \\ 0 & \text{if } F_{p,n}(0) > 0 \end{cases} \tag{1.3}$$

where $b_n(p) \equiv F'_{p,n}(0+)$ denotes the right-derivative of $F_{p,n}$ at 0. That is $b_n(p)$ denotes the

[8] This condition is analogous to the closed-loop property specified in [2]. I adopt the term *consistency condition* from [25], who define this condition for a general class of continuous-time games of timing.

[9] Formally, the formula is derived by applying Bayes Rule to a discrete-time approximation of the beliefs that obtain under this game. This derivation is presented in Section A.1.

instantaneous hazard rate of fake reports by a firm. This can be interpreted as the intensity with which a firm fakes at a particular (p, n).

This formula is intuitive. If $F_{p,n}(0) > 0$, there exists a point mass of reports at (p, n). However, because conclusive signals are continuously distributed over time, the probability with which a valid report is made at (p, n) is zero. Thus, the consumer and all competing firms know with certainty that a report made at (p, n) was fake, and thus assigns to it a credibility of zero. Meanwhile if there does not exist a point mass of reports at (p, n), credibility is assessed by comparing the instantaneous arrival rate of truthful reports (λp) to that of fake reports ($b_n(p)$), assigning higher credibility to reports made when the hazard rate of fake reports is comparatively low.

1.3 Equilibrium characterization

1.3.1 Properties of equilibrium

I begin by establishing two necessary conditions on the firm's equilibrium strategy that will guide the equilibrium characterization. Namely, I show that there are *no jumps* and *no gaps* in the distribution of fake reports whenever credibility is less-than-perfect. These two properties arise in other games with continuous strategy spaces, albeit in different forms.[10] In my setting, these properties hold even in the absence of competition. As I will illustrate below, this is because they are driven by the endogeneity of the firm's payoff.

These two properties are stated formally as Lemma 1:

Lemma 1. *In equilibrium, at any (p, n) on-path $F_{p,n}$ is*

(a) continuous at all t whenever $p < 1$;

(b) strictly increasing at any t such that $\alpha_n(p(t)) < 1$.

Let us begin by considering part (a) of Lemma 1, i.e., the "no jumps" property. This states that fake reports are distributed continuously over time whenever a firm is not certain that the story is

[10] In particular, similar properties have been established in war of attrition games ([26]) and all-pay auctions ([27]).

true. I.e., there can never be a point mass in faking when $p < 1$. This property holds even when competition is absent ($N = 1$). To see why, recall that reports that are made whenever there is a point mass in faking yield zero credibility. Meanwhile, faking while also not being certain than the story is true yields a strictly positive expected penalty $\beta(1 - p)$. Thus, a firm's value from faking at such a time is strictly negative. The firm could then profitably deviate by truth telling. Doing so would preclude the firm from making an error, ensuring a weakly positive payoff.

Next, let us turn to part (b) of Lemma 1, the "no gaps" property. This states that whenever the firm is less-than-fully credible, the hazard rate of fake reports must be strictly positive. I.e., firms must mix between faking at all times where $\alpha_n(p(t)) < 1$. Again, this results from the relationship between credibility and faking in equilibrium. A less-than-perfect credibility means that there is a positive hazard rate of fake reports ($b_n(p) > 0$). Notably, Lemma 1(b) has implications for the firm's incentives. It implies that whenever a firm's credibility is less-than-perfect, it must be indifferent between faking immediately and waiting some short increment of time before doing so.[11] As we will illustrate, this indifference condition will be crucial to the equilibrium characterization.

1.3.2 The monopoly benchmark and role of commitment

Before providing a full characterization, I will first consider the special case where there is a single firm, i.e. $N = 1$. This serves two purposes. First, it elucidates the forces at play when competition is absent. In particular, it shows that errors may occur even without competition, and that such errors are driven by a lack of commitment power by the firm. Second, I will use the no-competition case as a benchmark for understanding the marginal impact of competition.

I now equilibrium credibility under a monopoly. As there is only a single firm, I will drop the n index from all functions and parameters.

Claim 1. *Under a monopoly, for all p on-path*

$$\alpha(p) = \min\{\beta/k, 1\}.$$

[11]This indifference condition is formalized as Lemma 8 in the appendix.

Claim 1 establishes two facts about the monopoly equilibrium. First, credibility is constant over time. The static nature of credibility is often implies the firm strategy is dynamic in nature, exhibiting greater truthfulness over time. Second, the monopolist's credibility is weakly increasing in β, and is less-than-perfect whenever β is sufficiently small. Errors occur even without competition, whenever the ex-post penalty from erring is sufficiently small. In the remainder of this section, we will both provide intuition for these properties and discuss their implications for the firm's reporting strategy.

Let us first consider why credibility must be constant. Recall that a firm mixes between faking at all times in which its credibility is less-than-perfect (Lemma 1). Thus, whenever $\alpha_n(p) < 1$, the firm must be indifferent between two different strategies: (1) faking immediately and (2) remaining truthful until some short wait dt, and then faking.[12] By the martingale property of firm's belief p about the state, both of these strategies yield the same expected penalty from error $\beta(1-p)$. Then, in order to ensure that both strategies are optimal, they must yield the firm the same expected prize as well. Thus, credibility must be constant. Note that this reasoning implicitly assumesthat waiting is costless. Indeed, this is true under a monopoly. Not only is waiting not intrisnically costly (i.e., future payoffs are not discounted), a monopolist does does not incur the implicit cost to waiting that preemption entails. As we will illustrate below, such an implicit cost of waiting is precisely why credibility must increase under competition.

The static nature of the monopolist's credibility implies certain dynamics in reporting. Specifically, the hazard rate of faking (b) strictly decreases over time and tends to zero whenever credibility is less-than-perfect. A firm who fakes it will become gradually more truthful over time. This is illustrated by Figure 1.1, which graphs b over time. The decreasing nature of b follows directly from (1.3), and is intuitive. As time passes without the firm making a report, the consumer becomes increasingly skeptical of the story's truth. This declining belief is an artifact of the firm's one-sided Poisson learning process: the absence of a report means that the firm has not received a

[12]Implicitly, this line of reasoning assumes that $\alpha_n(p(t))$ is continuous in t: this ensures that if $\alpha_n(p) < 1$, then $\alpha_n(p(dt)) < 1$ for dt sufficiently small. While I do not discuss this here, I formally establish continuity in Subsection A.2.2 (see Lemma 9).

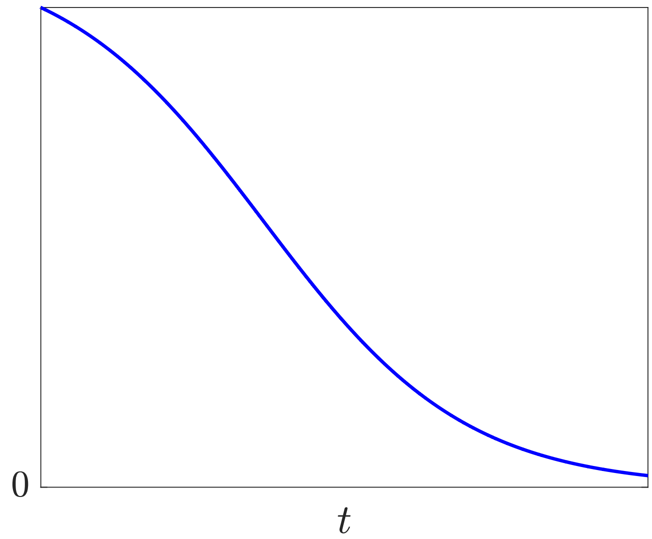

Figure 1.1: Hazard rate of fake reports in monopoly case when $\beta < k_n$.

conclusive signal, an event that is consistent with $\theta = 0$.[13] This means that the consumer believes that truthful reports will become increasingly less probable. To ensure that the firm's credibility remains constant, the hazard rate of fake reports must decline as well, and eventually vanish.

Finally, let us consider why the monopolist errs with positive probability when the penalty of error is sufficiently small. I argue that such errors are driven by a firm's inability to commit to a reporting strategy.

To provide some intuition, let us first consider why truth telling cannot be sustained when $k > \beta$. Recall that a firm that truth tells in equilibrium enjoys full credibility when making a report. Thus it is strictly optimal for the firm to report at some point, even if the story turns out to be false: the immediate payoff of the report, k, exceeds the maximal expected penalty (β) from error. That is, the firm can profitably deviate by faking. The profitability of this deviation is driven by the fact that consumers cannot discern whether a report is fake. Namely, the firm enjoys full credibility even when it were faking. If the consumer could observe whether a report is fake, it would allocate zero credibility to such reports, and the firm would achieve a negative expected payoff from faking. Instead, consumers assess credibility in equilibrium based on their knowledge

[13]Formally, in the monopoly case, $p(t) = \frac{pe^{-\lambda t}}{pe^{-\lambda t}+(1-p)}$.

of the firm's strategy. Even though these beliefs may be consistent with the firm's strategy, the firm can always deviate after its credibility has been assessed. In particular, the firm can in equilibrium deviate by faking more without damaging its credibility.

One may consider how monopolist firm would behave if it could commit to a reporting strategy. That is, suppose that the the firm may its strategy at the start of the game, and was unable to deviate from it once credibility had been assessed.[14] Under commitment, faking is more costly for the firm as it damages its credibility.

One can immediately see that given the ability to commit, the firm would always choose truth telling over its non-commitment strategy (even when $\beta < k$). By committing to truth telling, the firm is guaranteed payoff of k if $\theta = 1$, and 0 if $\theta = 0$. Meanwhile, under the no-commitment equilibrium, the firm will earn strictly less when $\theta = 1$, due to its strictly lower credibility. Meanwhile, it will also earn 0 when $\theta = 0$: though the firm may fake, it will break even from doing so, as its payoff from faking is exactly equal to the penalty of error. One can strengthen this result, showing that truth telling it is the unique commitment solution under a monopoly. That is, given the ability to commit, a monopolist would never commit errors. This result is presented formally in Section A.3. This also illustrates an important point about a firm's incentives: while commitment makes faking more costly, it leaves the firm better off in equilibrium.

1.3.3 Full model characterization

Now, I characterize the equilibrium under the full model (i.e., under an arbitrary n). I establish existence and uniqueness of an equilibrium. To this end, I show that any equilibrium is the solution to a recursive set of boundary value problems. Specifically, whenever the firm is not truthful, its credibility must satisfy an ODE and a boundary condition.

I first derive the set of conditions that determine whether or not a firm is truthful. This will both serve as a stepping stone to a full characterization, and elucidate how competition can deteriorate credibility and exacerbate faking. I state this result formally as Proposition 1.

[14] While we discuss the commitment solution informally here, a formal treatment is presented in Section A.3.

Proposition 1. *In equilibrium, at any (p, n) on-path, $\alpha_n(p) = 1$ if and only if:*

1. $k_n \leq \beta$
2. $p \leq p_n^* \equiv \min\{\frac{k_n - \beta}{\frac{k_n}{N-n+1} - \beta}, 1\}$.

This result provides two conditions, on the model parameters and the common belief, that are necessary and sufficient for truth telling. Recall that the first condition, $k_n \leq \beta$, was both necessary and sufficient for truth telling under a monopoly (Claim 1). However, Proposition 1 asserts that under competition, this condition alone is not enough to ensure truth telling. A second condition is required: the common belief must be sufficently low, lying below some threshold p_n^*. That is, firms must also be sufficiently skeptical about the story's validity.

The necessity of this second condition illustrates an important point: truth telling is harder to sustain under competition. Under a monopoly, this was true as long as the cost of an error (β) outweighed maximal the benefit from reporting (k). However competition also introduces a risk of being preempted. If a firm engages in truth telling, there is a risk that its opponent learns the story is true, and thus reports first. A firm can circumvent this risk by faking.

In the above reasoning, we took for granted that being preempted is costly whenever the firm is being truthful. Let us now explain why this is indeed the case. This is most obvious in the *winner takes all* case, where $k_n = 0$ for all $n > 1$. In this special case of the parameters all firms, with the exception of the first to report, earn a payoff of zero. Here, the costliness of being preempted is obvious: a firm who is preempted can at best earn a payoff of zero. Generally, however, the decreasing nature of the k_n alone does not guarantee that being preempted is costly: improved credibility for succeeding firms could endogenously counteract the decay in the k_n and make being preempted costless, or even valuable. However, one can see that conditional on the firm being truthful, being preempted must be costly in equilibrium: truthfulness guarantees full credibility, leaving no room for improvement for the succeeding firm.

Now, consider the second condition, namely that the firm is only truthful if it is sufficiently pessimistic about the story's truth. While truth telling entails a risk of being preempted, faking entails a different kind of risk. This is the of making an error and incurring penalty β. Both of these risks depend on the belief p about the state. A higher belief p imples both a lower risk of error and a higher risk of being preempted, both of which make faking relatively more appealing to the firm, and thus, make truth telling more difficult to sustain. While it is immediate that a greater p implies a smaller risk of error, that it implies a greater risk of being preempted is less obvious. Note that if the story is true, a firm faces a greater risk of being preempted, because its competitor may learn the story and make a preempting report. Thus, a firm with a higher p will perceive its risk of being preempted. While Proposition 1 pins down the conditions under which the firm is fully credible, it remains to characterize the firm's behavior when truth telling does not hold. To this end, we obtain a key result: the firm is faking, credibility must satisfy a particular ODE and limit condition.

Proposition 2. *In equilibrium, at all (p, n) on-path where $k_n \geq \beta$ or $p > p_n^* \equiv \frac{k_n-\beta}{k_n/n-\beta}$, the following ODE must be satisfied:*

$$\alpha_n'(p) = -\frac{1}{k_n(1-p)\alpha_n(p)}\frac{N-n}{N-n+1}[k_n\alpha_n(p) - V_{\tilde{p},n+1} - \beta(1-\alpha_n(p))(1-p)], \qquad \text{(ODE)}$$

where $\tilde{p} \equiv \alpha_n(p) + (1-\alpha_n(p))p$.

In addition, $\lim_{p\to 0+}\alpha_n(p) = \beta/k_n$ *must hold if $k_n > \beta$ and* $\lim_{p\to p_n^*+}\alpha_n(p) = 1$ *if $k_n \leq \beta$.*

The proof for Proposition 2 relies on the property that whenever a firm is less-than-fully credible, it must be indifferent between faking immediately and after some short wait. To state this formally, let δ_s denote the pure strategy distribution that places full mass on faking after s time has passed. The indifference condition can then be written as follows:

$$V_{p,n}(\delta_0) = V_{p,n}(\delta_{dt}),$$

where $V_{p,n}(\cdot)$ denotes the firm's value from playing a particular strategy at (p, n), and dt is some short increment of time.

Now, taking a Taylor approximation of the firm's value from waiting, $V_{p,n}(\delta_{dt})$, yields the following:

$$V_{p,n}(\delta_{dt}) - V_{p,n}(\delta_0) = [\frac{dp}{dt}(k_n\alpha_n'(p)) - \frac{\lambda p(N-n)}{\alpha_n(p)}(V_{\tilde{p},n} - V_{\tilde{p},n+1})]dt + o(dt^2). \tag{1.4}$$

(1.4) is intuitive. It states that waiting to fake, rather than faking immediately, has two implications for the firm's payoff. The first is that the firm's credibility $\alpha_n(p)$, and thus the payoff enjoyed from reporting, may change. This change in credibility is approximated by $\frac{dp}{dt}(k_n\alpha_n'(p))dt$. In addition, by waiting, the firm risks being preempted. Namely, with probability $\frac{\lambda p(N-n)}{\alpha_n(p)}dt$ the firm is preempted, in which case its expected payoff will decline by $V_{\tilde{p},n} - V_{\tilde{p},n+1}$. I interpret this decrease in value as the firm's cost from being preempted.

Let us now examine both the probability and cost of preemption more closely. As one might expect, the probability of being preempted is increasing in the number of rival firms $(N-n)$ and the expected rate at which these rivals are able to confirm the story (λp). It is also decreasing in equilibrium credibility. This is due to the fact that lower credibility firms are more likely to fake, and thus pose a greater preemptive threat.

As for the firm's cost of being preempted, let us begin by considering the second component of this expression, given by $V_{\tilde{p},n+1}$. This denotes the firm's continuation value in the event that it is preempted. Being preempted not only affects the firm's order but also the the common belief about the state. While the common belief was p prior to the rival firm's report, it increases to $\tilde{p} \equiv \alpha_n(p)+(1-\alpha_n(p))p$ following the report. This is due to observational learning. Specifically, a rival firm's report means one of two things: either the report was triggered by the arrival of a conclusive signal, in which case the story is certainly true and the belief would become p, or it was not, in which case the report provides no new information and the belief remains p. The common belief following this report is a weighted sum of these two conditional beliefs. In particular, the weight given to the rival firm's report being informed by a conclusive signal is precisely its credibility at the time of the report, $\alpha_n(p)$. It is precisely this new common belief that determines the firm's continuation value if it is preempted.

The cost of being preempted measures the impact of being preempted on the firm's continuation value. I.e., it measures how much the firm's continuation value from being preempted differs from that in which it is not. Importantly, both continuation values are assessed at the common belief after being preempted, $\tilde{p}$. In this sense, we can view the cost of being preempted as the firm's *ex-post regret* from being preempted.

In order for the indifference condition to be satisfied, the term of (ODE) that is linear dt must equal zero. This equality yields (ODE).

In addition to establishing (ODE), Proposition 2 imposes a limit condition on the firm's credibility. This limit condition always applies at the boundary of the region of beliefs in which the firm is faking. First consider the case where $k_n \leq \beta$. Recall from Proposition 1 that in this case, $\alpha_n(p) = 1$ whenever $p \leq p_n^*$. We must then have that $\alpha_n(p)$ limits to 1 as the belief approaches p_n^*. If it did not, then as the belief approached p_n^*, the firm could profitably deviate by not faking immediately, and rather waiting until p_n^* is reached to do so.

Let us next consider the case where $k_n > \beta$. In this case, the firm never truth tells in equilibrium, and thus the indifference condition must always be satisifed. As the common belief p approaches zero, a firm who fakes does so being increasingly certain that its report is erroneous, and will incur penalty β. Thus, the firm's payoff from faking limits to the following:

$$\lim_{p \to 0+} V_{p,n}(\delta_0) = k_n \lim_{p \to 0+} \alpha_n(p) - \beta.$$

Separately, even though the firm sometimes fakes, it must also never fake, i.e., play strategy δ_∞, with positive probability. This guarantees that the hazard rate of fake reporting remains low enough to ensure that credibility remains sufficiently high, and thus that the firm will indeed find it optimal to fake.[15] As $p \to 0+$, the value of truth telling tends to zero, as it becomes increasingly likely that the firm never reports. The limit condition in this case, $\lim_{p \to 0+} \alpha_n(p) = \beta / k_n$, is precisely what is needed to ensure indifference between faking and truth telling.

[15] I formalize this result in the Appendix as Lemma 7.

To take stock, Proposition 1 and Proposition 2 provide two necessary conditions on equilibrium credibility. They establish a region under which truth telling must occur in equilibrium (Proposition 1), and show that otherwise, credibility must satisfy a recursive boundary value problem (Proposition 2). I further show that these two conditions are not only necessary, but also sufficient, for an equilibrium.[16] There is some intuiton behind this sufficiency result. First, consider the region where truthfulness is necessary. Let us suppose by contradiction that is not an equilibrium strategy. Then, even though the firm's credibility would be perfect on this region, it could profitably deviate by choosing a strategy that involves faking. However, Proposition 1 establishes that such a strategy cannot be played in equilibrium. That is, the firm could profitably deviate by truth telling even when the the firm's opponents are also faking (i.e., the risk of being preempted is higher) and credibility is less-than-perfect (i.e., the benefit of reporting is lower). Such a strategy would thus not be more profitable than truth telling when the firm's opponents are not faking, and credibility is perfect. Meanwhile on the region where $\alpha_n(p) < 1$, the firm's strategy involves faking. Faking is optimal on this region because (ODE) guarantees it. In particular, it ensures the firm's indifference condition holds, i.e., that faking at any such time is optimal.

I thus establish that the equilibrium is fully characterized by the solution to a recursive boundary value problem. While I do not have a closed-form solution to this problem on the region where $\alpha_n(p) < 1$, I use the Picard Theorem to establish both existence and uniqueness. This result is stated as Theorem 1.

Theorem 1. *There is a unique equilibrium (where uniqueness applies at (p, n) on-path).*

1.4 Economic Implications

In this section, I consider the economic implications for (1) the dynamics of firm behavior and (2) the impact of competition on both credibility and the prevalence of errors.

[16]This result is stated formally in the Appendix as Lemma 10.

1.4.1 Equilibrium dynamics

Let us consider how a firm's credibility and reporting behavior evolves over the course of time. These results not only illustrates when firms are most prone to erring, but also offers us a understanding of the endogenous nature of the firm's incentives.

Dynamics take two separate forms in equilibrium: *continuous* changes and *discrete* changes. As I will show, continuous changes occur in the absence of any new reports, while discrete changes are triggered by a new report. Formally, let us denote a *subgame* by a pair (p, n), where p denotes a starting belief and n the order of the next firm to report. I claim that fixing a subgame, i.e., assuming that no new reports are made, the firm's credibility will change continuously over time. In particular it will gradually improve whenever preemptive concerns are present.

Proposition 3. *For all* (p, n) *on-path,* $\alpha_n(p(t))$ *is weakly increasing in* t*. Furthermore,*

1. *If* $\beta > k_N$*, then* $\frac{d}{dt}\alpha_n(p(t)) > 0$ *whenever* $\alpha_n(p(t)) < 1$.

2. *If* $\beta \leq k_N$*, then* $\alpha_n(p(t))$ *is constant in* t*.*

While $\alpha_n(p(t))$ must be constant under a monopoly, competition can introduce dynamics. Proposition 3 asserts that as long as $\alpha_n(p(t))$ has not reached its upper bound of 1, it strictly increases precisely when being preempted is costly.

Formally, this follows from (ODE). It is especially clear when we write (ODE) in the following form:

$$\frac{d}{dt}\alpha_n(p(t)) = \frac{\lambda p(N-n)}{\alpha_n(p(t))k_n}[V_{\tilde{p},n} - V_{\tilde{p},n+1}]$$

We can see that $\alpha_n(p(t))$ must strictly increase over time whenever the cost of preemption is strictly positive.

There is a clear intuition for this result. Recall again that whenever the firm is less-than-fully credible, it must be indifferent between faking immediately and waiting some period of time before doing so. However, if credibility remained constant, this indifference would fail whenever preemption is costly: the firm would obtain the same expected payoff from reporting in both cases, but by reporting immediately would avert being preempted. To ensure that indifference is preserved, the firm must somehow be compensated for waiting. This can only be achieved by means of strictly increasing credibility. While waiting presents a cost to being preempted, a strictly increasing $\alpha_n(p(t))$ ensures that the firm's report will be rewarded more in the event that it is not preempted. Thus, the increasing nature of $\alpha_n(p(t))$ is crucial to balancing the firm's equilibrium incentives: it endogenously mitigates the haste-inducing effects of preemptive risk.

Let us now consider the implications of this result. It asserts that news reports that are made with greater delay for research are generally more trustworthy in the eyes of consumers. That is, all else equal, consumers will have greater trust in a firm's journalistic standards when a report is not made quickly. In this sense, this result conforms with consumers' stated concerns about hasty reporting. This model provides a justification for such concerns that are grounded in the firm's incentives. Furthermore, by the same reasoning presented in the discussion of the monopoly case, the increasing nature of credibility within a subgame implies that firms become gradually more truthful over the course of the game. That is, $b_n(p(t))$ is strictly decreases over time whenever the firm is not fully credible.

Finally, while Proposition 3 asserts that $\alpha_n(p(t))$ must be strictly increasing when there is a cost to being preempted, this is not always the case. Specifically, when $k_N \geq \beta$, being preempted is costless in equilibrium (i.e., $V_{\tilde{p},n} - V_{\tilde{p},n+1} = 0$), and thus $\alpha_n(p(t))$ is constant. In other words, preemptive concerns endogenously disappear whenever the ex-post cost of error, β, is sufficiently small. Formally, the credibility function will adjust in such a way that ensures $k_n\alpha_n(p) = k_{n+1}\alpha_{n+1}(p)$ for all p. This highlights an notable feature of our model: competition alone does not imply preemptive concerns. Even if competition is present, credibility can change in such a way that makes preemption costless.

Let us now understand why α_n must be constant in equilibrium when $k_N \geq \beta$. We formally prove this by means of a backwards induction argument. To this end, we begin with the last (Nth) firm to report. It follows from the monopoly characterization that $k_N\alpha_N(p) = \beta$ at all p. Next, let us consider the second-to-last ($N-1$th) firm to report. One can show that in this case as well, $k_{N-1}\alpha_{N-1}(p) = \beta$. To see why, let us first suppose by contradiction that at some p, $k_{N-1}\alpha_{N-1}(p) > \beta$. In this case, the firm would find it strictly optimal to fake at some point, as it would not only guarantee the firm a positive payoff regardless of the true state, but also preclude the firm from being preempted. This would amount to a failure of the firm's indifference condition, and thus the equilibrium. If on the other hand $k_{N-1}\alpha_{N-1}(p) < \beta$ at some p, the firm would want to deviate by being truthful. This is because preemption would *benefit* the firm: by being preempted, the firm would be guaranteed the continuation prize of β, exceeding the prize it would earn from reporting immediately and avoiding preemption. One can proceed inductively in a similar fashion, establishing that for all n, $k_n\alpha_n(p) = \beta$.

Let us now consider discrete changes in the firm's credibility and faking. While credibility changes continuously within a subgame, a rival report will cause the firm's subgame to change. That is, a report made at (p, n) will cause the order of the next reporter to increase to $n+1$ and the common belief to increase to $\tilde{p}$. This will in turn result in discrete jumps in the firm's credibility and hazard rate of faking (b). These discrete jumps are apparent in Figure 1.2, which plots a simulation of α and b over the course of the game. As these graphs illustrate, jumps in both α and b are not monotonic. An opponent report may trigger either a boost or decline in α and b. This can be seen in Figure 1.2, while the first four reports cause credibility to decrease and faking to increase, the fifth report causes credibility to decrease and faking to increase. These first four reports illustrate a *copycat effect*, in which one firm's report causes an immediate a surge in the rate at which others fake.

Let us now consider what is responsible for this copycat effect. First note that the discrete change in credibility that happens when a firm makes the nth report under common belief p is given by the following:

$$\alpha_{n+1}(\tilde{p}) - \alpha_n(p),$$

where again $\tilde{p} > p$ denotes the common belief in the immediate aftermath of the report. This expression shows that a report by one firm affects credibility by imposing two different changes to the environment. First, it impacts the order of the next firm, i.e., by ensuring that the next firm to

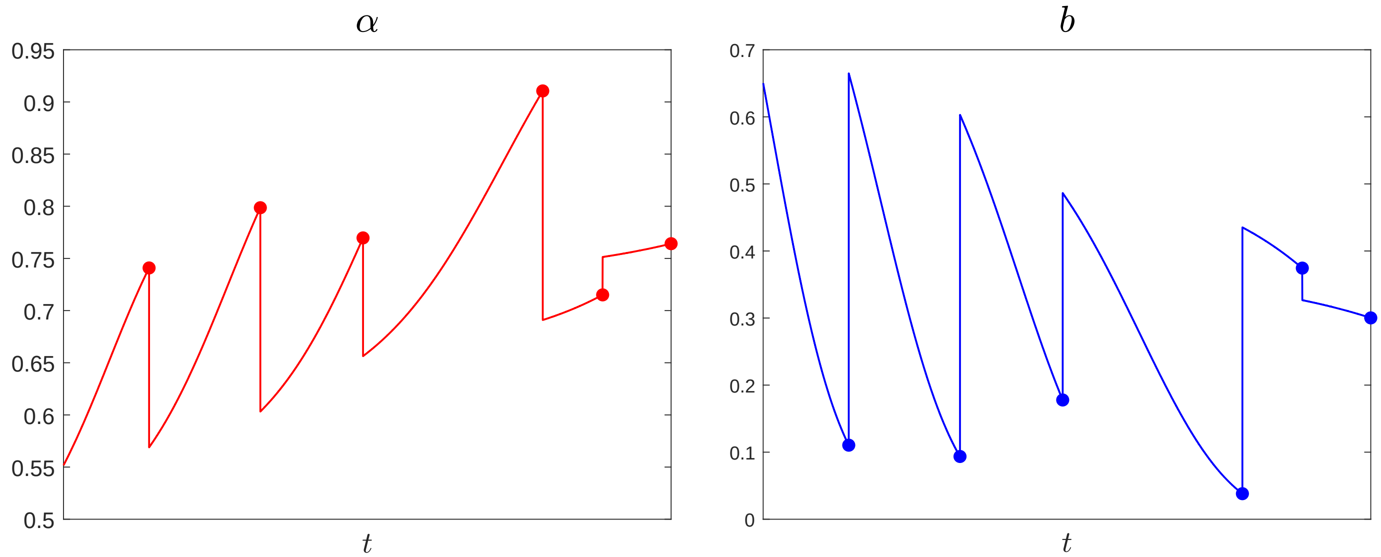

Figure 1.2: Simulations of crediblity and the hazard rate of fake reports, respectively, over the course of the game. Discrete jumps in both graphs signify that a firm has made a report.

report will be the $n+1$th firm to report, rather than the nth. Second, it causes a discrete upwards jump in the common belief: firms will learn observationally from the report of their opponent, and thus become more confident that the story is true. The following decomposition isolates the respective impacts of these two changes:

$$\alpha_{n+1}(\tilde{p}) - \alpha_n(p) = \underbrace{[\alpha_{n+1}(\tilde{p}) - \alpha_{n+1}(p)]}_{\text{change in belief}} + \underbrace{[\alpha_{n+1}(p) - \alpha_n(p)]}_{\text{change in order}}$$

The effects of a change in order alone, $\alpha_{n+1}(p) - \alpha_n(p)$, can have an ambiguous impact on firms' credibility. For instance, a large gap between k_n and k_{n+1} may be consistent with an improvement in credibility, whereas a smaller gap may be consistent with a deterioration in credibility.

However, observational learning will always cause a deterioration in credibility. Formally, $\alpha_{n+1}(\tilde{p}) - \alpha_{n+1}(p)$ will always be negative in equilibrium, and strictly so whenever premeptive concerns are present (i.e., whenever $k_N < \beta$). The negative correlation between credibility and the firm's belief that the story is true is also apparent in Proposition 3: later reports are associated with lower common beliefs about the story being true, and also with higher credibility. There is also a clear intuition to this: the more pessimistic the firm is about the story's validity, the higher its expected penalty from faking will be. This in turn yields the firm less willing to fake, and thus more credible in equilibrium. We thus see that the downwards jumps in credibility are caused, at least in part, by observational learning.

1.4.2 Effects of Competition

In this section, I consider the impact of competition on both credibility and faking in equilibrium. I assess the impact of competition by comparing the equilibrium under competition ($n \geq 2$) to that under the monopoly benchmark.

In order to isolate the effects of competition, I assume that the total ability of the market to learn is constant across these two cases. In particular, I assume that if each firm has ability λ under competition, then the firm has ability $n\lambda$ under the monopoly benchmark. In making this normalization, one ensures that our comparison accounts for only the impact of competition per se and does not confound this with the effects of an increased aggregate ability to learn that firm entry may entail. I do however consider the effects of market entry in the comparative statics section below, in which we do not normalize the total ability to learn.

These findings are shown in Figure 1.3, which depicts both credibility and the hazard rate of faking in the market ($nb_n(p(t))$) within a subgame, i.e., fixing a p and an n. The top and bottom row show the case where $\beta \in (k_N, k_n)$ and $\beta > k_n$, respectively. In both cases, we see that *competition causes a deterioration in credibility and an increase in faking*. The effect of competition in this case is driven by the cost of preemption that it induces. Firms are more inclined to fake, and thus less credible because the cost of preemption makes truth telling more costly. When $\beta \in (k_N, k_n)$, faking occurs even under a monopoly, but moreso under competition. That being said, the effects of competition dissipate over time, as the competition level of credibility limits to the monopoly level as time passes. Meanwhile, in the case where $\beta > k_n$, a monopolist firm will *never fake*, faking does temporarily occur under competition. Again, the effects of competition are greatest early on with firms faking gradually less as time passes.

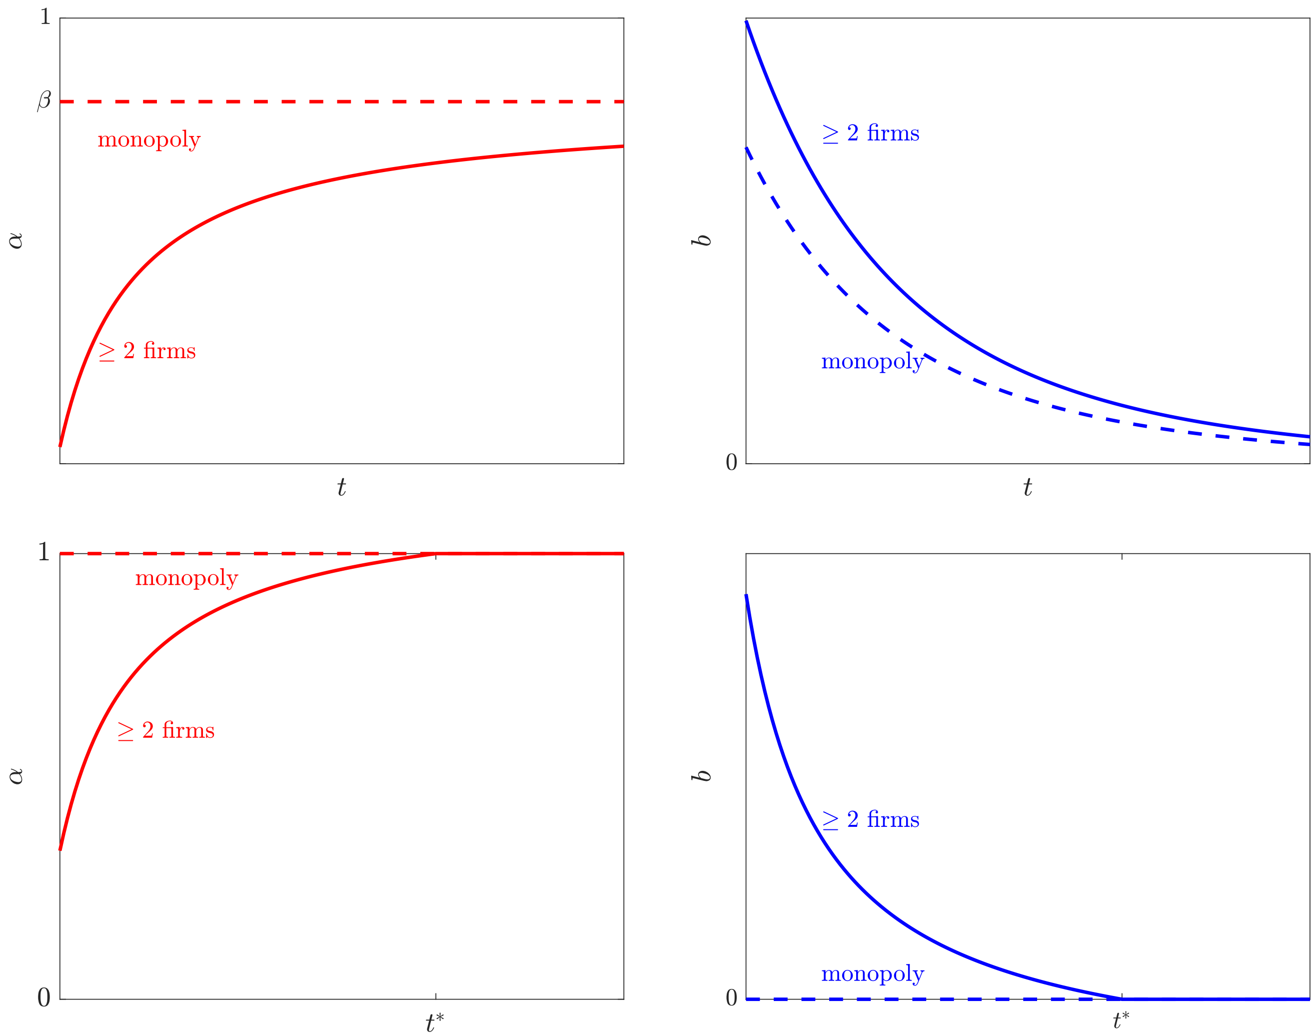

Figure 1.3: Credibility $\alpha_n(p(t))$ (left) and the hazard rate of faking in the market $nb_n(p(t))$ (right) under competition and a monopoly. Top row depicts case where $k_n > \beta$, while bottom row depicts case where $\beta \in (k_N, k_n)$.

1.5 Comparative Statics

In this section, I consider how the equilibrium changes with the parameters of the model. This will shed light on how various features of the news market can either exacerbate or curb erroneous reporting. These findings are stated as Proposition 4.

Proposition 4. *In any equilibrium, for any n,* $\alpha_n(p(t))$ *is*

(a) weakly increasing in β*, and strictly so whenever* $\alpha_n(p(t)) < 1$.

(b) weakly increasing in λ*, and strictly so for* $t > 0$ *whenever* $\alpha_n(p(t)) < 1$ *and* $k_N < \beta$.

(c) weakly decreasing in N, and strictly so whenever $\alpha_n(p(t)) < 1$*, when* $t \in [0, \bar{t}]$ *for some* $\bar{t} > 0$.

Part (a) states that no matter when a firm reports, it will be more credible under high β. This result is intuitive: a higher ex-post cost of error means firms are less likely to fake, and thus more credible. This result is a consequence of the firm's equilbrium incentives: a higher β makes faking more costly. This will either induce the firm to resort to truth telling instead, or require that it is compensated for this coster faking with greater credibility.

Now, let us consider the comparative static on λ. This result is also intuitive: it states that credibility is higher whenever firms have a greater ability to learn. Let us now understand what is driving this result. We first note that at any belief p the firm may hold, a change in λ will have *no effect* on $\alpha_n(p)$ in equilibrium. This is due to the fact that λ does not enter the boundary value problem which dictates the firm's credibility, and thus changes in λ have no effect $\alpha_n(p)$. However, changes in λ will have an effect on the time path of the common belief $p(t)$. Under a higher λ, firms learn about the state more quickly, and thus $p(t)$, the belief that $\theta = 1$ conditional on no reports, will decay faster. That is, firms will be more pessimistic about the story's validity at any time $t > 0$ when λ is higher. This greater pessimism about the story translates to a higher expected cost of erring, which thus makes faking more costly. As was true of the comparative static on β, this increased cost of faking must be counterbalanced by a higher credibility $\alpha_1(p(t))$ at every time

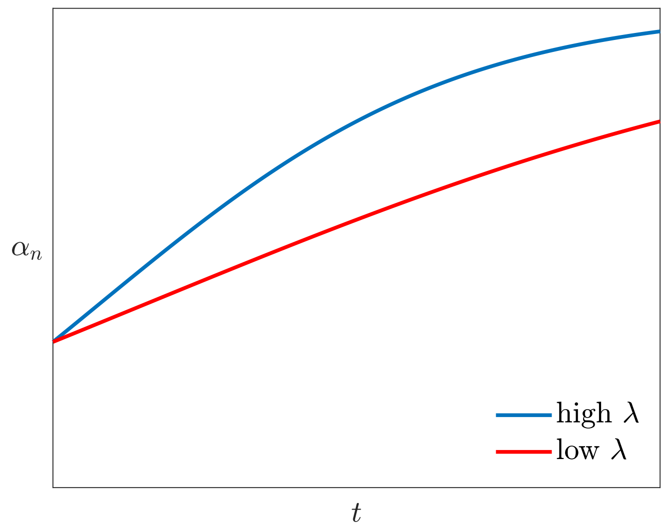

Figure 1.4: A simulation of $\alpha_n(p(t))$ when $\lambda = 1$ (blue line) and $\lambda = 0.5$ (red line). For the remaining parameter values, the following specifications were made: $\beta = 0.5$, $p_0 = 0.7$, $N = 8$, $k_n = 0.7^{(N-n)}$.

$t > 0$. This comparative static is illustrated by Figure 1.4, which shows simulations of the firm's credibility function under both high and low values of λ.

Let us finally consider the comparative static on the total number of firms, N. While it pertains to the level of competition, this exercise is notably distinct from our analysis in the previous section. Therein, we studied the overall impact of competition on equilibrium outcomes. This was done by comparing the case where competition is present ($N > 1$) to the monopoly case ($N = 1$) while holding constant the total learning ability of the market, $N\lambda$. With this comparative static, we are instead considering the *marginal* impact of an additional firm entering the market. In particular, we do not hold fixed the total learning ability of the market. Rather, I assume that this additional firm adds to the total learning ability of the market. In doing so, one can study the effect of *proliferation* in the news industry.

Proposition 4 states that adding a firm to the market will guarantee a deterioration in credibility, but only for a limited amount of time. In fact, the addition of a firm may result in an improvement in credibility during later periods. This phenomenon is captured by Figure 1.5. This figure plots simulations of α_n under $N = 5$ and $N = 6$, respectively, holding all other parameters fixed. While

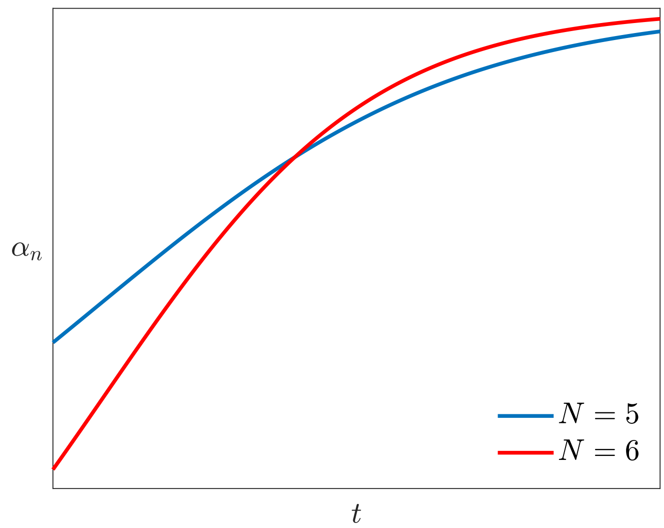

Figure 1.5: A simulation of $\alpha_n(p(t))$ when $N = 5$ (blue line) and $N = 6$ (red line). For the remaining parameter values, the following specifications were made: $\beta = 0.5$, $p_0 = 0.7$, $\lambda = 1$, $k_n = 0.7^{(N-n)}$.

the addition of a firm lowers credibility in early periods, it improves credibility in later periods.

To understand this result, note that an additional firm will effect two separate changes to the market. First, each firm faces greater competition, and thus a greater risk of being preempted. This change is precisely what was captured in our earlier exercise regarding the effects of competition. As illustrated by Figure 1.3, this change will cause a deterioration in credibility. However, an additional firm also increases the market's total ability to learn. This change is captured by our comparative static on λ, which shows that an increase in learning ability will cause an improvement in credibility. Thus, the effect of an additional firm can be understood as the combination of two countervailing forces: higher competition and a higher ability to learn within the market.

To understand why the credibility-diminishing effect of higher competition must dominate in early periods, we must compare the relative magnitudes of these the two countervailing forces. Figure 1.4 illustrates that while credibility is pointwise higher at every $t > 0$ under high λ, this difference is negligible in early periods. This is due to the fact that firms learn gradually over time, and thus it takes time for differences in learning ability to substantially impact firms' beliefs. Meanwhile, as illustrated by Figure 1.3, an increase in competition will have a non-negligible impact on credibility even when $t = 0$. For this reason, the impact of higher competition must dominate in early periods, resulting in a net reduction in credibility. However, as time passes and the effect of faster learning grows, a reversal may take place, i.e., there may be a net improvement in credibility. Such a scenario is precisely what is depicted by Figure 1.5.

1.6 Extension: Heterogeneous Ability

In this section, I consider an extension in which firms are heterogeneous in their abilities to learn. Doing so will shed light on how a firm's credibility correlates with it's ability in equilibrium.

Formally, this extended model is identical to the model above except for three changes. First, rather than assuming that each firm is endowed with the same ability λ, I assume that each firm i is endowed with an individual-specific ability λ^i. As with all other parameters, I assume that these individual-specific abilities are common knowledge. Second, to simplify our analysis for this exercise, I will restrict attention to a winner-takes-all setting: i.e., I assume that $k_n = 0$ for all $n > 1$. Finally, I relax our assumption that the equilibrium is symmetric. Thus, different firms (and in particular, firms with different abilities) may play different strategies in equilibrium and are thus a firm's credibility is individual-specific. Accordingly, I let α^i denote the credibility of firm i.

I obtain an intuitive result: a firm's ability correlates positively with its credibility in equilibrium. This is stated formally as Proposition 5.

Proposition 5. *For all (i, j) such that $\lambda^i < \lambda^j$, $\alpha_1^i(p(t)) \leq \alpha_1^j(p(t))$. Furthermore, this inequality is strict whenever $\alpha_1^i(p(t)) < 1$.*

Proposition 5 states that regardless of when a report is made, a firm with higher ability will be more credible.[17] Furthermore, a high ability firm will be strictly more credible than a low ability firm whenever firms are not fully truthful.

Let us now consider why this correlation arises. First, note that high ability firms are able to confirm a story more quickly and thus, all else equal, pose a greater preemptive threat in equi-

[17] This claim restricts attention to the first firm to report, because by the winner-takes-all assumption, all following senders will never fake, i.e., $\alpha_n^i(p) = 1$ whenever $n > 1$.

librium. This in turn implies that in comparison to a high-ability firm, a low-ability firm faces a greater preemptive threat. Thus, the low-ability firm finds immediate faking more advantageous. In light of this, the firms' credibilities must adjust in such a way to preserve their respective indifference conditions. This is achieved endogenously by means of a lower credibility for the low-ability firm, which ensures that it has less to gain from faking immediately.

1.7 Conclusion

In this paper, I presented a dynamic model of breaking news to understand the nature of reporting errors. I sought to explain how strategic forces that could induce firms to err. In this setting, errors were driven by two qualities of the breaking news environment: a firm's lack of commitment power as well as competition. I find that competition induces firms to err through two separate channels: preemptive motives and observational learning. While preemptive motives can give rise to errors by encouraging firms to report hastily, observational learning can cause an existing error to propagate through the market.

The second key objective was to understand the dynamics of reporting errors. In equilibrium, these dynamics take two forms. First, firms become gradually more truthful over time as long as no new reports are made. Furthermore, a firm's credibility gradually increases whenever preemptive motives are at play. Importantly, this improvement in credibility incentivizes firms to take their time, and thus counteracts the haste-inducing effects of preemption. Dynamics also take the form of discrete changes in the firm's behavior and credibility which are triggered by a rival report. In particular, I document a copycat effect, where a report by one firm can induce a surge in faking by other firms in the market.

While I consider breaking news specifically, this model provides broader insight into how preemptive concerns can affect the quality of information provided by experts. To understand how preemption impacts information provision more broadly is a topic that warrants further investigation.

Chapter 2: Reputation and Misreporting in News Media

2.1 Introduction

The recent addition of social media and non-traditional outlets to the news media lanscape has brought with it concerns about misinformation. However, factual errors have often been a characteristic of traditional media as well. Many such errors by traditional media have regarded some of the most crucial news stories in the United States: the 2000 presidential election [1], the 2013 Boston bombings [2], the Sandy Hook massacre [3], the $9/11$ attacks, the John F. Kennedy assassination [4], etc. The ubiquity of such errors is reflected in the beliefs of consumers: in a Pew Research Center survey from 2009, only 29% of respondents said that news organizations often "get the facts straight", with 63% of respondents expressing a belief that news stories are "often inaccurate"[5] What is responsible for these errors? Could self-serving motives be at play?

In this paper, we consider how *reputational concerns* impact news firms' reporting behavior, and in particular the incidence of misreports. This approach is driven by the observation that news firms' vitality relies heavily on their reputation. A news firm with a reputation for both skilled reporting and journalistic integrity is able to attract and retain consumers who value these attributes. While reputation may matter in a variety of industries, it is particularly salient in the news industry, where given the frequency of exchange between firms and consumers, sustaining these interactions is critical to the firm's welfare.

To this end, we present a model of a reputationally concerned news firm (sender), who dynamically learns about an unknown state, and reports to a consumer. Firms wish to maximize their

[1] Howard Kurtz. *Washington Post,* December 22, 2000.
[2] David Carr. *The New York Times,* April 21, 2013.
[3] Paul Farhi. *Washington Post,* December 18, 2012
[4] Rebecca Greenfield. *The Atlantic,* September 16, 2013
[5] "Press Accuracy Rating Hits Two Decade Low." Pew Research Center, Washington, D.C. (September 13, 2009) https://www.pewresearch.org/politics/2009/09/13/about-the-survey-373/.

reputation for being *good*, which entails both a high ability for learning and a sense of integrity (i.e., honest reporting). We further incorporate a key feature of the news environment: senders decide not only what to report, but when to report.

Our analysis gives rise to three key findings. First, in equilibrium, senders "fake" with positive probability, which entails reporting despite being completely uninformed about the state. This behavior is responsible for a higher incidence of misreporting than that which would prevail if the sender were truthful, i.e., only reporting when she is informed. Importantly, the equilibrium is such that the uninformed sender is indifferent between faking and truth-telling. Notably, we find that faking, and the resulting misreporting, is ceaseless in equilibrium, i.e, it occurs with positive probability at any given time.

Second, we find that in equilibrium, a sender's reputation is rewarded for two separate characteristics of her report: speed and accuracy. That is, while the sender is intrinsically interested solely in maximizing her reputation, she behaves *as if* she wishes to maximize some combination of speed and accuracy. Furthermore, for the sender who is ignorant about the state, these two objectives are at odds, and thus she faces a *speed accuracy tradeoff*: while faking will prevent her from incurring the reputational deterioration that comes with further delaying a report, truth telling saves her from incurring the reputational harm from making a mistake. In equilibrium, the desire for speed and accuracy precisely counterbalance each other in a way that preserves the uninformed sender's aforementioned indifference between faking and truth telling.

This result is compelling for a number of reasons. With regards to our application, it shows that reputational concerns alone may be responsible for the all-familiar speed-accuracy tradeoff faced by newsrooms. Rather than explicitly modeling this tradeoff, we have provided a microfoundation for it, by showing that it arises endogenously when news media are reputation-concerned. Beyond our application, this result demonstrates speed and accuracy, which are canonically assumed to be of intrinsic value to an individual decision maker, may be of *signalling* value in a delegated learning setting. Finally, because the importance of speed and accuracy is endogenous, this allows us an opportunity to understand *how* they dynamically impact the sender's reputation. With regards to accuracy, we find that while making an accurate report causes an improvement in the sender's reputation and making an inaccurate report causes a decline, the magnitudes of these changes are not equal: the reputational harm from inaccuracy strictly exceeds the reputational gain from accuracy. The interpretation of this result in a news media setting is intuitive: while accurate reporting may cause a modest improvement in a news firm's reputatability, the reputational harm from misreporting is much more consequential.

Thirdly, we consider dynamics in the sender's reporting behavior, specifically how her propensity for misreporting changes over time. We find that if the sender is of sufficently low ability, she becomes strictly *more truthful* as time passes, and thus becomes less likely to misreport. This is despite the fact that speed continues to benefit her reputationally over the course of the game. With regards to our application, this result implies that a news firm is more hasty with its reporting, and thus more likely to misreport, when her research process is still in its early stages, but becomes more scrupulous as time passes.

This paper exists at the intersection of two separate literatures: reputational games and games of timing, and specifically those in which decision timing is of signalling value.

With regards to reputational games, there is a broad literature in static settings. [28] consider reputational cheap talk generally, where senders wish to maximize a receiver's belief that they have a high ability to learn. Their key result is that in such games, it is generically impossible for senders to truthfully report their ifnormation. In more specific environments ([29], [30], [31], [22], [32]) this deviation from truth-telling takes various forms, but universally involves low-ability senders manipulating their messages or actions to mimic the behavior of high-ability senders. Two of these papers, [32] and [31], consider dynamic settings. [31] model information transmission in financial markets, finding that the degradation of information driven by reputational concerns makes it impossible for prices to converge to the market value, no matter how much time passes. Meanwhile, [32] model a reputation-concerned investment manager. They find low-ability senders exaggerate their information in early periods, while discounting newer information in later periods.This key result stems from the fact that high-ability senders learn more quickly, and thus act more decisively on their information early in the game. As we will illustrate below, this same force is at play in our model: low-ability senders endogenously quicken their reports to partially mimic the high-ability senders' ability to learn faster. Relatedly, [33] consider an investment setting, finding that low-ability individuals will act suboptimally due to reputational concerns, where in their setting, consists of herding on the decisions of prior investors. While these latter papers consider dynamic settings, they are not games of timing, i.e., players choose how to act, but not *when* to do so. This is a vital feature of the media application we consider, as newsrooms decide both what and when to report.

The literature on endogenous decision timing is vast. What is relevent to our work, however, is specifically that in which decision timing is of signalling value. This sub-literature considers a variety of different applications ([34], [35], [36], [37]). These papers also come to varying conclusions regarding whether early or later timing is favorable for the sender. In [34], managers decide whether and when to disclose information about their firm to the market, and they obtain that later disclosures are viewed more favorably by the market, and thus of greater signalling value to the firm. Meanwhile, [36] consider a cash-constrained firm that endogenously chooses when to make an investment, which in turn serves as a signal to other investors regarding the project's viability. They come to the opposite conclusion: early investment by the firm serves as a signal of project viability, which in turn will boost the firm's ability to obtain additional investment for its project. We obtain a similar outcome to [36], in that in the equilibrium of our model, early reporting is reputationally advantageous for the sender. Notably, however, while in their setting this incentive to act early is driven by a desire to influence the behavior of investors, in our setting, it is *purely* reputational.

As noted, our paper lies at the intersection of these two separate literatures, as both reputational concerns, and endogenous report timing, are key features of news media setting. However, we are not the first to explore this intersection. [38] do the same, but in the context of a professional forecasting setting. This difference in application, and the associated modeling assumptions, accounts for a reversal in the key predictions we obtain. While we find a reputational benefit from speed arises endogenously, they find the opposite, namely, that the sender's reputation climbs the more she delays. This reversal is similarly true of our finding regarding the dynamics of the sender's reporting behavior. While we find that the sender's reports become more reliable as time passes, they find that the sender errs more as the game progresses.

The remainder of this paper will proceed as follows. In section 2, we present a dynamic model of a reputation-concerned news firm. In section 3, we characterize the equilibrium in a static version of the model presented in section 2, in order to fix ideas and provide intuition for the dynamic model. In section 4, we provide an equilibrium characterization for the full, dynamic model, by building on the static characterization. In section 5, we analyze the equilibrium reputation function, showing that it rewards the sender for both speed and accuracy in reporting. In section 6, we examine dynamics in the sender's reporting behavior, showing that if she is of sufficiently low ability, she becomes more truthful as time passes. Finally, section 7 concludes.

2.2 Model

Here, we present a model of a reputation-concerned sender who learns and reports dynamically.

2.2.1 The game

There is one sender and one receiver. A binary, time-invariant state $\theta \in \{0, 1\}$ is initially unknown to both players. We assume that at the start of the game, both sender and receiver hold prior $Pr(\theta = 1) = \frac{1}{2}$.

The sender has access to a private signal about the state, the quality of which depends on her type. Specifically, the sender's type, denoted by i, may either be "good" ($i = G$) or "bad" ($i = B$). This type is private information: i.e., it is known by the sender, but not by the receiver. We let $R_0 \in (0, 1)$ denote the sender's initial reputation, i.e., the receiver's prior belief that the sender is good. The sender's signal takes a particular form: in any period $t \in \{1, 2, ..., T\}$, θ is privately revealed to the sender with probability λ_i. I.e., in each period, the sender observes a signal r_t:

$$r_t = \begin{cases} \theta & \text{with probability } \lambda_i \\ \emptyset & = \text{with probability } 1 - \lambda_i \end{cases}$$

where $r^t = \emptyset$ indicates that the state was not revealed to the sender at t. Crucially, we impose that good senders are more able learners than bad senders by assuming $\lambda_G > \lambda_B \geq 0$.

At every period t in which the game has not yet ended, after observing r_t, the sender may choose to *report* to the receiver, which consists of choosing a message $m \in \{0, 1\}$ to send to the receiver. The sender can report at most once over the course of the game, i.e., once she reports, the game ends. Alternatively, the sender can choose to *abstain*, which consists of sending message $m = \emptyset$. If the sender abstains, the game continues to the next period (unless $t = T$, in which case the game ends regardless). The sender is not obligated to report, i.e., she can choose to abstain even in the last period, T. Thus, we can interpret time T not as a deadline by which the sender must report, but rather a final opportunity to do so. We let τ denote the time at which a report is made, and denote the absence of a report by $\tau = \emptyset$.

At $\tau + 1$ (i.e., after observing the sender's report, or lack thereof) the receiver observes a private signal $s \in \{0, 1\}$ about θ. Specifically, $Pr(s = \theta) = \pi \in (\frac{1}{2}, 1)$. We assume that this private signal is not fully correlated with the state to avoid the possibility of "off-path" occurences, which will become clear when we define our notion of equilibrium. However, this assumption is for analytical convenience only, and may be relaxed.

The sender's payoff is given by her reputation, i.e., the receiver's belief that she is the good type, after the receiver observes her private signal. We do not specify payoffs for the receiver, as they are immaterial.

2.2.2 Equilibrium

At any given time, p denotes the sender's belief that $\theta = 1$. A (mixed) strategy for the bad firm at time t is given by $\sigma_t(m, p)$, and specifies the probability that the sender sends message m under belief p, conditional on not having yet reported. To ensure that this is a proper distribution, we impose that for all t,

$$\sum_{m \in \{0,1,\emptyset\}} \sigma_t(m, p) = 1.$$

In defining the bad sender's strategy in this manner, we are implicitly restricting the strategy of a sender who has learned the state to not depend on *when* she learned the state. For instance, if it is period t sender who learned that $\theta = 1$ in period t can behave no differently from a sender who learned that $\theta = 1$ in period $t - 1$. Under the notion of equilibrium we will employ, this assumption is without loss, and is used only for notational convenience.

Because the game we present is one of cheap talk, we impose discipline on the equilibrium by restricting attention to equilibria in which the good sender is *truthful*: she sends message θ if and only if she knows the state. Formally, she follows strategy σ^G, where for all t,

$$\sigma_t^G(1, 1) = \sigma_t^G(0, 0) = \sigma_t^G(\emptyset, \frac{1}{2}) = 1$$

This assumption is a selection criterion and not a formal restriction on the good sender's behavior. I.e., the good sender is not a commitment type, we are rather restricting attention to equilibria in which the good sender is truthful. This selection criterion also has an economic rationale: we wish to examine equlibria in which news firms' ability is tied to their integrity, i.e., their judiciousness when reporting.

Next, we consider the the receiver's beliefs about the sender, i.e., the sender's reputation. These beliefs are denoted by a *reputation function*. In particular, $R_t(m, s)$ denotes the receiver's belief that the sender is good, upon observing message $m \in \{0, 1\}$ at time t and private signal s. We let $R(\emptyset, s)$ denote the sender's reputation in the case that she never sends a report. We refer to the collection of these functions as R, or the reputation function.

An *equilibrium* of our model is a perfect Bayesian equilibrium that satisfies our selection criterion. Specifically, two conditions must hold. First, at any given t, all player beliefs (p and R) are consistent with Bayes Rule given the bad (good) sender's strategy σ (σ_G), and in the case of the sender, her private information $r_1, ..., r_t$. Second, the senders' strategies (both type B and type G) must maximize their respective expected reputations at all t and beliefs p they may hold.

In all analysis that proceeds, we will frequently refer to the bad sender simply as the sender, as hers is the only strategy that is not pre-determined in equilibrium.

2.3 Static case: characterization

Before analyzing the full dynamic model above, we will provide a characterization for the static case, in which $T = 1$. This will elucidate certian key results which will extend to the dynamic setting, without having to grapple with the additional analytical complications that dynamics introduce. We will show that there is a unique equilibrium in which a sender who knows the state reports it truthfully, by a sender that does not know the state mixes non-trivially between reporting truthfully (sending message $\emptyset$) and "faking" (sending message 0 or 1). We will also show that in equilibrium, senders are rewarded for accuracy.

For the remainder of this section, we will drop the time index from all functions. All formal proofs for claims presented below will be relegated to the Appendix. We will instead use the main text to outline and provide intuition for these arguments.

2.3.1 Link between strategy and reputation

The below analysis relies heavily on the relationship between the sender's strategy and the reputation function, R. Because the receiver is Bayesian, the reputation function is computed using Bayes Rule as follows:

$$R(m, s) = \frac{1}{1 + L(m, s)\frac{1-R_0}{R_0}}$$

Where $L(m, s)$ is the likelihood ratio of outcome (m, s) for bad senders compared to good senders under σ:

$$L(m, s) \equiv \frac{Pr(m, s|B)}{Pr(m, s|G)}$$ [6]

Thus, the reputation a receiver ascribes to a particular outcome hinges on how likely the outcome is for the good sender as compared to the bad sender. This highlights the key force behind our results: the more weight the bad sender's strategy σ places on a particular outcome occuring, the lower the reputation assigned to that outcome will be. In particular, any outcome that is relatively less likely for the bad sender compared to the good sender will cause her reputation to *improve* compared to her initial reputation R_0, while an outcome that is relatively more likely for the bad sender will cause her reputation to *deteriorate*.

2.3.2 Behavior when informed

In this section, we show that the sender truthfully reports arrivals. I.e., she reports 1 (0) if she has learned that the state is 1 (0). To this end, we begin by establishing two lemmas, pertaining to the sender's strategy and the reputation function, respectively.

We begin by establishing a fundamental difference in reporting behavior between the good sender and bad sender. While a good sender reports 0 (1) only if she has learned that $\theta = 1$ ($\theta = 0$), respectively, this is not the case for the bad sender. I.e., the bad sender's reports are not always truthful: she will, with strictly positive probability, send message 1 (0) even when $\theta \neq 1$ ($\theta \neq 0$).

Lemma 2. *In any equilibrium, for* $\theta \in \{0, 1\}$

$$\sum_{p \neq \theta} \sigma(\theta, p) > 0$$

To see why this must hold, suppose, for example, that the bad sender reports 1 only if she

[6] Formally, $Pr(m, s|i)$ is a function of the sender's strategy, σ^i:

$$Pr(m, s|i) = \frac{1}{2} \sum_{\theta \in \{0,1\}} [\sigma^i(m, \theta)\lambda_i + \sigma^i(m, \frac{1}{2})(1 - \lambda_i)][\pi + (1 - 2\pi)\mathbb{I}(\theta \neq s)]$$

has learned that $\theta = 1$. Because the bad type learns that $\theta = 1$ with strictly lower probability, this means that reporting 1 is relatively *more likely* for the good type. Furthermore, this holds regardless of the receiver's private signal. Thus, a report of 1 *guarantees* a strictly improved reputation for the sender, regardless of the receiver's private signal. Of course, senders would in equilibrium take advantage of this opportunity to guarantee themselves an improved reputation, meaning that message 1 must serve as a profitable deviation at some p.

Next, we establish a property of the reputation function: the sender's reputation is rewarded for being *accurate*, namely, for matching the receiver's private signal.

Lemma 3. *In any equilibrium, for $m \neq m' \in \{0, 1\}$*

$$R(m, m) > R(m, m')$$

Again, this result is driven by the connection between the sender's strategy and the reputation function. Showing that the sender is rewarded for being accurate is equivalent to showing that that conditional on reporting, a good sender is more likely to be accurate than the bad sender. Because the good type's reports are strictly truthful, they are maximally correlated with s. However, this is not the case for the bad type: by Lemma 2, the sender will report $m \in \{0, 1\}$ with strictly positive probability even when the state is not m, and her reports are consequently less correlated with s. It follows that reports by good senders are indeed more likely to be accurate, and thus reputation function must rewards accuracy.

With these two observations in hand, we can now show that the sender must truthfully report arrivals.

Proposition 6. *In any equilibrium, for $\theta \in \{0, 1\}$,*

$$\sigma(\theta, \theta) = 1.$$

Before proceeding, let us discuss the intuition behind this claim. Consider a sender who has learned that $\theta = 1$ who must then decide which message to send. Note that by reporting 1, the sender's report will be more accurate than if she were to report 0. Because by Lemma 3 accuracy strictly benefits the sender, reporting 1 strictly dominates reporting 0. Now let us consider why the sender would not want to abstain. Because the good type is committed to truthful reporting, sending an accurate report signals that the sender knows the state, whereas abstaining signals ignorance. Because the bad type is less likely to know the state than the good type, an informed report makes her appear relatively more reputable than silence. For this reason, $\emptyset$ is dominated as well.

2.3.3 Behavior when uninformed

Proposition 6 tells us that in equilibrium, the bad sender mimics the good sender's strategy when she is informed. I.e., she reports truthfully. Below, we will demonstrate that under non-arrival, this is not the case. In particular, while the good sender abstains (i.e., reports $\emptyset$) when she is uninformed, the bad sender *mixes* non-trivially between *faking*, i.e., reporting despite not knowing the state, and abstaining. Furthermore, we show that she fakes the two messages $m \in \{0, 1\}$ with equal probability.

Proposition 7. *In any equilibrium, for all $m \in \{0, 1, \emptyset\}$,*

$$\sigma(m, \frac{1}{2}) > 0$$

Furthermore, $\sigma(1, \frac{1}{2}) = \sigma(0, \frac{1}{2})$.

To see why the sender must report both 0 and 1 with strictly positive probability when uninformed, recall that by Proposition 6 above, the sender cannot misreport arrivals (i.e., $\sigma(1, 0) = \sigma(0, 1) = 0$). Thus, in order to satisfy Lemma 2, the sender must be reporting both 0 and 1 with strictly positive probability when $p = \frac{1}{2}$. To see why the sender must report $\emptyset$ when uninformed, let's consider what would transpire if she didn't. Proposition 6 would then imply that the bad sender *never* abstains, regardless of her information. Because the good sender does so with strictly positive probability, abstaining would then guarantee a perfect reputation (i.e., $R(\emptyset, 1) = R(\emptyset, 0) = 1$).

Of course, abstaining would then serve as a profitable deviation for the bad sender, a contradiction.

Next, let's consider the second part of the proposition, which states that the uninformed sender must mix equally betwen 0 and 1. To see why this must be true, suppose by contradiction that the bad type reports one message $m \in \{0, 1\}$ with greater probability than the other, m'. Because the good type reports the two messages with equal probability, m' is more likely than m to have come from the good type. This means that m' will on average yield a strictly higher reputation for the uninformed sender than m. This in turn implies that m' will serves as a profitable deviation from m when she is uninformed, violating our result that the sender must be indifferent between the two messages.

2.3.4 Equilibrium existence and uniqueness

In the two subsections above, we showed that in any equilibrium, the sender truthfully reports arrivals, and when she experiences non-arrival, mixes nontrivally between all three potential messages 0, 1, and $\emptyset$. We claim that there exists a unique strategy in this class that constitutes an equilibrium.

Formally, let σ^b for $b \in (0, 1)$ denote the strategy such that

$$\sigma^b(1, 1) = \sigma^b(0, 0) = 1 \text{ and } \sigma^b(1, \frac{1}{2}) = \sigma^b(0, \frac{1}{2}) = b/2$$

Proposition 8. *There exists a unique equilibrium, consisting of a strategy σ^{b^*}, where $b^* \in (0, 1)$*

Existence of this equilibrium is immediate. To see why the equilibrium must be unique, consider what happens to the reputation function, and consequently the sender's value function, as b changes. Specifically, let $R^b(m, s)$ denote the reputation function under strategy σ^b, and $V^b(m, p)$ the sender's value from sending message m under belief p. Note that this value function is a function of the reputation function:

$$V^b(m, p) = \tilde{p} R^b(m, 1) + (1 - \tilde{p}) R^b(m, 0)$$

where $\tilde{p} \equiv p\pi + (1-p)(1-\pi)$ denotes the probability that $s = 1$ is realize, given belief p about the state. Note that in order for σ^b to constitute an equilibrium, by Proposition 7, it must be that

$$V^b(1, \frac{1}{2}) = V^b(0, \frac{1}{2}) = V^b(\emptyset, \frac{1}{2})$$

To show that there exists a unique b such that the value functions satisfy this conditions, we make two observations. First, we consider the value function at the two extreme cases of b. When $b = 0$, i.e., when the bad sender is truth-telling, her lower arrival rate compared to the good sender means that reporting will always strictly improve the sender's reputation, whereas abstaining will always harm the sender's reputation. Thus, reporting yields a strictly higher value for the uninformed sender:

$$V^b(1, \frac{1}{2}) = V^b(0, \frac{1}{2}) > V^b(\emptyset, \frac{1}{2})$$

At the other extreme, when $b = 1$, the bad sender always reports and never abstains, regardless of her belief. Because the good sender abstains with strictly positive probability, this implies that abstaining will guarantee a perfect reputation. So in this case, abstaining will yield a strictly higher value for the uninformed sender:

$$V^b(1, \frac{1}{2}) = V^b(0, \frac{1}{2}) < V^b(\emptyset, \frac{1}{2})$$

Next, let's consider what happens to the value functions as b increases. Note that as b increases, the bad sender is on average reporting more and abstaining less. Thus, the reputation function will respond by assigning increasingly higher reputation to senders who abstain, and lower reputation to senders who report. This will in turn be reflected in the value functions: the value of reporting is strictly decreasing in b, whereas the value of abstaining is strictly increasing in b (see Figure 1). Given our observations about V^b at the endpoints, this implies that there exists a *unique* point $b^* \in (0, 1)$ such that indifference is achieved.

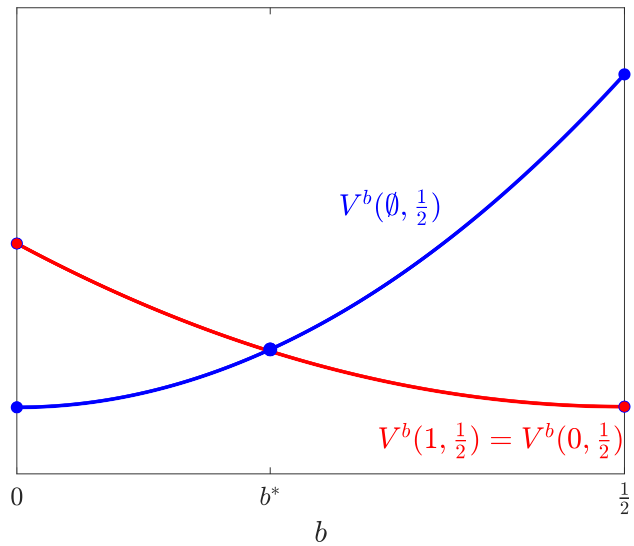

Figure 2.1: Value function associated with σ^b. b^* denotes unique value such that indifference condition is satisfied.

2.3.5 Discussion

To summarize, we have shown that the static equilibrium is one in which the bad sender truthfully reports arrivals, but when she does not know the state, she mixes between abstaining and "faking", i.e., making an uninformed report. Furthermore, this equilibrium is one in which senders are endogenously rewarded for making *accurate* reports (Lemma 3).

In fact, this reward for accuracy is key to sustaining the uninformed sender's indifference between reporting and abstaining. Because good types are better learners and thus more likely to make an informed reports, senders are rewarded for reporting. However, the sender is also penalized for making an inaccurate report (an outcome which is relatively more likely for the uninformed sender), which deters the sender from reporting. In equilibrium, the reputation function is such that these incentives precisely counteract each other when the sender is uninformed, yielding indifference.

As we will discuss below, these same incentives and reporting behaviors prevail under the equilibrium of the dynamic model, in particular the reputation function's endogenous rewarding of accuracy. The dynamic model, however, further sheds light on how the passage of *time* impacts both the sender's behavior and her reputation. Understanding these dynamics will be the focus of the sections that follow.

2.4 Dynamic model: Faking in equilibrium

Above, we showed that in any equilibrium, the sender truthfully reports arrivals, and when she experiences non-arrival, mixes nontrivally between all three potential messages 0, 1, and $\emptyset$. In this section, we show that the equilibrium in the dynamic setting must also take this form. In particular we show that the sender fakes with positive probability in every period.

2.4.1 Correlation between state and type

Showing that the dynamic characterization takes the same qualitative form as in the static setting presents unique challenges. One of these involves potential correlation between the state θ and the sender's type i. In the static model, we relied on the assumption that they were uncorrelated, and this assumption was responsible for the resulting symmetry of the sender's strategy across messages (for instance, that $\sigma_t(1, \frac{1}{2}) = \sigma_t(0, \frac{1}{2})$). However, one cannot assume this to be true in a dynamic setting. This is due to the fact that the sender's strategy may *induce correlation* between type and state in periods beyond the first. Such correlation will in turn induce behavior in the subsequent periods that violates this symmetry across messages that we obtained in the static setting. In this subsection, we demonstrate that although possible in general, such correlation between the state and the type can never arise in equilibrium.

To this end, we begin by formalizing a notion we refer to as *silence symmetry*.

Definition 2. A strategy is silence symmetric if at time t

$$\sigma_t(\emptyset, 1) = \sigma_t(\emptyset, 0)$$

In words, this condition stipulates that the sender if the sender chooses to withhold a report despite knowing the state (i.e., remaining silent), she must be equally likely to do so in each state. We are concerned with silence symmetry because correlation between the sender's type and the state arises precisely when it is violated.

To fix ideas, let's consider examples of strategies which both satisfy and fail this condition.

A simple example of silence symmetry is the case in which the sender is truthful: regardless of the state, an informed sender withholds a report with probability 0. Another simple example of silence symmetry is one in which the sender *never* reports in a given period, regardless of her information. In this case, she withholds a report with probability 1, regardless of the state.

Let us now consider an example of a strategy in which silence symmetry is violated. Consider the strategy $\hat{\sigma}_t$ in which the sender truthfully reports 1, but abstains otherwise, i.e.,

$$\hat{\sigma}_t(1,1) = 1, \hat{\sigma}_t(\emptyset, 0) = \hat{\sigma}_t(\emptyset, \frac{1}{2}) = 1 \tag{2.1}$$

In this extreme example of silence asymmetry, the sender *always* withholds a report when she knows the state is 0, but *never* does so when she knows the state is 1. To understand why the failure of silence symmetry can cause correlation between the sender's type and the state, consider the case where $T = 2$ and the bad sender plays $\hat{\sigma}_1$ in the first period. This implies that the bad sender is more likely to survive into the second period (i.e., not report in the first period) under $\theta = 0$. Specifically, when $\theta = 0$, she survives with probability 1, but if the state is 0, she only survives with probability $1 - \lambda_B$. However, this is not the case for the good sender: because she is truthful in equilibrium, she is equally likely to reach period 2 regardless of the state, specifically with probability $1 - \lambda_G$. This in turn implies that at the beginning of the second period, there exists a correlation between the sender's state and the type: $\theta = 0$ with greater probability if the sender is bad than if the sender is good.

2.4.2 Strategy in equilibrium

Having illustrated the importance of silence symmetry, we begin by establishing a set of necessary conditions that must hold in any equilibrium of the dynamic setup. Among these conditions is silence symmetry.

Lemma 4. *In equilibrium, at all t,*

1. σ_t *is silence symmetric.*

2. $\sigma_t(1, 0) = \sigma_t(0, 1) = 0$.

3. $\sigma_t(1, \frac{1}{2}) = \sigma_t(0, \frac{1}{2}) > 0$.

Proof. See appendix. □

In addition to silence symmetry, we establish two facts that also held in the static setup: namely, that the sender never reports the "wrong" message when she is informed about the state (point 2) and that the sender "fakes" a report with positive probability in every period (point 3). However, unlike the static setting where we were able to show both these points directly, in the dynamic setting, we must take a forward induction approach. For instance, silence symmetry at period t is necessary to show that conditions 2 and 3 hold in $t + 1$. To see why this is the case, recall that a failure of silence symmetry in period t would imply that the state is correlated with the sender's type in $t + 1$. This could make one of the two reports (in particular, the one that was withheld less by the bad sender in t) better for the sender's reputation in $t + 1$, all else equal. This could in turn cause violations of conditions 2 and 3 above.

While these necessary conditions substantially narrows the set of candidate equilibria, it still leaves open the possibility that the sender withholds reports with positive probability in a given period, i.e., that $\sigma_t(\emptyset, 1) = \sigma_t(\emptyset, 0) > 0$. We next rule this out as a possibility, establishing that the sender's equilibrium strategy in every period takes the same form as in the static model.

To this end, we begin by defining an T-period equivalent to the class of strategies we defined in the static setting. For any $b \equiv (b_1, ..., b_T) \in (0, 1)^T$, let σ^b denote the strategy such that

$$\sigma_t^b(1, 1) = \sigma_t(0, 0) = 1 \text{ and } \sigma_t^b(1, \frac{1}{2}) = \sigma_t^b(0, \frac{1}{2}) = b_t/2$$

This strategy is one in which, in every period, the sender truthfully reports arrivals, and mixes nontrivially between all three messages when she does not, specifically reporting 0 and 1 with equal probability. We now claim that any equilibrium must belong to this class.

Proposition 9. *In equilibrium, the sender's strategy is given by σ^b, for some $b \in (0,1)^T$.*

Proof. See appendix. □

Let us now take stock of this result. First, note that the sender in equilibrium fakes a report with positive probability in *each* period setting. I.e., faking never ceases. Second, while the equilibrium strategy here is a T-period extension of the static equilibrium strategy, it is dynamic in nature. Specifically, the probability with which the sender reports when she has not learne the state, b_t changes over time. We will explore the dynamics in b_t later in the section. However, we will begin by examining the equilibrium reputation function, and the dynamics it entails.

2.5 Reputation: accuracy and speed

Here, we consider how reputation is assessed dynamically in equilibrium. We begin by decomposing the sender's reputation into two components: one capturing the impact of time's passage on her reputation, the other denoting the impact of the report itself, and its associated accuracy, on her reputation. First, we establish that, as in the static setting, the sender is rewarded for accuracy. We specifically show that making an accurate report will cause the sender's reputation to strictly increase from where it was before she reported, whereas an inaccurate report will cause her reputation to strictly decrease. We further examine the relative magnitudes of these changes, and find that the reputational decline that comes from making an inaccurate report strictly exceeds the reputational gain from making an accurate one. Finally, we find that the sender's reputation is rewarded for speed: with every time increment that passes without the sender making a report, her reputation deteriorates.

In order to facilitate the aforementioned decomposition of the sender's reputation function, we begin by defining R_t for $t \in \{0, 1, ..., T\}$, the sender's *interim reputation*. This denotes the receiver's belief about the sender's type conditional on the event that she has not reported at any $s \leq t$. When $t = 0$, R_0 is merely the sender's prior reputation. For all $t \geq 1$, R_t is an equilibrium object, and can be computed recursively using Bayes' Rule, given the sender's strategy. Under the equilibrium characterization we obtain above, this recursive definition takes the following form:

$$R_t = \frac{1}{1 + \frac{1-R_{t-1}}{R_{t-1}} \frac{(1-\lambda_B)(1-b_t)}{1-\lambda_G}}$$

Now let us interpret this object. While in our model the sender is solely concerned with her reputation at the end of the game, R_t captures how the sender's reputation dynamically evolves as the game progresses. Specifically, R_t specifies the sender's reputation at periods preceding that in which she reports. For instance, if the sender were to report in period $t = 5$, R_4 denotes where her reputation stood immediately before her report was made.

By linearly separating the sender's interim reputation from her reputation function, it can be written in the following form:

$$R_t(m, s) = R_{t-1} + \alpha_t \mathbb{I}(m = s) + \beta_t \mathbb{I}(m \neq s) \tag{2.2}$$

where, like R_t, α_t and β_t are equilibrium objects and functions of the sender's equilibrium strategy. Note that (2.2) implicitly assumes that the sender's reputation depends on the content of her report (whether $m \in \{0, 1\}$) only to the extent that it impacts her accuracy (i.e., whether $m = s$). While we have omitted a formal proof of this, the reputation function can be written in this way due to the fact that in equilibrium, the sender's strategy exhibits symmetry across messages $m \in \{0, 1\}$ and states $\theta \in \{0, 1\}$ in every period.

Now, let us interpret this decomposition. As discussed above, the first component, R_{t-1}, denotes the sender's reputation immediately prior to her time t report. This component thus captures all the dynamic change in her reputatation that happened prior to her report, and specifically, the impact that *delay* had on her reputation. Meanwhile, the residual component, $\alpha_t \mathbb{I}(m = s) - \beta_t \mathbb{I}(m \neq s)$, denotes the change that occurs in the sender's reputation the moment in which she reports. This component accounts for the impact that accuracy has on the sender's reputation. Note that the α_t and β_t are time-contingent: this is due to the fact that the magnitude of impact accurate (or inaccurate) reporting has on the sender's reputation, like R_t, is dynamic in nature. Figure 2 below illustrates how the sender's reputation evolves dynamically, and sheds further light on this decomposition: the dotted lines denote time paths of the sender's reputation, when she reports at $t = 5$, either accurately (blue line) or inaccurately (red line). Her reputation at all periods prior to her report are given by R_t, and the final jump in her reputation that occurs in the period in which she reports is given by α_t or β_t, if her report was accurate or inaccurate, respectively.

2.5.1 Accuracy

With this decomposition in hand, we first seek to understand the role that accuracy plays in the sender's reputation. Our findings are summarized by the following proposition.

Proposition 10 (Accuracy and reputation)**.** *In any equilibrium*

1. $\alpha_t > 0$ *and* $\beta_t < 0$.
2. $-\beta_t > \alpha_t$.

Proof. See appendix. □

This proposition makes two separate claims. Let us begin by understanding the first, i.e., that $\alpha_t > 0$ and $\beta_t < 0$ for all t. This claim asserts that at any given point in time, a correct report will cause an increase in the sender's reputation compared to where it stood immediately prior to the report, whereas an incorrect report will cause a decrease. In part, this is an extension of Lemma 3 from the static model, in that it implies that in any period, a sender's reputation from making a correct report, $R_t(1, 1)$, strictly exceeds her reputation from making an incorrect report, $R_t(1, 0)$. As in the static model, this is driven by the fact that the bad sender fakes with strictly positive probability in any given period, while the good sender does not. Because faking is associated with inaccurate reporting, inaccuracy is thus reputation-damaging.

However, this claim goes a step further than this: it additionally asserts that the sender's reputation following a correct report must strictly exceed her interim reputation immediately prior to that report, whereas her reputation following an incorrect report must lie strictly below her interim reputation. That is, accuratereporting causes an dynamic boost in the sender's reputation, whereas inaccurate reporting causes a deterioration. The intuition for this is clear: suppose by contradiction that both α_t and β_t were positive. Then, all time-t senders could guarantee an improvement in their reputations by reporting. It would thus follow that at the end of time, the sender's reputation will improve with probability 1, regardless of her ability. Thus, the reputation function must not be consistent with Bayes Rule, a contradiction. If, instead we assume that both α_t and β_t were negative, no sender would choose to report at time t, as it will cause her reputation to decline with probability 1, even if her report is accurate. Instead, she would elect to abstain, in which case her reputation would evolve to R_t, which must, in order to be consistent with a negative α_t and β_t, lie above R_{t-1}.

Let us now consider the second component of the claim: namely, that $-\beta_t > \alpha_t$. This claim asserts that the reputational deterioration resulting from an inaccurate report *strictly exceeds* the reputational growth that occurs with an accurate report. Economically, this is a compelling claim: it tells us that that the sender has more to lose from an erroneous report than she has to gain from a correct one. Consequently, if an uninformed sender chooses to fake a report, she is choosing to take a binary gamble where the cost incurred when losing (i.e., making an incorrect report) outweigh the benefit earned when winning (i.e., making a correct report).

This result is due entirely to both the features of the sender's strategy and the reputation function in equilibrium. We begin by observing that the equilibrium is one that is strictly informative about the sender's type, i.e., the reputation function is not constant. It follows that reporting behavior exhibited more often by the bad type will cause her reputation to *decline* on average. This holds in particular when the sender *fakes*, as it is done with positive probability by the bad sender, but never by the good one. Because the uninformed sender holds a belief $\frac{1}{2}$ about the states, a sender who fakes will be accurate half of the time. Thus, her expected reputation from faking at time t is given by

$$R_{t-1} + \frac{\alpha_t - \beta_t}{2}$$

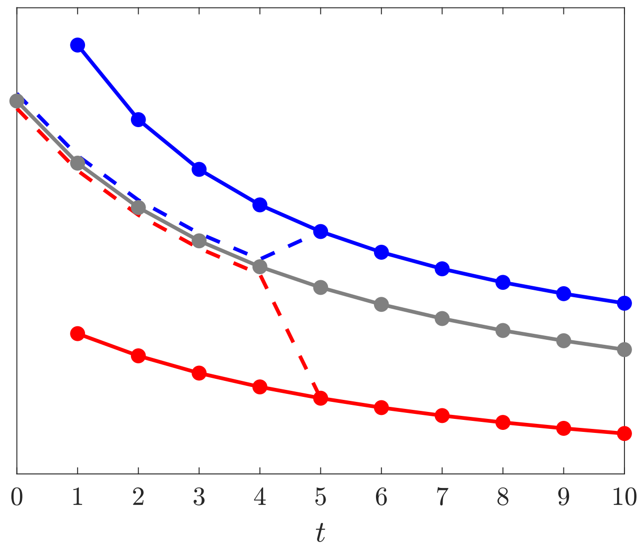

Figure 2.2: Dynamics of the sender's reputation function. Grey line denotes R_t, solid red (blue) line denotes reputation from making an accurate (inaccurate report) at t. Dotted red (blue) line shows dynamics of reputation for a sender who makes an inaccurate (accurate) report in period 5.

In order to ensure that the faking sender's reputation declines, it must be that $\alpha_t < -\beta_t$. Inuitively, the sender must be penalized for faking, otherwise, she would be guaranteed to sustain a reputation of at least R_{t-1} by the end of the period t, which is not possible in any equilibrium that is informative about the sender's type.

2.5.2 Speed

Next, we consider how the sender's *timing* affects her reputation. Our key claim is the following, which asserts that delaying reporting strictly damages the sender's reputation.

Proposition 11 (Speed and reputation)**.** *In equilibrium, R_t is strictly decreasing in t.*

Proof. See appendix. □

The proof for this claim follows a backwards induction argument. We begin by arguing that $R_T < R_{T-1}$. To see why this must hold, note that $R_T = R_T(\emptyset, 1) = R_T(\emptyset, 0)$, i.e., R_T equals the reputation the sender enjoys in the event that she stays silent through the course of the entire game.

If we assume by contradiction that $R_T \geq R_{T-1}$, abstaining at time T when uninformed serves as a profitable deviation: as we demonstrated above, this is due to the fact that, as we showed above, faking at time T yields a reputation strictly less than R_{T-1}.

Next, let's consider an arbitrary period t, and assume by contradiction that $R_t \geq R_{t-1}$. Unlike T, it does not directly follow that uninformed sender can profitably deviate by abstianing at t. This is because the game does not end at $t+1$, and her reputation may decline following this. In particular, even if the sender starts period $t+1$ with a relatively high reputation R_t, if b_{t+1} is sufficiently high, her value from reporting in the next period may be relatively low. This is because the more the bad sender reports on average (which is associated with a higher b_{t+1}), the worse reporting in the next period will be for her reputation, regardless of her accuracy. Thus, in order to show that the sender can profitably deviate by abstaining in period t when she is uninformed, we must show that b_{t+1} remains relatively small.

As we will now show, this in fact follows directly from making the inductive assumption that $R_t > R_{t+1}$. Specifically, $R_t > R_{t+1}$ implies that b_{t+1} lies below some threshold $\bar{b}$. [7]

There is a clear intuition for this: the more the bad sender fakes (i.e., the higher b_{t+1} is), the *better* abstaining must be for the sender's reputation, as it is less probable for the bad sender. This in turn results in a higher R_{t+1}. Thus, b_{t+1} must lie below the above threshold to ensure that R_{t+1} does not exceed R_t. As we show formally in the appendix, this bound we obtain is sufficient to show that the uninformed sender at t can profitably deviate by abstaining at t, thus showing that we cannot have $R_t \geq R_{t-1}$.

2.5.3 Discussion

In this section we have shown that the equilibrium reputation function rewards the sender for two aspects of her report: speed and accuracy. Thus, although the sender is reputation-maximizing, she behaves as if she is maximizing some combination of speed and accuracy. Notably, in canonical settings in which an individual who is intrinsically interested in learning must decide how long to

[7] Formally, $\bar{b} \equiv \frac{\lambda_G - \lambda_B}{1-\lambda_B}$. It is the unique value of b_t such that the good and bad sender abstain with equal probability, i.e., it is the solution to the following equality: $1 - \lambda_G = (1-\lambda_B)(1-b)$.

learn for, it is often assumed that her payoff is an increasing function of speed and accuracy (e.g., [39]). We have shown that in a reputational setting, the same holds, but arises endogenously. That is, while speed and accuracy may be of intrinsic value in the canonical setting, they are in equilibrium of signalling value in a reputational setting.

The dynamics of the sender's reputation, as well as the importance of both speed and accuracy, are shown in Figure 2, which illustrates our findings from this section. As shown, the sender's reputation is dynamically discounted until the time in which she reports (Proposition 11). Once she reports, her reputation will exhibit a strict upwards jump if her report was accurate, and a downwards jump if it was inaccurate (Proposition 10). Also in line with Proposition 10, we see that the loss in reputation the sender sustains from making an inaccurate report exceeds the gain she enjoys when she makes an accurate one.

2.6 Dynamics in strategy

Having examined the equilibrium reputation function, we now turn to the sender's strategy. In this section, we examine dynamics in the sender's strategy. We show that b_t is strictly decreasing in t, provided that λ_B lies below some bound. We now state the result:

Proposition 12 (Declining b_t)**.** *In equilibrium, there exists a* $\underline{\lambda} \in (0, \lambda_G)$ *such that if* $\lambda_B < \underline{\lambda}$, b_t *is strictly decreasing in* t.

Proof. See appendix. □

To understand why this result holds, first recall Proposition 11, which tells us that conditional on not having reported, the sender's reputation at the beginning of period $t+1$, R_t, must be strictly less than her reputation at the beginning of period t, R_t. This means that all else equal, and in particular if the bad sender were to employ the same strategy in the both periods (i.e., if $b_t = b_{t+1}$), reporting at time t provides the sender with a higher reputation than reporting at time $t+1$ regardless of whether or not the report was accurate. For the sender who is uninformed about the state at t, this makes faking at time t particularly attractive, as it ensures that she will enjoy a higher reputation on average than she would by faking at time $t + 1$ instead.

Thus, in order to ensure that the uninformed sender at t does not have a strict incentive fake at time t, thus violating the indifference condition, she must be compensated for waiting through two separate channels. First, by waiting, she will with positive probability learn the state in the next period, and increase her chances of making an accurate report, which she is rewarded for. If however this probability (λ_B) is sufficiently small, she must instead be compensated through a different channel: the reputation function. In particular, although the sender starts off $t + 1$ with a lower reputation, she must exhibit a greater boost in reputation from actually reporting. To see why this implies that $b_t > b_{t+1}$, consider the relationship between b_t and the reputation function. Because a higher b_t is associated with a higher probability that the bad sender will report, the reputation function will respond by assigning a *lower* reputation to reporting at time t (regardless of the sender's accuracy). Thus, a greater boost in reputation from reporting at $t + 1$ will happen only if $b_t > b_{t+1}$.

Economically, this result tells us that a sender who is of sufficiently low ability becomes *more truthful* as time passes, i.e., is less likely fake a report. This indicates that the sender's approach to reporting changes as the game progresses. As we have described above, the sender has two separate means of convincing the receiver she is of high ability: the speed of her report, and its accuracy. While both factors are present throughout the course of the game, this result indicates that in earlier periods, the bad sender relies relatively more on the speed approach to convince the receiver she is good, accomplished through a higher b_t. However, as time passes, she gradually shifts her behavior towards the accuracy approach, by trying to convince the receiver she is good through her reluctance for dishonesty.

2.7 Conclusion

In this paper, sought to understand the nature of misreporting on the part of news firms when firms are reputation-maximizing. To this end, we have presented sender-receiver model tailored to the newsroom environment. In particular, senders learn dynamically and must choose when to time their report. Furthermore, given the fast-paced nature of the news cycle, we assume that sender's exclusively seek to maximize their reputations at the end of the game.

Our analysis gave rise to several findings, regarding both the sender's reporting behavior as well as her incentives in equilibrium. First, we find that in equilibrium, the sender will fake a report with positive probability in every period, which entails reporting in the absence of knowledge about the state. This behavior gives rise to a higher level of misreporting than would happen if instead the sender's reports were exclusively informed. Next, we examine how reputation is endogenously assessed in equilibrium, finding that the sender is endogenously rewarded two separate qualities of the sender's report: speed and accuracy. We thus provide a reputational microfoundation for the speed-accuracy tradeoff in the newsroom setting. We make a further observation about the nature of accuracy's importance: senders are in equilibrium penalized strictly more for inaccurate reporting than they are rewarded for accurate reporting, meaning that erroneous reporting is much more consequential for the sender's reputation. Finally, we explore dynamics in the sender's reporting behavior, finding that if the sender is of sufficiently low ability, she will become strictly more truthful as time passes. This implies that misreports by a sender who has a low a capability for learning are most probable in the immediate aftermath of obtaining a lead, and becomes less probable as time passes.

While our model provides a simple framework for understanding reputation-driven misreporting by news media, some of our findings warrant further investigation in more general frameworks. In particular, it is not clear that the endogenous reputational reward for both speed and accuracy obtained under our model will exist under a general learning structure on the part of the sender, and even relaxing our selection assumption regarding the good sender's reporting behavior. Furthermore, it is unclear if additional reporting incentives arise in a more general framework, beyond just speed and accuracy. These questions are the subject of continued research.

Chapter 3: Dynamic Reputation-Driven Media Bias

3.1 Introduction

What causes media bias? This is a question that has been generously explored in a literature surveyed by [40]. One explanation, offered by [22], is that this bias may be driven by a news firm's efforts to appear reputable. They show that a low-ability firm, i.e. one with access to low-quality information, will bias its reports in response to reputational concerns. In particular such a firm will in equilibrium bias its reports in favor of a consumer's prior.

This model is static. That is, it considers a setting where news firms report once and do not communicate repeatedly with a consumer. However, the news media setting is inherently dynamic, with firms making numerous reports throughout the day and establishing long-term relationships with consumers.

In this paper, I seek to understand how the dynamic nature of news reporting may impact reputation-driven media bias. Specifically, I consider how a reputation-concerned firm will behave when it makes repeated, incremental reports. I further consider how a firm may behave differently if it is concerned with its long term, rather than immediate, reputation.

To this end, I extend the model in [22], hereafter GS, model to a dynamic setting. In my model, rather than receiving a single signal and reporting once as in GS, firms learn incrementally and report repeatedly over time. Crucially, because high-ability firms have access to better information, their signals will also be more correlated over time. The firm seeks to maximize its current and future reputations, i.e., the consumer's belief that it is of high ability.

I begin the analysis by characterizing the equilibrium in each period under the assumption that the firm is *myopic*, i.e., at every time is exclusively interested in maximizing its immediate reputation. I show that under appropriate selection criteria, there exists a unique equilibrium. The firm's reporting behavior in any given period will involve biasing its signals, as in GS. However, unlike GS where bias is driven purely by an appeal to the consumer's prior, firm bias is driven by two separate factors. The first is the bias in the firms prior, precisely the factor driving bias in GS's static setting. However, unlike GS, there is a second firms also bias their reports in order to appear consistent with their past reports. Firms engage in such behavior in equilibrium becuase high-ability firms have access to higher quality information, and thus will receive signals that are more correlated with each other over time. A bad firm will thus mimic such correlation by biasing is reports to be consistent with its previous history of reports. Notably, while role of the consumer's prior is static in nature, the report history is dynamic, and thus this consistency bias will change over time. Furthermore, I illustrate the role of the firm's prior will become gradually less important over time as the firm accumulates a richer history of reports. Economically, this result indicates that in the long term, reputation-driven media bias may not take the form of pandering as it does in GS, but rather that of self-corroboration.

I then relax the assumption that the firm is myopic, and characterize the equilibrium under a forward-looking firm. Such a firm seeks not only to maximize its immediate reputation but also future ones. With relative ease, I am able to show that there exists a unique equilibrium in this case as well, and it is precisely the myopic equilibrium. That is, a forward-looking firm will behave no differently than a myopic one in equilibrium, and thus assuming myopia is without loss in this setting. The proof for this result relies critically on the fact that the myopic equilibrium is pooling, and thus the firm's report reveals nothing about its ability. This proof is not only of interest in the context of this model, but more broadly. It illustrates that the assuming myopia, which is sometimes done in dynamic games of reputation to ensure tractability (e.g., [32]), such an assumption may be without loss in settings when the equilibrium is pooling.

I then test the limits of this result. Specifically, I ask whether one can show that it is without loss to assume myopia under a similar model that gives rise to equilibrium separation, i.e., one where the firm's reputation is sensitive to its reporting behavior in equilibrium. To explore this question, I extend the baseline model by incorporating lying costs. Formally, in addition to seeking reputation, firms incur a fixed cost from misreporting their signals. Under myopia, there again exists a unique equilibrium. While similar in nature to the equilibrium of the baseline model, firms will unsurprisingly bias their reports less in the face of lying costs, and will be truthful whenever the lying cost is sufficiently high. Importantly, this equilibrium involves separation: the firm's reputation will react to its reporting behavior. There is a clear intuition for why lying costs introduce such separation. In the baseline model, if one report yields the firm a higher reputation than the other, the bad firm will exclusively report the reputation-maximizing message, thus causing the equilibrium to collapse. However, the same cannot be said if lying costs are present: even if one report yields a higher reputation, as long as the cost of lying is high enough, the firm may not find it optimal to remain truthful. I then consider a forward-looking firm, and establish that assuming myopia can be with loss in this setting. While this result is specific to the model at hand, it indicates in general, it may be with loss to assume that senders are myopic in dynamic games of reputation.

Related Literature This paper relates broadly to the literature on media bias, of which [40] provide a review. More specifically, it is a model of reputation-driven media bias, a literature which includes GS and others (e.g. [41], [42], [43], [21]). In such models, bias emerges not due to an effort by firms to mislead, but rather as a response to consumer preferences and beliefs. For instance, one can conceive that consumers possess certain beliefs and receive utility in part from making the right decisions, but also receiv utility from having their beliefs confirmed. Such psychological utility models of bias formation include [44] and [21]. Other papers demonstrate that such bias can arise even when firms face rational consumers who receive no utility from having their beliefs confirmed. In one such paper, [43] consider a setting where news parties compete electorally, and news firms make reports regarding their respective news platforms. They find that given the discretion to send cheap talk messages may result in biased reporting by a news firm, but that such cheap-talk messages will discipline the political parties platforms in such a way that is beneficial to consumers. In contrast, GS show that biased reporting may arise even in the absence of political competition, provided that the news firm is reputationally concerned. Unlike [43], bias is strictly harmful to consumers in this setting, as it causes a deterioration in the informativeness of the firm's reporting. Similar reputation-driven bias arises in [45], who consider a labor market setting. This paper departs from much of the current litature, including GS, by considering a dynamic setting. While this would not be the first paper to consider media bias in a dynamic setting (e.g., [15]), to my knowledge it is the first to consider dynamics in the reporting behavior of news firms.

Beyond the literature on media bias, this paper also relates to the literature on career concerns. Notably, [46] considers an employee whose type is unobservable to an employer, but who can indirectly reveal her type through performance. More directly relevant to the model I present is [32]. They consider a dynamic reputation model in which a portfolio manager biases her investment decisions in order to appear knowledgeable. Similar to the verison of the model I present which incorporates lying costs, managers incur a cost from making investing decisions which are not optimal. In equilibrium, the manager's decision involves overreacting to new information in the short-run and under-reacting to new information in the long-run. This long-run underreaction is similar in nature to the consistency effect I document in this paper: the manager underreacts to new information because high-ability firms will find it optimal to be more consistent with their investing decisions. This paper also relates to [32] in that it explores the consequences of their assumption that the portfolio manager is myopic. While this assumption was made in the name of tractibility, its consequences remain unclear. In this paper, I show that in a binary dynamic reputation setting, myopic and forward-looking equilibria need not align. Finally, an effort to be consistent is also demonstrated in [47]. In their setting, a principal who wishes to only hire high-ability forecasters must design a mechanism for evaluating forecasters. Under the optimal mechanism, firm engage in strategic behavior to improve their hiring prospects, specifically by striving to appear more consistent in its behavior than it is in reality.

The remainder of this paper will proceed as follows. In section 2, I present the baseline model. In section 3, I characterize the equilibrium under myopia, and consider the economic implications of this equilibrium in section 4. In section 5, I relax the assumption that the firm is myopic, and show that assuming myopia is without loss in the baseline model. In section 6, I augment the baseline model to incorporate lying costs, characterizing the equilibrium under myopia and showing that the assuming myopia is with loss in this setting. Finally, in section 7 concludes.

3.2 The Model

3.2.1 The game

There is a binary, time-invariant state $\theta \in \{R, L\}$. We can interpret $\theta = L$ ($\theta = R$) as denoting that a left (right) leaning candidate or policy platform is superior to a right (left) leaning one.

There are two players: one firm and one consumer. Time is discrete, and there are T periods: $t \in \{1, ..., T\}$. At the start of the game, the consumer and firm hold a common prior p_0 regarding the state, where $p_0 = Pr(\theta = R)$. The firm dynamically learns and reports about the state to the consumer. Namely, in each period t, the firm privately observes a signal $s_t \in \{r, l\}$ and makes a report $\hat{s}_t \in \{r, l\}$. The s_t are a sequence of Bernoulli draws. Specifically, $Pr(s_t = \theta) = \pi_i$, where $\pi_i \in [\frac{1}{2}, 1)$ and $i \in \{G, B\}$ is firm's type. This type denotes whether the firm is *good* ($i = G$) or *bad* ($i = B$). I assume that $\pi_G > \pi_B$, and thus that good firms have access to more informative signals than bad firms. This type is time-invariant and private information: it is known to the firm but not to the consumer, who is endowed with prior $\lambda_0 \in (0, 1)$ that $\theta = 1$.

In each period, the firm's objective is to maximize an increasing function of its future reputations. Formally, let λ_t denote the firm's reputation in time t, i.e., the consumer's belief at the end of time t (i.e., after having observed the $\hat{s}_t$ that $i = G$. The firm's time-t payoff is then given by $F_t(\lambda_t, ..., \lambda_T)$, where F_t is continuous and weakly increasing in λ_s for all $s \in \{t, .., T\}$, and strictly increasing in λ_s for some $s \in \{t, ..., T\}$.

3.2.2 Defining equilibrium

Let $\hat{\mathcal{H}}_t \equiv \{r, l\}^{t-1}$ denote the space of time-t report histories. In each period $t \in \{1, ..., T\}$, the bad firm chooses a strategy $m_t : \{r, l\} \times \hat{\mathcal{H}}_t \to [0, 1]$. I.e., a strategy specifies, for every possible current signal and report history, the probability with which the firm reports r in that period.[1] Meanwhile, a time-t reputation function $\lambda_t : \hat{\mathcal{S}}_{t+1} \to [0,1]$ assigns a reputation to each report history for the firm.[2]

An equilibrium is a triple of functions $(m_t^B, m_t^G, \lambda_t)$ for each t such that: (1) m_t^i maximizes the type-i firm's time-t payoff at each t, given the future strategies $(m_{t+1}, ..., m_T)$ and current and reputation functions $(\lambda_t, ..\lambda_T)$ and (2) λ_t is consistent with Bayes Rule under every report history that occurs with positive probability given the $\{m_1, ..., m_T\}$.

Finally, following in the precedent set by GS, I impose the following selection criterion: the good firm is truthful in equilibrium. This criterion selects the equilibrium where the good firm's reports are maximally informative about the state, and thus one where the bad firm seeks to mimic this informative signal. I formalize this selection criterion below as Criterion 1.

Criterion 1. *The good firm is truthful. I.e., at all t, for any current signal s_t and report history $\hat{h}_t = (\hat{s}_1, ..., \hat{s}_{t-1})$,*

$$m_t^G(s_t, \hat{h}_t) = \begin{cases} 1 & \text{if } s_t = r \\ 0 & \text{if } s_t = l. \end{cases}$$

3.3 Equilibrium Under Myopia

I begin by characterizing the equilibrium under the special case where the firm is *myopic*, i.e., the bad firm maximizes its immediate reputation. Formally, F_t is strictly increasing in λ_t, but constant in λ_s for all $s > t$. We begin with the myopic case for two reasons. First, it illustrates the core intuition behind the equilibrium characterization without having to consider the complications that a forward-looking firm may introduce. But more importantly, I will ultimately rely on the myopic characterization to characterize the forward-looking equilibrium, and show that they are equivalent.

[1] I am implicitly assuming that the firm can only condition its reporting behavior on its current signal, and not the history of its signals. This memorylessness assumption is made for analytic and notational simplicity.

[2] While a time-t report history consists only past reports (and does not include$\hat{s}_t$), the time-t reputation function λ_t assigns a reputation to each history of past and current reports (including $\hat{s}_t$).

3.3.1 Formal characterization

Before proceeding, let us introduce some notation. First, fix any time t report history $\hat{h} = (\hat{s}_1, ..., \hat{s}_{t-1})$. Then, let $p^i(\hat{h})$ denote the consumer's (interim) posterior belief that $\theta = 1$, conditional on both history $\hat{h}$ and the firm being of type i. Further, let $q^i(\hat{h})$ denote the consumer's belief that she will receive an r report in period t, again conditional on history $\hat{h}_t$ and the firm being of type i. Because, the good firm is truthful, in any equilibrium q^G follows simply from Bayes Rule:

$$q^G(\hat{h}) = p^G(\hat{h})\pi_G + (1 - p^G(\hat{h}))(1 - \pi_G). \tag{3.1}$$

Meanwhile, because the bad firm is not necessarily truthful, an equilibrium q^B depends on not only the quality of the bad firm's signal, but also its strategy.[3] Specifically, Bayes Rules yields the following:

$$q^B(\hat{h}) = p^B(\hat{h})(\pi_B m_t(r, \hat{h}) + (1-\pi_B)m_t(l, \hat{h})) + (1-p^B(\hat{h}))(\pi_B m_t(l, \hat{h}) + (1-\pi_B)m_t(r, \hat{h})). \tag{3.2}$$

One can characterize the equilibrium behavior in the first period in terms of the q^i. Specifically, I show that under myopia, $(m_t, \lambda_t)_{t=1,...,T}$ is an equilibrium if and only if $q^B(\hat{h}) = q^G(\hat{h})$ at every report history $\hat{h}$. That is, from the perspective of the consumer, good and types must be equally likely to report r. I formalize this result as Lemma 5.

Lemma 5. *$(m_t, \lambda_t)_{t=1,...,T}$ is a myopic equilibrium if and only if λ_t is consistent with m_t at every t, and at every report history $\hat{h}$, $q^B(\hat{h}) = q^G(\hat{h})$.*

There is clear intuition behind this result (a formal proof is relegated to the appendix). Let us first consider why $q^B(\hat{h}) = q^G(\hat{h})$ is necessary in equilibrium. One can see that a failure of this equality would guarantee a profitable deviation for the firm. Suppose for instance that at some $\hat{h}$, $q^G(\hat{h}) > q^B(\hat{h})$. Note that such a $q^B(\hat{h})$ is consistent with a strategy in which the bad firm reports l with positive probability under report history $\hat{h}$. Furthermore, good firms in this case

[3] While q^B is in general a function of the bad firm's strategy, I supress this dependence for simplicity.

on average report r more than bad firms. The consumer's beliefs are consistent with Bayes Rule, and would thus reflect this fact: an r report would guarantee an immediate improvement in the firm's reputation in period t, whereas an l report guarantees a decline in its reputation. Since a myopic firm is exclusively concerned with its current reputation, it could then profitably deviate by exclusively reporting r in the first period.

Next, let us consider why having $q^B(\hat{h}) = q^G(\hat{h})$ at every report history $\hat{h}$ is sufficient for an equilibrium to hold (provided that the λ_t are consistent with Bayes Rule). In this case, at every history, good and bad firms produce identical distibutions over reports (unconditional of the state). Because the consumer is assessing reputation based solely on the firm's reports, it must assign the same reputation to a firm that reports r in the first period as one that reports l. Thus the firm must be indifferent between reporting r and l, ensuring that any m_t is optimal, including the equilibrium strategy.

Lemma 5 provides a condition on the m_t to ensure that they are comprise an equilibrium. However, it is easy to see that the m_t are not unique. This can be seen by examining equations (3.1) and (3.2): there exist a continuum of m_t that will ensure $q_t^B(\hat{h}) = q_t^G(\hat{h})$. There is also a clear intuition behind this. Lemma 5 only pins down the unconditional distribution of reports for the bad type. Namely, it places no restriction on the distribution of reports conditional on the firm's signal, i.e., how the firm chooses to corrolate its reports with its private signals.

Thus, I introduce a second selection criterion. Namely, I restrict attention to the *most informative* of these equilibria. This is the m_t where the bad firm garbles its signals as little as possible at every history.

Criterion 2. *An equilibrium $(m_t, \lambda_t)_{t=1,\ldots,T}$ is a most informative equilibrium if for each t and report history $\hat{h}_t = (\hat{s}_1, \ldots, \hat{s}_{t-1})$,*

$$m_t^B(h_t, \hat{h}_t) = \begin{cases} \arg\max_{m_t \in \mathcal{M}_t(\hat{h})} m_t(s_t, \hat{h}) & \text{if } s_t = r \\ \arg\min_{m_t \in \mathcal{M}_t(\hat{h})} m_t(s_t, \hat{h}) & \text{if } s_t = l, \end{cases}$$

where $\mathcal{M}_t(\hat{h})$ is the set of functions $m_t(\cdot, \hat{h})$ such that there exists an equilibrium in which $m_t(\cdot, \hat{h})$ is played at $\hat{h}$.

Note the restrictiveness of this selection assumption: in order for an equilibrium strategy to be the most informative, it must minimize how much the firm garbles either of its signals at every report history. As I will illustrate below, an equilibrium that satisfies this criterion exists despite its restrictiveness. This characterization is presented formally as Proposition 13.

Proposition 13. *There exists a unique most informative equilibrium $(m_t^*, \lambda_t^*)_{t=1,...,T}$. Under this equilibrium, for any time-t report history $\hat{h}_t = \{\hat{s}_1, ..., \hat{s}_{t-1}\}$, if $s_t = 0$*

$$m_t^*(s_t, \hat{h}_t) = \begin{cases} 0 & \text{if } p(\hat{h}_t) \le \frac{1}{2} \\ b(\hat{h}_t) & \text{if } p(\hat{h}_t) > \frac{1}{2} \end{cases}$$

and if $s_t = 1$

$$m_t^*(s_t, \hat{h}_t) = \begin{cases} 1 - b(\hat{h}_t) & \text{if } p(\hat{h}_t) < \frac{1}{2} \\ 1 & \text{if } p(\hat{h}_t) \ge \frac{1}{2}, \end{cases}$$

where $b(\hat{h}_t) \equiv \frac{[p_t^G(\hat{h})\pi_G + (1-p_t^G(\hat{h}))(1-\pi_G)] - [p_t^B(\hat{h})\pi_B + (1-p_t^B(\hat{h})(1-\pi_B)]}{\pi_B(1-p^B(\hat{h}_t)) + (1-\pi_B)p^B(\hat{h}_t)}$.

Proposition 13 states that under a myopic equilibrium, firms bias their reports in each period in a similar fashion as in GS's static setting. Specifically, if the consumer holds a right-biased belief at a particular history ($p(\hat{h}_t) > \frac{1}{2}$) the firm will engage in right-biased reporting: it will truthfully report its r signals but if it observes an l signal, will instead report r with some positive probability $b(\hat{h}_t)$. Likewise, if the consumer is left-biased in that period ($p(\hat{h}_t) < \frac{1}{2}$), the firm will truthfully report l signals, but garble r signals with some positive probability. Despite the similarity to GS's static characterization, this equilibrium exhibits dynamics: the direction and magnitude of the firm's biasing ($b(\hat{h}_t)$) depends crucially on the report history. In the section 4, we will explore these dynamics further.

Before doing so, it is worth noting that while the firm's strategy exhibits dynamics in equilibrium, the same is not true of its reputation. In fact, the consumer learns nothing about the firm's type in equilibrium, i.e., reputation remains constant. I state this result formally as Corollary 1.

Corollary 1. *In a myopic equilibrium, the reputation function is constant. Specifically, for all t and $\hat{h}_t \in \hat{\mathcal{H}}_t$, $\lambda_t(\hat{h}_t) = \lambda_0$.*

This is an immediate corollary of Proposition 13: the firm garbles its signals in equilibrium in precisely such a way that either message yields an equal reputation. There is an intuition behind this: the consumer assesses reputation by observing the firm's signals and nothing else. Thus, if different reports yielded the firm different reputations, the firm could profitably deviate by exclusively reporting the reputation-maximizing message, because it would incur no cost from doing so. As I will illustrate later, this result will fail to hold if the firm faces lying costs.

3.4 Economic implications

In this section I consider the economic implications of this equilibrium as they relate to both dynamics and comparative statics of the firm's strategy.

3.4.1 Dynamics

While I show above that the firm biases its reports in every period in a similar fashion to GS's static model, both the magnitude and direction of this bias may change in over time. Here, I explore the nature of these dynamics. Specifically, reporting bias in a dynamic setting does not only consist of confirming the firm's prior as is the case in the static setting (*prior bias*). The firm also engages in *consistency bias*, which is dynamic in nature: as time passes and the firm accumulates a richer history of reports, the firm will bias its reports in such a manner to appear consistent with this report history. Furthermore, I illustrate below that consistency bias will gradually overtake prior bias as time passes. The remainder of this subsection will be dedicated to formalizing these concepts.

Recall that in the previous section, $p^G(\hat{h}_t)$ was defined as the consumer's posterior around the state, given that $i = G$ and having observed history $\hat{h}_t$. Recall that this value determines $q^G(\hat{h}_t)$, thus determines $b_t(\hat{h})$. Now for any time-t history $\hat{h}_t$, let $n(\hat{h}_t)$ determine the number of r elements in $\hat{h}_t$ (i.e., how many times the firm has reported r). It follows from Bayes Rule that

$$p^G(\hat{h}_t) = \frac{1}{1 + C(\hat{h}_t)P}$$

where

$$C(\hat{h}_t) \equiv \left(\frac{1-\pi_G}{\pi_G}\right)^{2n(\hat{h}_t)-t+1}$$

and

$$P \equiv \frac{1-p_0}{p_0}.$$

Thus, we see that $p^G(\hat{h}_t$ is decreasing function of two terms: P and $C(\hat{h}_t)$. Note that $P \in [0, \infty)$ and is decreasing in p_0. This is the inverse prior likelihood ratio, and I will refer to it as the *prior bias*. As its name suggests, this measures the extent to which the consumer's prior leans leftward. Notably, this value is static: it only depends on the consumer's prior belief regarding the state.

Let's now consider $C(\hat{h}_t)$. Like P, $C(\hat{h}_t) \in [0, \infty)$. However, unlike P, C depends on the firm's report history. I will refer to this expression as *consistency bias*. It measures the extent to which the firm's reports, conditional on it being good, are consistent with L being the true state of the world. I will formalize this interpretation below.

Let us start by examining the exponent on the expression for $C(\hat{h}_t)$: $2n(\hat{h}_t) - t + 1$. I refer to this value as the *r-surplus*, as it measures the surplus of past $\hat{r}$ reports over $\hat{l}$ reports. In particular, a positive (negative) r-surplus indicates that there have been more (fewer) r reports than l reports. Thus, holding π_G constant, C must be strictly decreasing in the r-surplus. This is intuitive: the more r reports a firm has received as compared to l reports, the less consistent the firm's reports are with L being the state.

Now, let now consider how C depends on π_G, the quality of the good firm's signal. To illustrate this, let us first suppose that the r-surplus is positive. In this case, C is decreasing in π_G. Next, let us consider a $\hat{r}$-surplus which is negative. In this case, C is increasing in π_G. In other words, for a given set of reports, a higher π_G draws the consistency bias further away from $\frac{1}{2}$, making it more extreme. This is due to the fact that a higher π_G that the good firm has a more accurate signal. Thus, for a set of past reports that have been more left-leaning, they will be more consistent with L being the true state when π_G is higher, because the these reports are more likely to be driven by L being the true state, rather than random noise. To illustrate this point further, let us consider two extreme cases in turn: (1) $\pi_G = \frac{1}{2}$ and (2) $\pi_G = 1$.

First, suppose that $\pi_G = \frac{1}{2}$. This indicates that the good firm's signals are completely uninformative about the state. Thus, the consumer's belief around the state should be unaffected by any reports she has received. Indeed, we see that $C(\hat{h}_t) = 1$ for all $\mathcal{I}_t$. This means that all report histories will be equally consistent with L being the true state. Next, let's suppose that $\pi_G = 1$. This indicates that the good firm's signals are fully informative about the state, as assumed in GS. One would thus expect the consistency bias to be highly sensitive to the report history. Again, this hypothesis is confirmed. Consider C in the second period, after the firm has only made one prior report. Then the r-surplus is either -1 or 1, corresponding to an l or r report in the first period, respectively. In the first case, C achieves its lower bound, 0. That is, the firm's report, assuming that it is good, is consistent with R being the true state. In the second case, where an l report is made in the first period, C approaches its upper bound, infinity. In this case, the good firm's report is consistent with L being the true state. This would indicate that if $\pi_G = 1$, the equilibrium strategy of the bad firm would be to replicate its report in the the prior period. I.e., $\hat{s}_t = \hat{s}_{t-1}$ for all $t > 1$. Formally, is due to the fact that the consistency effect dictates the equilibrium in all periods after the first. Intuitively, this because contradicting one of its prior reports, the bad firm would be revealing its type to the consumer. This was precisely the motivation for allowing $\pi_G < 1$. Otherwise, the equilibrium in all periods beyond the first would be degenerate; the firm would perpetually repeat its first report, thus eliminating the dynamic aspect of the equilibrium.

3.4.2 Comparative Statics

Let us now consider how the firm's biasing behavior changes with the parameters of the model. Specifically, let us ask how $b(\hat{h})$ changes with the consumer's prior bias (p_0), the abilities of both the good and bad firm (π_B and π_G), as well as the r-surplus, i.e., the report history. I formalize these findings as Proposition 14.

Proposition 14. *Holding all other parameters fixed, the equilibrium bias $b(\hat{h})$ is*

- *strictly increasing in p_0.*
- *strictly decreasing π_B and strictly increasing in π_G if $p(\hat{h}) > \frac{1}{2}$.*
- *strictly increasing in π_B and strictly decreasing in π_G if $p(\hat{h}) < \frac{1}{2}$.*
- *strictly decreasing in the r-surplus $2n - t + 1$, where t denotes the length of the history $\hat{h}$, and n the number of r reports that were made under $\hat{h}$.*

A formal proof of Proposition 14 is presented in the appendix, which follows from the expression for $b(\hat{h})$ above. Here, we will discuss the intuition behind these results. Let us begin by considering first why bias must be increasing in the consumer's prior. Holding a report history $\hat{h}$ fixed, a consumer who holds a higher prior p_0 will hold a higher interim posterior $p(\hat{h})$. If the equilibrium strategy did not change in response to this higher $p(\hat{h})$, bad firm's report distribution ($q^B(\hat{h})$) would react less to this change in belief than the good firm's report distribution ($q^G(\hat{h})$). This is due to the fact $\pi_B < \pi_G$: the bad firm's signals are less corrolated with the state than the good firm's signals (in the extreme case where $\pi_B = 1$, they are entirely uncorrolated). In order to ensure that the equilibrium is preserved, the bad firm must adjust its strategy in order to preserve the equality of these distributions, and would have to do so by reporting r more often, i.e., increasing $b(\hat{h})$.

Similar intuition explains why $b(\hat{h})$ must increase in π_G and decrease in π_B when $p(\hat{h}) > \frac{1}{2}$. When $p(\hat{h}) > \frac{1}{2}$, the consumer's belief at that history is right-leaning. The bad firm must bias its reports to the right in equilibrium in order to compensate for the fact that it has a lower π_i, and thus all else equal would be less likely to make an r report. Increasing π_G will only exacerbate the difference in π_i between the two types, and thus in order to preserve indifference, the firm must introduce even more right-leaning bias in its reporting strategy (i.e., increase $b(\hat{h})$). Likewise, increasing π_B will reduce the gap between the good and bad firm's ability, and thus necessitate a lower bias $b(\hat{h})$ in order to preserve indifference. Economically, this comparative static sheds light on how variation in firm ability can affect bias in the market. Specifically, it states that one should expect more bias in a setting where firms differ greatly in their abilities (π_G and π_B are far apart) and less so when firms are more homogenous in ability.

The last comparative static relates not to the parameters of the model, but rather to the report history. It state thats the more r reports the bad firm has had relative to l reports, the more it must bias its reports in equilibrium. This is driven by the consistency effect discussed above: good firms' signals are more likely to be consistent with one another than bad firms' signals. Thus, in period t, to make itself indistinguishable from a good firm, a bad firm must bias its reports towards r in order to replicate the good firm's level of consistency.

3.5 Relaxing Myopia

In this section, I relax the assumption that the firm is myopic in equilibrium. Specifically, I argue that even if the firm is forward looking, there exists a unique equilibrium, and this equilibrium is identical to the myopic one. I.e., a forward-looking firm will behave no differently in equilibrium than a myopic firm. I formalize this result as Proposition 15, and provide a proof below.

In what follows, a *forward looking* firm refers to a firm with a payoff function that is not necessarily myopic. That is, at each t, the firm's payoff is given by $F_t(\lambda_t, ..., \lambda_T)$, where F_t is weakly increasing in the λ_s for all $s \geq t$, and strictly increasing in λ_s for some $s \geq t$.

Proposition 15. *There exists a unique forward-looking equilibrium. This equilibrium is the myopic equilibrium* $(m_t^*, \lambda_t^*)_{t=1,...,T}$.

Proof. First, I show that the myopic equilibrium $(m_t^*, \lambda_t^*)_{t=1,...,T}$ is indeed an equilibrium under any forward-looking payoff functions F_t. To show this, it suffices to show that at every t, m_t^* must be a best response given λ_t^*. To this end, let us recall from Corollary 1 that under $\lambda_t^*(\hat{h}) = \lambda_0$ under all $\hat{h}$. Thus, any reporting strategy m_t is trivially a best response, because all such strategies must yield the same reputation λ_0 in all future periods. Thus, the firm could never profitably deviate from m_t^*.

Next, I will show that the myopic equilibrium is the unique equilibrium under a forward-looking firm. To this end, suppose that there exists some forward-looking equilibrium $(m_t, \lambda_t)_{t=1,...,T}$. We wish to show that for t, $m_t = m_t^*$. I will proceed by backwards induction. In period T, the firm is trivially myopic: F_T can only be a function of λ_T. Thus, the firm's equilibrium strategy at time T must also be the myopic one, m_t^* Now, fix a period t, and assume by induction that $m_s = m_s^*$ for all $s > t$. By identical reasoning to the proof of Corollary 1 it follows that $\lambda_s = \lambda_t$ for all $s > t$. Thus,

$$F_t(\lambda_t, ..., \lambda_T) = G_t(\lambda_t),$$

where G_t is some strictly increasing function of λ_t. This is precisely a myopic payoff function, and thus the myopic equilibrium strategy m_t^* must hold in period t as well. □

The intuition behind this proof relies crucially on the fact that under myopia, the consumer learns nothing about the firm's type, i.e., its reputation in all future periods must equal its reputation today. Thus, as long as the myopic strategy is played in future periods, even a forward-looking firm will operate as if it is only maximizing its current reputation.

This result is of interest for two reasons. First, it provides us with a characterization of the equilibrium even when the firm is myopic. Second, it sheds light more generally on the restrictiveness of the myopia assumption. While this assumption made in dynamic reputation games to simplify otherwise intractible problems (e.g., in [32]), here we have identified a setting in which this assumption is without loss.

However, it is notable that this result is not necessarily robust to a number of potential changes in the model. The proof relied crucially on the fact that the myopic equilibrium is pooling, i.e. that the firm's report reveals nothing about its type (Corollary 1). In a setting where there is even partial separation of types (i.e., in which the firm's reputation does react to its reporting behavior), this equivalence need not hold. In what follows, we will consider such a setting. In particular, I assume that lying is not costless for the firm, and that it incurs some non-zero cost from misreporting its signals. As I will illustrate, introducing lying costs will give rise to such a partially separating equilibrium, and the equivalence between the forward-looking and myopic equilibrium fails.

3.6 Incorporating Lying Costs

Suppose that the bad firm's payoff incorporates a lying cost. In particular, suppose that at every t, the firm's payoff is given by

$$F_t(\lambda_t, ..., \lambda_T) - \tau \mathbb{I}(\hat{s}_t \neq s_t), \tag{3.3}$$

where F_t is the payoff function assumed in the baseline model, and $\tau \geq 0$ denotes the cost of error.

Before proceeding with this augmented model, let us discuss the motivation behind introducing lying costs. This motivation is two fold. First, it may seem unrealistic that the firm faces no consequences from misreporting its signals. In fact, this lack of such consequences for lying drives the equilibrium multiplicity in the baseline model in the absence of selection, where completely uninformative reporting strategies are viable equilibria. If a firm faces a lying cost, even an infinitesimally small one, it will never lie out of indifference. As I will show below, this will grant us equilibrium uniqueness. Furthermore, such a payoff function allows us see how the dual objectives of reputation maximization and truthfulness interact with each other. Of course, a lying cost is not the only means of achieving these objectives. For instance, GS induce truthfulness by incorporating a feedback probability, in which there is some non-zero probability that the consumer learns the state, and uses this information to assess the firm's type. I restrict attention to lying costs for the tractibility it grants in a dynamic setting. The second motivation for incorporating a lying cost is that it will allow for separation in equilibrium. In particular, the consumer will engage in some learning about the firm's type. As discussed earlier, considering a setting with separation will shed light on the robustness of the earlier result establishing equivalence between a myopic and forward-looking equilibrium. As discussed earlier, PS assume myopia due to the intractibility of the forward-looking problem. Incorporating lying costs into this simple binary model can serve as a first step to understanding in what way such assumptions may be with loss, and how concerns regarding long-term reputation may alter a firm's behavior in equilibrium.

3.6.1 Equilibrium Under Myopia

As with the baseline model, I begin by characterizing the firm's behavior under myopia. Before proceeding with the equilibrium characterization, will reconsider the need for selection criteria in this setting. I will continue to assume Criterion 1, i.e., restrict attention to equilibria where the good firm reports truthfully. However, one can illustrate that unlike the setting without lying costs, there is no need for additional criteria (e.g., Criterion 2) in order to obtain equilibrium uniqueness. In fact, one can show that under lying costs, whenever the bad firm is garbling its l signals, it can never misreport its r signals. Likewise, if it is garbling its r, it can never misreport its l signals. This claim is formalized as Lemma 6.

Lemma 6. *Suppose a myopic firm faces payoff function* (3.3) *and* $\tau > 0$*. Then, for any* $\hat{h} \in \hat{\mathcal{H}}_t$*, if* $m_t(l, \hat{h}) > 0$*, then* $m_t(r, \hat{h}) = 1$*. Likewise, if* $m_t(l, \hat{h}) < 1$*, then* $m_t(r, \hat{h}) = 0$*.*

A formal proof of this claim is relegated to the appendix. The intuition is as follows: if a firm did at some time t misreport its r signals (i.e., report l despite receiving an r signal), it would then find it strictly optimal to truthfully report l signals. This is because reporting r (or l) will yield the firm the same reputation regardless of what the state is. However, while reporting l incurs a lying cost under signal r, it does not incur a lying cost under signal l. So, a firm that finds it weakly optimal to report l under an r signal, it must find it strictly optimal to report l under an l signal.

Let us now proceed with the equilibrium characterization. To fix ideas, let us first consider how the equilibrium under lying costs might differ from that under the baseline model. Intuitively, one would expect the lying cost τ to induce the bad firm to be more truthful in equilibrium than under the baseline equilibrium. Indeed, one can see that the baseline equilibrium does not in general hold whenever the lying cost is strictly positive ($\tau > 0$).

This can be illustrated by a brief formal argument. Fix any t and consider two different time $t+1$ report histories $\hat{h}_{t+1} \equiv \{\hat{s}_1, ..., \hat{s}_{t-1}, r\}$ and $\hat{h}'_{t+1} \equiv \{\hat{s}_1, ..., \hat{s}_{t-1}, l\}$. Note that these report histories are identical, except that under $\hat{h}_{t+1}$ an r report is made in period t and under $\hat{h}'_{t+1}$ an l report is made in period t. It follows from Corollary 1 that

$$\lambda_t(\hat{h}_{t+1}) = \lambda_t(\hat{h}'_{t+1}) = \lambda_0.$$

Now, suppose that $p(\hat{h}_t) > \frac{1}{2}$, where $\hat{h}_t = (s_1, ..., s_t)$. It then follows from Proposition 13 that $m_t(l, \hat{h}_t) > 0$. I.e., upon receiving an l signal, the firm mixes between sending messages l and r. Consider the bad firm's payoff when receiving signal l. Reporting r in period t would yield a payoff of $\lambda_0 - \tau$ whereas reporting l would yield a payoff of λ_0. Thus, the firm must find it strictly optimal to report l upon receiving signal l. This is a contradiction.

Some intuition underlies the formal argument above. The baseline equilibrium requires that when $p(\hat{h}_t) > \frac{1}{2}$ a firm who observes signal l indifferent between truthfully reporting its signal and lying (i.e., reporting r). However, holding all else equal and incorporating a lying cost, lying will yield the firm strictly worse off than being truthful, thus causing the indifference condition to fail.

While I have shown that the baseline equilibrium can never hold when $\tau > 0$, I have not yet characterized the equilibrium under lying costs. Before presenting a formal characterization, let us discuss the intuition behind this characterization. Let us suppose without loss that under some report history $\hat{h}$, $p(\hat{h}) > \frac{1}{2}$. Intuitively, one should expect that the firm still biases its reports as long as τ is sufficiently small. This is because if the firm were truthful, sending an r report would be reputationally beneficial for the firm even under singal l, and as long as the cost of lying is small enough, would be the strictly profitable response for the firm. However, the presence of the lying cost would mean that the two reports could not yield the firm identical reputations: the firm's reputation from lying must be strictly higher in order to compensate it for the cost from lying. Meanwhile, one would intuitively expect that that if τ is sufficiently large, the firm would never lie. In this case, the reputational gain from lying could never be large enough to compensate the firm for the cost from doing so.

In order to describe the equilibrium characterization, let is begin by introducing some notation.

First, fix a t and $(m_s, \lambda_s)_{s=1,...,t-1}$. Then, let $\bar{m}_t$ denote the fully truthful time-t strategy, i.e., for any $\hat{h} \in \hat{H}_t$,

$$\bar{m}_t(r, \hat{h}) = 1, \ \bar{m}_t(l) = 0.$$

Further, let $\bar{\lambda}_t^{(}\hat{h}_{t+1})$ denote a Bayesian consumer's belief about the firm's type at time t under truthful reporting under history $\hat{h}_t$, assuming that report r was made in period t, and given $(\lambda_s)_{s=1,...,t-1}$.[4]

Proposition 16. *Suppose a myopic firm faces payoff function* (3.3) *in each period, where* $\tau > 0$. *There exists a unique equilibrium* $(m_t^*, \lambda_t^*)_{t=1,...,T}$ *that satisfies Criterion 1. This equilibrium can be constructed recursively as follows. For any t, fix any* $\hat{h}_t = (s_1, ..., s_{t-1})$*, and let* $\hat{h}_{t+1} = (s_1, ..., s_{t-1}, r)$ *and* $\hat{h}'_{t+1} = (s_1, ..., s_{t-1}, l)$*. Then*

(a) If $\tau \geq |F_t(\bar{\lambda}_t(\hat{h}_{t+1})) - F_t(\bar{\lambda}_t(\hat{h}'_{t+1}))|$*, then the firm is truthful, i.e.:*

$$(m_t^*, \lambda_t^*) = (\bar{m}_t, \bar{\lambda}_t).$$

(b) If $\tau < |F_t(\bar{\lambda}_t(\hat{h}_{t+1})) - F_t(\bar{\lambda}_t(\hat{h}'_{t+1})))|$*, then*

$$\tau = |F_t(\lambda_t^*(\hat{h}_{t+1})) - F_t(\lambda_t^*(\hat{h}'_{t+1}))|. \tag{3.4}$$

The proof of this claim is included in the appendix. This result is of interest for a few reasons. First, unlike the baseline model, we are guaranteed that even if the lying cost is infinitesimally small, the equilibrium must be unique even without imposing Criterion 2. This uniqueness results directly from Lemma 6. Second, this claim establishes that with a sufficiently high lying cost, the bad firm will be strictly truthful. Finally, as shown by part (b), if τ is small enough, the firm does still bias its reports, though it does so at a lower rate than in the baseline equilibrium. Specifically, the firm faces a tradeoff: either it truthfully reports its signal at a cost to its reputation, or lies and incurs a lying cost. In equilibrium, the firm is exactly indifferent between being truthful and lying.

[4]Note that $\bar{\lambda}_t$ is a function of $(\lambda_s)_{s=1,...,t-1}$. However, we drop this dependence for brevity.

3.6.2 Relaxing Myopia

Having established the equlibrium under myopia, as with the baseline model, I now consider what the equilibrium may look like in the presence of lying costs. I show that we do not achieve equivalence between myopic and forward-looking equilibria under lying costs.

Proposition 17. *The myopic equilibrium is not necessarily also a forward-looking equilibrium.*

The formal proof of this claim is included in the appendix, in which I construct an example where the myopic equilibrium fails to also serve as a forward-looking equilibrium. There is a clear reason why the equivalence proof of the baseline model (i.e., the proof of Proposition 15) falls through when one introduces lying costs. As illustrated by our above characterization, under a myopic equilibrium of the model with lying costs, the firm's reputation does not react to its reporting behavior. This ensured that even if it is forward looking, a firm could not profitably deviate from the myopic equilibrium strategy. Meanwhile, as demonstrated by Proposition 16, under lying costs, the firm's reputation does respond to its reporting behavior.

3.7 Conclusion

In this paper, I extended the GS model of reputation-driven media bias to a dynamic setting to understand how the nature of reputation-driven media bias may change over time. I have shown that, under a myopic firm, the equilibrium involves bias by the bad firm that is driven by both an appeal to the consumer's prior regarding the state and an effort to appear consistent in its reports. I was able to capitalize on the pooling nature of this equilibrium to establish an equivalence between forward-looking and myopic equilibria. That is, I show that assuming a firm is myopic is without loss. Next, I considered an extension of this baseline model where in addition to seeking a favorable reputation, firms incur a cost from lying. I showed that the myopic equilibrium in this setting is unique, and involves strictly less bias in equilibrium than in the model without lying costs. Finally, I demonstrated that the myopic and forward-looking equilibrium need not equal each other in this extended model, a result is driven by the partially separating nature of the equilibrium. While

I have shown that the forward-looking and myopic equilibrium need not equal each other under lying costs, I have not provided a full characterization of the forward-looking equilibrium. This is the subject of future work.

References

[1] D. Fudenberg, R. Gilbert, J. Stiglitz, and J. Tirole, "Preemption, Leapfrogging and Competition In Patent Races," *European Economic Review*, vol. 22, no. 1, pp. 3–31, 1983.

[2] D. Fudenberg and J. Tirole, "Preemption and Rent Equalization In the Adoption of New Technology," *The Review of Economic Studies*, vol. 52, no. 3, pp. 383–401, 1985.

[3] S. R. Grenadier, "The Strategic Exercise of Options: Development Cascades and Overbuilding in Real Estate Markets," *The Journal of Finance*, vol. 51, no. 5, pp. 1653–1679, 1996.

[4] D. Abreu and M. K. Brunnermeier, "Bubbles and Crashes," *Econometrica*, vol. 71, no. 1, pp. 173–204, 2003.

[5] H. A. Hopenhayn and F. Squintani, "Preemption Games with Private Information," *The Review of Economic Studies*, vol. 78, no. 2, pp. 667–692, 2011.

[6] C. Bobtcheff, J. Bolte, and T. Mariotti, "The Researcher's Dilemma," *The Review of Economic Studies*, vol. 84, no. 3, pp. 969–1014, 2017.

[7] C. Bobtcheff, R. Levy, and T. Mariotti, "Negative results in science: Blessing or (winner's) curse?," 2022.

[8] F. Gul and R. Lundholm, "Endogenous Timing and the Clustering of Agents' Decisions," *Journal of Political Economy*, vol. 103, no. 5, pp. 1039–1066, 1995.

[9] C. Chamley and D. Gale, "Information Revelation and Strategic Delay in a Model of Investment," *Econometrica: Journal of the Econometric Society*, pp. 1065–1085, 1994.

[10] D. Levin and J. Peck, "Investment Dynamics with Common and Private Values," *Journal of Economic Theory*, vol. 143, no. 1, pp. 114–139, 2008.

[11] C. Lin, "Speed Versus Accuracy in News," 2014.

[12] A. Pant and F. Trombetta, "The Newsroom Dilemma," *Available at SSRN 3447908*, 2019.

[13] G. Andreottola, A.-I. de Moragas, *et al.*, *Scandals, Media Competition and Political Accountability*. CSEF, Centre for Studies in Economics and Finance, Department of Economics ..., 2020.

[14] M. Gentzkow and J. M. Shapiro, "Competition and Truth in the Market for News," *Journal of Economic Perspectives*, vol. 22, no. 2, pp. 133–154, 2008.

[15] A. Liang, X. Mu, and V. Syrgkanis, “Dynamically Aggregating Diverse Information,” 2021.

[16] S. Galperti and I. Trevino, “Coordination Motives and Competition for Attention in Information Markets,” *Journal of Economic Theory*, vol. 188, p. 105 039, 2020.

[17] H. Chen and W. Suen, “Competition for Attention and News Quality,” Working paper, Tech. Rep., 2019.

[18] J. Perego and S. Yuksel, “Media Competition and Social Disagreement,” *Working paper*, 2018.

[19] V. P. Crawford and J. Sobel, “Strategic Information Transmission,” *Econometrica*, pp. 1431–1451, 1982.

[20] E. Kamenica and M. Gentzkow, “Bayesian Persuasion,” *American Economic Review*, vol. 101, no. 6, pp. 2590–2615, 2011.

[21] S. Mullainathan and A. Shleifer, “The Market for News,” *American Economic Review*, vol. 95, no. 4, pp. 1031–1053, 2005.

[22] M. Gentzkow and J. M. Shapiro, “Media Bias and Reputation,” *Journal of Political Economy*, vol. 114, no. 2, pp. 280–316, 2006.

[23] R. Boleslavsky and C. R. Taylor, “Make It ’Til You Fake It,” *Economic Research Initiatives at Duke (ERID) Working Paper Forthcoming*, 2020.

[24] Y.-K. Che and J. Hörner, “Recommender Systems as Mechanisms for Social Learning,” *The Quarterly Journal of Economics*, vol. 133, no. 2, pp. 871–925, 2018.

[25] R. Laraki, E. Solan, and N. Vieille, “Continuous-Time Games of Timing,” *Journal of Economic Theory*, vol. 120, no. 2, pp. 206–238, 2005.

[26] K. Hendricks, A. Weiss, and C. Wilson, “The War of Attrition in Continuous Time with Complete Information,” *International Economic Review*, pp. 663–680, 1988.

[27] M. R. Baye, D. Kovenock, and C. G. De Vries, “The All-Pay Auction with Complete Information,” *Economic Theory*, vol. 8, no. 2, pp. 291–305, 1996.

[28] M. Ottaviani and P. N. Sørensen, “Reputational cheap talk,” *The Rand journal of economics*, vol. 37, no. 1, pp. 155–175, 2006.

[29] M. Ottaviani and P. N. Sørensen, “The strategy of professional forecasting,” *Journal of Financial Economics*, vol. 81, no. 2, pp. 441–466, 2006.

[30] M. Ottaviani and P. N. Sørensen, "Professional advice," *Journal of Economic Theory*, vol. 126, no. 1, pp. 120–142, 2006.

[31] A. Dasgupta and A. Prat, "Information aggregation in financial markets with career concerns," *Journal of Economic Theory*, vol. 143, no. 1, pp. 83–113, 2008.

[32] C. Prendergast and L. Stole, "Impetuous youngsters and jaded old-timers: Acquiring a reputation for learning," *Journal of political Economy*, vol. 104, no. 6, pp. 1105–1134, 1996.

[33] D. S. Scharfstein and J. C. Stein, "Herd Behavior and Investment," *American Economic Review*, pp. 465–479, 1990.

[34] I. Guttman, I. Kremer, and A. Skrzypacz, "Not only what but also when: A theory of dynamic voluntary disclosure," *American Economic Review*, vol. 104, no. 8, pp. 2400–2420, 2014.

[35] G. Gratton, R. Holden, and A. Kolotilin, "When to drop a bombshell," *The Review of Economic Studies*, vol. 85, no. 4, pp. 2139–2172, 2018.

[36] C. Bobtcheff and R. Levy, "More haste, less speed? signaling through investment timing," *American Economic Journal: Microeconomics*, vol. 9, no. 3, pp. 148–86, 2017.

[37] M. Halac and I. Kremer, "Experimenting with career concerns," *American Economic Journal: Microeconomics*, vol. 12, no. 1, pp. 260–88, 2020.

[38] A. Smirnov and E. Starkov, "Timing of predictions in dynamic cheap talk: Experts vs. quacks," *University of Zurich, Department of Economics, Working Paper*, no. 334, 2019.

[39] A. Wald, "Foundations of a general theory of sequential decision functions," *Econometrica, Journal of the Econometric Society*, pp. 279–313, 1947.

[40] M. Gentzkow, J. M. Shapiro, and D. F. Stone, "Media bias in the marketplace: Theory," in *Handbook of media economics*, vol. 1, Elsevier, 2015, pp. 623–645.

[41] D. Strömberg, "Mass media competition, political competition, and public policy," *The Review of Economic Studies*, vol. 71, no. 1, pp. 265–284, 2004.

[42] W. Suen, "The self-perpetuation of biased beliefs," *The Economic Journal*, vol. 114, no. 495, pp. 377–396, 2004.

[43] J. Chan and W. Suen, "Media as watchdogs: The role of news media in electoral competition," *European Economic Review*, vol. 53, no. 7, pp. 799–814, 2009.

[44] N. Schulz and J. Weimann, "Competition of newspapers and the location of political parties," *Public Choice*, vol. 63, no. 2, pp. 125–147, 1989.

[45] C. Prendergast, “A theory of" yes men",” *The American Economic Review*, pp. 757–770, 1993.

[46] B. Holmström, “Managerial incentive problems: A dynamic perspective,” *The review of Economic studies*, vol. 66, no. 1, pp. 169–182, 1999.

[47] R. Deb, M. M. Pai, and M. Said, “Evaluating strategic forecasters,” *American Economic Review*, vol. 108, no. 10, pp. 3057–3103, 2018.

Appendix A: Omitted Proofs and Formalities from Chapter 1

A.1 Beliefs in equilibrium

In this section, I will present certain properties and notation regarding the players' beliefs. These will be relevant to the analysis that follows.

First, I remark that at all times t and histories H, all players, with the exception of those who have already reported, must hold a common belief about the state. I omit a formal proof as this follows directly from the selection assumption (SC). This assumption implies that it is common knowledge that all firms who have not yet reported have not observed a conclusive signal. Thus, all such firms, as well as the consumer, share the same information set, and thus a common belief about the state.

Next, fixing an initial common belief p, and number of remaining firms n, let us define two conditional beliefs: $p(s)$ and $p^i(s)$. First, let $p(s)$ denote the common updated belief, conditional on no new reports being made after s time passes. It follows from Bayes Rule that

$$p(s) = \frac{pe^{-n\lambda s}}{pe^{-n\lambda s} + (1-p)} \tag{A.1}$$

Next, let $p^i(s)$ denote the common belief, conditional on the event that player i reported after s time has passed, and no other reports were made. Again, $p^i(s)$ follows directly from Bayes Rule, given belief α:

$$p^i(s) = \alpha_n(p(s)) + (1 - \alpha_n(p(s))p(s) \tag{A.2}$$

To understand how $p^i(s)$ is computed, note that if a report is made after time s has passed, conditioning on the event that i's report was informed, the common belief will update to 1. However, conditioning on the event that i was uninformed when making the report, the report would have no impact on the common belief, which would thus be given by $p(s)$. Thus, $p^i(s)$ is given by the weighted sum of these two beliefs, where the weighting is specified by the belief that the report was informed, i.e., $\alpha_n(p(s))$.

Next, I formally justify equation (1.3) by showing that it is the limit of Bayes-consistent beliefs under a discretized version of the game presented in Section 1.2. To this end, for any $\epsilon > 0$, let the *ϵ-approximation* of the game be identical to the game presented in section (1.2), except with the following modification: any report made by a firm on $[0, \epsilon]$ is observed by all other players (including the consumer) at ϵ. Formally, rather than observing t_i, the players observe $\tilde{t}_i$, where

$$\tilde{t}_i \equiv max\{t_i, \epsilon\}$$

At any (p, n) that is on-path, let $\alpha_n^\epsilon(p)$ denote the firm's credibility, i.e., the consumer's belief that $s_i \leq \epsilon$ given that $\tilde{t}_i = \epsilon$, under the ϵ approximation of the game. Let α_n^ϵ denote the right-limit of the α_ϵ. Formally:

$$\alpha_n(p) \equiv \lim_{\epsilon \to 0+} \alpha_n^\epsilon(p)$$

I now establish that $\alpha_n(p)$ is given by (1.3) at any (p, n) on-path.

Claim 2. *For any (p, n) on-path,*

$$\alpha_n(p) = \begin{cases} \frac{\lambda p}{\lambda p + F'_{p,n}(0+)} & \text{if } F_{p,n}(0) = 0 \\ 0 & \text{if } F_{p,n}(0) > 0 \end{cases}$$

Proof. For any $\epsilon > 0$, $\alpha_n^\epsilon(p)$ is uniquely determined by Bayes Rule and given by

$$\alpha_n^\epsilon(p) = \frac{p(1 - e^{-\lambda\epsilon})}{p(1 - e^{-\lambda\epsilon}) + F_{p,n}(\epsilon)e^{-\lambda\epsilon}}.$$

First, consider the case where $F_{p,n}(0) = 0$. In this case, it follows from L'Hôpital's Rule that:

$$\lim_{\epsilon\to 0+} \alpha_n^\epsilon(p) = \frac{\lambda p}{\lambda p + F'_{p,n}(0+)}$$

Next, consider the case where $F_{p,n}(0) > 0$. In this case, we obtain

$$\lim_{\epsilon\to 0+} \alpha_n^\epsilon(p) = \frac{0}{0 + \lim_{\epsilon\to 0+} F_{p,n}(\epsilon)} = 0$$

where the final equality follows from the fact that $\lim_{\epsilon\to 0+} F_{p,n}(\epsilon) = F_{p,n}(0) > 0$. □

A.2 Equilibrium characterization

A.2.1 The firm's problem

Here, I formally define the firm's problem. Before proceeding, I define a useful object, the *first report distribution* Ψ. Formally, fixing a (p, n), let us index the players who have not yet reported by $i \in \{1, ..., n\}$. Then, $\Psi^i(s)$ denotes the probability that player i reported at or before s and was not preceded by any of the remaining firms in doing so. Fixing a strategy profile $(F^1_{p,n}, ..., F^n_{p,n})$, this is given by:

$$\Psi^i(s) = p\int_0^s e^{-\lambda r(N-n)}\prod_{j\neq i}(1-F^j_{p,n}(r))d(e^{-\lambda r}(F^i_{p,n}(r)-1))+(1-p)\int_0^s\prod_{j\neq i}(1-F^j_{p,n}(r))dF^i_{p,n}(r)$$

The first integral of the expression denotes the probability that i reports before s time has passed, and was the first of the remaining firms to do so, conditional on $\theta = 1$. Meanwhhile the second integral denotes the same probability conditional on $\theta = 0$. $\Psi^i(s)$ is then the weighted sum of these two probabilities, where the weight is given by the common belief p about θ. Note that while Ψ is a function of F, p, and n, I omit this dependences for brevity.

The firm's problem is defined recursively as follows. Fix an n, p, i, α, and continuation value $V_{p,n+1}$ for each p. Trivially, $V_{p,0} = 0$ for all p. Assume all firms $j \neq i$ play the same strategy F, and let $-i$ refer to a generic $j \neq i$. Then i's expected payoff from playing strategy F^i at (p, n) is given by:

$$V_{p,n}(F^i) = \int_0^\infty [k_n \alpha_n(p(s)) - \beta(1 - p^i(s))] d\Psi^i(s) + (N - n) \int_0^\infty V_{p^{-i}(s),n+1} d\Psi^{-i}(s) \quad \text{(A.3)}$$

Note that first integral of this expression is firm i's expected payoff from reporting, when it is the first of the remaining firms to do so. Meanwhile, the second integral is the firm's expected payoff conditional on being preempted. The firm's problem at (p, n) is then given by the following:

$$\max_{F^i \in \mathcal{F}} V_{p,n}(F^i),$$

where $\mathcal{F}$ denotes the set of permissible distributions, i.e., those that are piecewise continuously differenetiable, right-differentiable, and that satisfy the selection criterion (SC). Further, let $V_{p,n} \equiv \sup_{F^i \in \mathcal{F}} V_{p,n}(F^i)$.

A.2.2 Proofs

Proof of Lemma 1. Let us begin by showing that at all (p, n) on-path such that $p < 1$, $F_{p,n}$ is continuous at 0. To this end, suppose by contradiction that $F_{p,n}$ is discontinuous at 0. By the right-continuity of $F_{p,n}$, this implies that $F_{p,n}(0) > 0$. Because (p, n) is on path, by (1.3), $\alpha_n(p) = 0$. Furthermore, it follows by (A.2) that $p^i(0) = p$. Recalling that we are restricting attention to symmetric equilibria, let Ψ denote the first-report distribution at (p, n) under the equilibrium strategy profile $F_{p,n}$. Because $F_{p,n}(0) > 0$, $\Psi^j(0) > 0$ for all j who have not yet reported.

Now define the following deviation $\hat{F}_{p,n}$. This strategy is identical to $F_{p,n}$, except that all the mass that $F_{p,n}$ places on 0 is shifted to ∞:

$$\hat{F}_{p,n}(s) = \begin{cases} F_{p,n}(s) - F_{p,n}(0) & \text{if } s < \infty \\ 1 & \text{if } s = \infty \end{cases}$$

Now, fix some i who has not yet reported. Let $\hat{\Psi}$ denote the first-report distribution at (p, n) under the strategy profile where i plays $\hat{F}_{p,n}$ and all $j \neq i$ play $F_{p,n}$. By definition, for all $s \geq 0$,

$$\hat{\Psi}^i(s) = \Psi^i(s) - \Psi^i(0).$$

Then,

$$\int_0^\infty [k_n\alpha_n(p(s)) - \beta(1-p^i(s))]d\hat{\Psi}^i(s) = \int_0^\infty [k_n\alpha_n(p(s)) - \beta(1-p^i(s))]d\Psi^i(s) + \beta(1-p^i(0))\Psi^i(0)$$
$$> \int_0^\infty [k_n\alpha_n(p(s)) - \beta(1-p^i(s))]d\Psi^i(s).$$

Again by definition, for all $s \geq 0$,

$$\hat{\Psi}^{-i}(s) = \Psi^{-i}(s) + X(s),$$

where

$$X(s) \equiv \Psi^i(0)[p\int_0^s (1-F_{p,n})^{n-2}(1-\hat{F}_{p,n}(r))e^{-\lambda r}d(e^{-\lambda r}(F_{p,n}(r)-1))$$
$$+(1-p)\int_0^s (1-F_{p,n}(r))^{n-2}(1-\hat{F}_{p,n}(r))dF_{p,n}(r)]$$

Then, we have

$$\int_0^\infty V_{p^{-i}(s),n+1}d\hat{\Psi}^{-i}(s) - \int_0^\infty V_{p^{-i}(s),n+1}d\Psi^{-i}(s) = \int_0^\infty V_{p^{-i}(s),n+1}dX(s) \geq 0.$$

where the final inequality follows from the fact that $X(s)$ is increasing in s and $V_{p^{-i}(s),n+1} \geq V_{p^{-i}(s),n+1}(\delta_\infty) \geq 0$.

Combining the above two inequalities we have

$$V_{p,n}(\hat{F}_{p,n}) = \int_0^\infty [k_n\alpha_n(p(s)) - \beta(1-p^i(s))]d\hat{\Psi}^i(s) + (N-n)\int_0^\infty V_{p^{-i}(s),n+1}d\hat{\Psi}^{-i}(s)$$
$$> \int_0^\infty [k_n\alpha_n(p(s)) - \beta(1-p^i(s))]d\Psi^i(s) + (N-n)\int_0^\infty V_{p^{-i}(s),n+1}d\Psi^{-i}(s) = V_{p,n}(F_{p,n})$$

Thus, i can profitably deviate at (p, n). Contradiction.

We will now show that for all (p, n) on-path such that $p < 1$, $F_{p,n}$ must be continuous at all t. Suppose by contradiction that it is not. Let t denote the time at which a discontinuity occurs. Because $F_{p,n}$ is increasing and right-differentiable by assumption, this must be a jump discontinuity, i.e.,

$$\lim_{r \to t-} F_{p,n}(r) < F_{p,n}(t)$$

By (1.2),

$$F_{p(t),n}(0) = \frac{F_{p,n}(t) - \lim_{r \uparrow t} F_{p,n}(r)}{1 - \lim_{r \uparrow t} F_{p,n}(r)} > 0.$$

But then, this implies that $F_{p(t),n}$ is discontinuous at 0, contradicting the above.

Part (b) of the statement follows directly from (1.3). □

Lemma 7. *For any (p, n) on-path,*

- $\alpha_n(p) \geq \bar{\alpha}_n(p) \equiv \min\{\beta(1 - p)/k_n, 1\}$
- $F'_{p,n}(0+) \leq \bar{f} \equiv \lambda p(\frac{1}{\bar{\alpha}_n(p)} - 1)$

Proof of Lemma 7. We begin by showing the first point above. The second point follows by definition of $\alpha_n(p)$.

First, suppose by contradiction that there exists a (p, n) on-path such that

$$\alpha_n(p) < \min\{\beta(1 - p)/k_n, 1\}$$

Recalling that $p(s)$ is given by (A.1), we begin by claiming that for all s sufficiently small, $(p(s), n)$ is on-path. Suppose not by contradiction. Since (p, n) is on-path by assumption, this implies that $F_{p,n}(s) = 1$, which contradicts Lemma 1. It thus follows from (1.3), combined with the piecewise twice differentiability and right-differentiability of $F_{p,n}$, that $\alpha_n(p(s))$ is continuous in some right-neighborhood of $s = 0$. Formally, there exists an $\epsilon > 0$ such that for all $s \in [0, \epsilon]$,

$$k_n \alpha_n(p(s)) < \beta(1 - p).$$

Next, I claim that $F_{p,n}(\epsilon) > 0$. Suppose this is not true by contradiction. Then, it follows that $F_{p,n}(s) = 0$ for all $s \in [0, \epsilon]$, implying by definition of α that $\alpha_n(p) = 1$, contradicting our assumption that $\alpha_n(p) < 1$.

Now, define the following deviation $\tilde{F}_{p,n}$, which shifts the mass $F_{p,n}$ places on $[0, \epsilon]$ to ∞:

$$\tilde{F}_{p,n}(s) = \begin{cases} 0 & \text{if } s \in [0, \epsilon] \\ F_{p,n}(s) - F_{p,n}(\epsilon) & \text{if } s \in (\epsilon, \infty) \\ 1 & \text{if } s = \infty \end{cases}$$

The admissibility (i.e., right-continuity and piecewise twice-differentiability) of $\tilde{F}_{p,n}$ follows from the admissibility of $F_{p,n}$. We now wish to show that $\tilde{F}_{p,n}$ is a profitable deviation at (p, n). Let Ψ denote the first-report distribution under the strategy profile where all players play $F_{p,n}$, and let $\tilde{\Psi}$ denote the first-report distribution under the strategy profile where i plays $\tilde{F}_{p,n}$ and all $j \neq i$ play $F_{p,n}$.

By definition of Ψ,

$$\tilde{\Psi}^i(s) = \Psi^i(s) - X(s)$$

where

$$X(s) = \begin{cases} p \int_0^s e^{-\lambda r(N-n)}(1 - F_{p,n}(r))^{N-n} d(e^{-\lambda r}(F_{p,n}(r) - 1)) + (1-p) \int_0^s (1 - F_{p,n}(r))^{N-n} dF_{p,n}(r) & \text{if } s \in [0, \epsilon] \\ X(\epsilon) & \text{if } s > \epsilon \end{cases}$$

Now, note that $X(s)$ is weakly increasing in s. Note further that because $F_{p,n}(\epsilon) \in (0, 1]$, it follows that $F_{p,n}(s)$ strictly increases on $[0, \epsilon]$. Thus, $X(s)$ is strictly increasing at some $s \in [0, \epsilon]$. Now, by the above definition:

$$\int_0^\infty [k_n \alpha_n(p(s)) - \beta(1 - p^i(s))] d\tilde{\Psi}^i(s) - \int_0^\infty [k_n \alpha_n(p(s)) - \beta(1 - p^i(s))] d\Psi^i(s)$$
$$= \int_0^\epsilon [k_n \alpha_n(p(s)) - \beta(1 - p(s))] dX(s) > 0$$

where the strict inequality follows from the fact that $X(s)$ is strictly increasing on $[0, \epsilon]$ and the above-established fact that $k_n \alpha_n(p(s)) < \beta(1 - p(s))$ for all $s \in [0, \epsilon]$.

Next, let us consider $\tilde{\Psi}^{-i}(s)$. It again follows from the definition of Ψ that

$$\tilde{\Psi}^{-i}(s) = \Psi^{-i}(s) - Y(s)$$

where

$$Y(s) = p \int_0^s [e^{-\lambda r}(1 - F_{p,n}(r))]^{n-2} F(\min\{r, \epsilon\}) d(e^{-\lambda r}(F_{p,n}(r) - 1)) + \\ (1 - p) \int_0^s (1 - F_{p,n}(r))^{n-2} F_{p,n}(\min\{r, \epsilon\}) dF_{p,n}(r)$$

Thus,

$$\int_0^\infty V_{p^{-i}(s),n+1} d\tilde{\Psi}^{-i}(s) - \int_0^\infty V_{p^{-i}(s),n+1} d\Psi^{-i}(s) = \int_0^\infty V_{p^{-i}(s),n+1} dY(s) \geq 0$$

where the final inequality follows for from the fact that $Y(s)$ is increasing in s and $V_{p^{-i}(s),n+1} \geq 0$. Combining the previous two inequalities, we obtain that

$$V_{p,n}(\tilde{F}_{p,n}) > V_{p,n}(F_{p,n})$$

and thus i can profitably deviate at (p, n). Contradiction. □

Lemma 8. *If $\alpha_n(p) < 1$ and (p, n) is on-path, then there exists an $\epsilon > 0$ such that*

$$V_{p,n} = V_{p,n}(\delta_s) \text{ for all } s \in [0, \epsilon) \cup \infty.$$

Proof of Lemma 8. Assume that $\alpha_n(p) < 1$. Note that by the right twice-differentiability of $F_{p,n}$, and by (1.3), that $\alpha_n(p(s))$ is right-continuous in s. Thus, there exists an $\epsilon > 0$ and $d > 0$ such that

$$\alpha_n(p(s)) < 1 - d \text{ for all } s \in [0, \epsilon].$$

I claim that for all $s \in [0, \epsilon)$, $V_{p,n} = V_{p,n}(\delta_s)$. Suppose to the contrary that for some $s \in [0, \epsilon)$,

$$V_{p,n}(\delta_s) < V_{p,n}$$

Now, I claim that $V_{p,n}(\delta_s)$ is right-continuous in s. To see why this is the case, note that by definition,

$$V_{p,n}(\delta_s) = \int_0^s k_n \alpha_n(p(r)) d\Psi^i(r) + (N-n) \int_0^s V_{p^i(r),n} d\Psi^{-i}(r) + \\ (1 - \sum_j \Psi^j(s))[k_n \alpha_n(p(s)) - \beta(1 - p(s))]$$

Where $\Psi^j(s)$ is the first-report distribution that arises when i plays δ_∞ and all $j \neq i$ play $F_{p,n}$. The right-continuity with respect to s then follows from the absolute continuity of Ψ^j (which follows from Lemma 1), as well as the right-continuity of $\alpha_n(p(s))$ with respect to s, which follows from the right-continuity of $F_{p,n}(s)$ by assumption.

Given the right continuity of $V_{p,n}(\delta_s)$, there exists some $\epsilon' \in (0, \epsilon - s)$ and $x > 0$ such that

$$V_{p,n} - V_{p,n}(\delta_r) > x \text{ for all } r \in [s, s + \epsilon']$$

Now I claim that there must exist some $s^* \in [0, \infty]$ such that $V_{p,n} = V_{p,n}(\delta_{s^*})$. Suppose by contradiction that $V_{p,n} > V_{p,n}(\delta_{s^*})$ for all $s^* \in [0, \infty]$. Letting $F_{p,n}$ denote the firm's equilibrium strategy, it follows that $V_{p,n} = V_{p,n}(F_{p,n})$. It follows from (A.3) that

$$V_{p,n}(F_{p,n}) = \int_0^\infty V_{p,n}(\delta_s) dF_{p,n}(s) + (1 - \lim_{s \to \infty} F_{p,n}) V_{p,n}(\delta_\infty) < V_{p,n}$$

where the strict inequality follows from the assumption that $V_{p,n} > V_{p,n}(\delta_{s^*})$ for all s^*. Contradiction.

Now, define the following deviation $\tilde{F}_{p,n}$ which shifts all the mass from $[s, s+\epsilon']$ to s^*. Specifically, when $s^* < s$:

$$\tilde{F}_{p,n}(t) = \begin{cases} F_{p,n}(t) + F_{p,n}(s+\epsilon) - F_{p,n}(s) & \text{if } t \in [s^*, s] \\ F_{p,n}(s+\epsilon) & \text{if } t \in (s, s+\epsilon'] \\ F_{p,n}(t) & \text{otherwise.} \end{cases}$$

Meanwhile, when $s^* > s + \epsilon$:

$$\tilde{F}_{p,n}(t) = \begin{cases} F_{p,n}(s) & \text{if } t \in [s, s+\epsilon] \\ F_{p,n}(t) - [F_{p,n}(s+\epsilon') - F_{p,n}(s)] & \text{if } t \in (s+\epsilon', s^*) \\ F_{p,n}(t) & \text{otherwise.} \end{cases}$$

Now, by definition:

$$V_{p,n}(\tilde{F}_{p,n}) = V_{p,n}(F_{p,n}) + \int_s^{s+\epsilon'} [V_{p,n}(\delta_{s^*}) - V_{p,n}](\delta_r) dF_{p,n}(r) \geq V_{p,n}(F_{p,n}) + x\epsilon' > V_{p,n}(F_{p,n})$$

Thus, $\tilde{F}_{p,n}$ is a profitable deviation. Contradiction.

It remains to show that $V_{p,n} = V_{p,n}(\delta_\infty)$. Suppose by contradiction that $V_{p,n} > V_{p,n}(\delta_\infty)$. It follows that $\lim_{t\to\infty} F_{p,n}(t) = 0$, because otherwise, the firm could profitably deviate by placing no mass on $t = \infty$. But this implies that for some $s \in (0, \infty]$,

$$\lim_{t\to s-} b_n(p(t)) = \infty \Rightarrow \lim_{t\to s-} \alpha_n(p(t)) = 0,$$

which contradicts Lemma 7. □

Lemma 9. *$\alpha_n(p(s))$ is continuous in s for all (p, n) on path such that $s > 0$.*

Proof of Lemma 9. Fix a (p, n) on-path. I first claim that for all $s \geq 0$,

$$\alpha_n(p(s)) = \frac{\lambda p(s)}{\lambda p(s) + \frac{F'_{p,n}(s+)}{1-F_{p,n}(s)}} \tag{A.4}$$

To see why, note that it follows from Lemma 7 that $(p(s), n)$ is on-path for all $s \geq 0$. Thus, by Lemma 1, $F_{p(s),n}(0) = 0$, and by (1.3)

$$\alpha_n(p(s)) = \frac{\lambda p(s)}{\lambda p(s) + F'_{p(s),n}(0+)}.$$

Next, it follows from (1.2) that

$$F'_{p(s),n}(0+) = \frac{F'_{p,n}(s+)}{1 - F_{p,n}(s)}.$$

Combining the previous two equations yields (A.4). It thus follows from the right-differentiability and piecewise twice-differentiability of $F_{p,n}$ that $\alpha_n(p(s))$ is right-continuous in s. It remains to show that it is left-continuous. Suppose by contradiction there exists an s such that $\alpha_n(p(s))$ is left-discontinuous. Then there exists some $d > 0$ such that for all $\epsilon > 0$, there exists an $s_\epsilon \in (s - \epsilon, s)$ such that

$$|\alpha_n(p(s_\epsilon)) - \alpha_n(p(s))| > d.$$

First consider the case where for all $\epsilon > 0$, there exists an $s_\epsilon \in (s - \epsilon, s)$ such that $\alpha_n(p(s_\epsilon)) - \alpha_n(p(s)) > d$. I begin by claiming that for all $\epsilon > 0$,

$$V_{p(s_\epsilon),n} = V_{p(s_\epsilon),n}(\delta_{s-s_\epsilon}). \tag{A.5}$$

To this end, first note that there exists some $s^* \in (s, \infty]$ such that $V_{p(s_\epsilon),n} = V_{p(s_\epsilon),n}(\delta_{s^*})$. To see why this must hold, suppose not, by contradiction. Then it must be that $F_{p(s_\epsilon),n}$ places full mass on $[s_\epsilon, s]$, and thus, either Lemma 1 or (7) would be violated. Thus, we have

$$V_{p(s_\epsilon),n} = \int_0^{s-s_\epsilon} k_n\alpha_n(p(r))d\Psi^i(r) + (N-n)\int_0^{s-s_\epsilon} V_{p^i(r),n+1}d\Psi^{-i}(r) +$$

$$(1-\sum_j \Psi^j(s-s_\epsilon))V_{p(s),n}(\delta_{s^*-s}) = \int_0^{s-s_\epsilon} k_n\alpha_n(p(r))d\Psi^i(r) + (N-n)\int_0^{s-s_\epsilon} V_{p^i(r),n+1}d\Psi^{-i}(r)$$

$$+(1-\sum_j \Psi^j(s-s_\epsilon))V_{p(s),n}(\delta_0) = V_{p(s_\epsilon),n}(\delta_{s-s_\epsilon})$$

where Ψ is the first-report distribution associated with the strategy profile in which i plays δ_∞ and all $j \neq i$ play $F_{p(s_\epsilon),n}$. Note that the equality follows from the fact that $\alpha_n(p(s)) < 1$, and thus by Lemma 8, $V_{p(s),n} = V_{p(s),n}(\delta_0)$. However, note that for all $\epsilon > 0$,

$$V_{p(s_\epsilon),n}(\delta_{s-s_\epsilon}) = \int_0^{s-s_\epsilon} k_n\alpha_n(p(r))d\Psi^i(r) + (N-n)\int_0^{s-s_\epsilon} V_{p^i(r),n+1}d\Psi^{-i}(r)$$

$$+(1-\sum_j \Psi^j(s-s_\epsilon))[k_n\alpha_n(p(s),n) - \beta(1-p(s))]$$

Because the Ψ^j are absolutely continuous,

$$\lim_{\epsilon\longrightarrow 0} V_{p(s_\epsilon),n}(\delta_{s-s_\epsilon}) = k_n\alpha_n(p(s),n) - \beta(1-p(s))$$

Then, by the assumption that $\alpha_n(p(s_\epsilon))-\alpha_n(p(s)) < d$, for all $\epsilon > 0$ sufficiently small $V_{p(s_\epsilon),n}(\delta_0) = k_n\alpha_n(p(s_\epsilon),n) - \beta(1-p(s_\epsilon)) > V_{p(s_\epsilon),n}(\delta_{s-s_\epsilon})$, contradicting (A.5).

Next, consider the case where for all $\epsilon > 0$, $\alpha_n(p(s)) - \alpha_n(p(s_\epsilon)) > d$. As noted above, $\lim_{\epsilon\to 0} V_{p(s_\epsilon),n}(\delta_{s-s_\epsilon}) = V_{p(s),n}(\delta_0)$. Thus, for ϵ sufficiently small,

$$V_{p(s_\epsilon),n}(\delta_{s-s_\epsilon}) > k_n\alpha_n(p(s_\epsilon)) - \beta(1-p(s_\epsilon)) = V_{p(s_\epsilon),n}(\delta_0)$$

However, since $\alpha_n(p(s_\epsilon)) < 1$ for all $\epsilon > 0$, by Lemma 8, $V_{p(s_\epsilon),n} = V_{p(s_\epsilon),n}(\delta_0)$. Contradiction.

□

Proof of Proposition 1. I begin by showing that $\alpha_n(p) = 1$ whenever $k_n < \beta$ and $p \leq p_n^* \equiv \frac{k_n-\beta}{k_n/n-\beta}$.

To this end, fix an n, and suppose that $k_n < \beta$. I first show that for all $q < \frac{\beta - k_n}{\beta}$, $\alpha_n(q) = 1$. Note that for all such q

$$V_{q,n}(\delta_0) = k_n\alpha_n(q) - \beta(1-q) \leq k_n - \beta(1-q) < k_n - \beta(1 - \frac{\beta - k_n}{\beta}) = 0.$$

Since $V_{q,n} \geq V_{q,n}(\delta_\infty) \geq 0$, it follows $V_{q,n} > V_{q,n}(\delta_0)$. Thus, by Lemma 8, $\alpha_n(q) = 1$. Now, let

$$q_n^* \equiv \sup\{p|\alpha_n(q) = 1 \text{ for all } q < p\}$$

It follows from the above that $q_n^* \geq \frac{\beta - k_n}{\beta} > 0$. Now suppose by contradiction that $q_n^* < p_n^*$. By Lemma 9, there exists an $\epsilon > 0$ such that for all $p \in (q_n^*, q_n^* + \epsilon)$, $\alpha_n(p) < 1$, and thus, by Lemma 8

$$V_{p,n} = V_{p,n}(\delta_0) = k_n\alpha_n(p) - \beta(1-p)$$

Thus, it follows from Lemma 9 that

$$\lim_{p \to q_n^*+} V_{p,n} = k_n - \beta(1 - q_n^*) \tag{A.6}$$

By definition of V, because by Lemma 1 $F_{p,n}$ is absolutely continuous, it follows that $V_{p,n}(\delta_\infty)$ is as well, and thus:

$$\lim_{p \to q_n^*+} V_{p,n}(\delta_\infty) = V_{q_n^*,n}(\delta_\infty) = \frac{k_n q_n^*}{n} \tag{A.7}$$

In order for δ_∞ to not serve as a profitable deviation for $p \in (q_n^*, q_n^* + \epsilon)$, it must be that for all such p, $V_{p,n}(\delta_0) \geq V_{p,n}(\delta_\infty)$. Taking a limit we obtain that

$$\lim_{p \to q_n^*+} V_{p,n}(\delta_0) \geq \lim_{p \to q_n^*+} V_{p,n}(\delta_\infty)$$

Substituting (A.6) and (A.7) above, we obtain that $\frac{k_n q_n^*}{n} \leq k_n - \beta(1 - q_n^*)$. However, $k_n \leq \beta$ and $q_n^* < p$ implies that $\frac{k_n q^*}{n} > k_n - \beta(1 - q^*)$. Contradiction.

Next, we show that $\alpha_n(p) < 1$ whenever $\beta \le k_n$ or $p > p_n^*$. To this end, assume $\beta \le k_n$ or $p > p_n^*$. Assume by contradiction that $\alpha_n(p) = 1$. Also assume by induction that if $n < N$, then the statement holds for $n + 1$.

First, consider the case where $\alpha_n(q) = 1$ for all $q < p$. By (1.3), this implies that $F'(q, n) = 0$. Furthermore, by Lemma 1, this implies that $F_{p,n}(s) = 0$ for all $s > 0$, i.e., $F_{p,n} = \delta_\infty$. However,

$$V_{p,n}(\delta_0) = k_n - \beta(1 - p) > \frac{k_n p}{n} = V_{p,n}(\delta_\infty),$$

where the above strict inequality follows from the above assumption that either $\beta \le k_n$ or $p > p_n^*$. Contradiction.

Next, consider the case where $\alpha_n(q) < 1$ for some $q < p$. By Lemma 9, for all $\epsilon > 0$ sufficiently small, there exists some $\bar{p} < p$ and $\bar{s} > 0$ such that $\alpha_n(\bar{p}) \in (1 - \epsilon, 1)$ and $\alpha_n(q)$ is strictly increasing on $[\bar{p}(\bar{s}), \bar{p}]$. By Lemma 8, there exists some $\Delta \in (0, s)$ such that

$$V_{\bar{p},n}(\delta_\Delta) = V_{\bar{p},n}(\delta_0).$$

By definition,

$$\begin{aligned} V_{\bar{p},n}(\delta_\Delta) = &\int_0^\Delta k_n \alpha_n(\bar{p}(s)) d\Psi^i(s) + (N - n) \int_0^\Delta V_{\bar{p}^i(s),n+1} d\Psi^{-i}(s) + \\ &(1 - \sum_j \Psi^j(\Delta))(k_n \alpha_n(\bar{p}(\Delta)) - \beta(1 - \bar{p}(\Delta)) \end{aligned}$$

where Ψ is the first-report distribution associated with the strategy profile where i plays δ_Δ and all $j = i$ play $F_{p,n}$. Meanwhile,

$$\begin{aligned} &V_{\bar{p},n}(\delta_0) = k_n \alpha_n(\bar{p}) - \beta(1 - \bar{p}) \\ &= \int_0^\Delta k_n \alpha_n(\bar{p}) d\Psi^i(s) + (N - n) \int_0^\Delta k_n \alpha_n(\bar{p}) - \beta(1 - \bar{p}^i(s)) d\Psi^{-i}(s) \\ &+ (1 - \sum_j \Psi^j(\Delta))(k_n \alpha_n(\bar{p}) - \beta(1 - \bar{p}(\Delta)) \end{aligned}$$

Thus, in order to preserve the above equality, for some $r \in (0, \bar{s})$,

$$k_n\alpha_n(\bar{p}) - \beta(1 - p^i(r)) < V_{\bar{p}^i(r),n+1}. \tag{A.8}$$

First, consider the case where $\alpha_{n+1}(\bar{p}^i(r))) < 1$. Then, for $\epsilon > 0$ sufficiently small

$$V_{\bar{p}^i(r),n+1} = V_{\bar{p}^i(r),n+1}(\delta_0) = k_{n+1}\alpha_{n+1}(\bar{p}^i(r)) - \beta(1 - \bar{p}^i(r)) < k_n\alpha_n(\bar{p}) - \beta(1 - \bar{p}^i(r))$$

where the first equality follows from Lemma 8. Thus, equation (A.8) is violated. Contradiction.

Next, consider the case where $\alpha_{n+1}(\bar{p}^i(r)) = 1$ and $\beta < k_n$. By the inductive assumption, it follows that $\alpha_{n+1}(q) = 1$ for all $q \leq \bar{p}^i(s)$. Thus, $F_{\bar{p}^i(s),n+1} = \delta_\infty$. So, we have that for ϵ sufficiently small:

$$\begin{aligned} V_{\bar{p}^i(r),n+1} &= V_{\bar{p}^i(r),n+1}(\delta_\infty) = \frac{k_{n+1}\bar{p}^i(r)}{N-n} \leq \bar{p}^i(r)k_n\alpha_n(\bar{p}) + (1 - \bar{p}^i(s))k_n\alpha_n(\bar{p}) - \beta) \\ &= k_n\alpha_n(\bar{p}) - \beta(1 - \bar{p}^i(s)) \end{aligned}$$

Again, this is a contradiction of (A.8).

Finally, consider the case where $\alpha_{n+1}(\bar{p}^i(r)) = 1$ and $\beta \geq k_n$. Recall by Proposition 1 that $\alpha_n(q) = 1$ for all $q \geq p_n^*$. Thus, because $\alpha_n(\bar{p}) < 1$, it follows from (9) that $\alpha_n(\bar{p}(s))$ must be strictly increasing in s for some $s > r$. Formally, let

$$r' \equiv \inf\{s > r | \alpha_n(\bar{p}(s)) \text{ is strictly increasing}\}.$$

First, I claim that

$$k_n\alpha_n(\bar{p}(r')) - \beta(1 - \bar{p}^i(r)) < V_{\bar{p}^i(r),n+1} \tag{A.9}$$

By the inductive assumption, since $\alpha_{n+1}(\bar{p}^i) = 1$, it must be that $\alpha_{n+1}(q) = 1$ for all $q < \bar{p}^i(r)$. Because $\alpha_n(\bar{p}(s))$ is weakly decreasing for $s \in [r, r']$, it follows by definition of $\bar{p}^i(s)$ that $\bar{p}^i(s) <$

$\bar{p}^i(r)$ for all $s \in [r, r']$. Thus, for all $s \in [r, r']$

$$V_{\bar{p}^i(s),n+1} = \frac{k_{n+1}\bar{p}^i(s)}{N-n}.$$

It follows from this that for all $s \geq r$,

$$\begin{aligned} k_n\alpha_n(\bar{p}(s)) - \beta(1-\bar{p}^i(s)) &< V_{\bar{p}^i(s),n+1} \\ \Leftrightarrow k_n\alpha_n(\bar{p}(s)) - \beta(1-\bar{p}^i(s)) &< \frac{k_{n+1}\bar{p}^i(s)}{N-n} \\ \Leftrightarrow \bar{p}^i(s) &< \frac{\beta - k_n\alpha_n(\bar{p}(s))}{\beta - k_{n+1}/(N-n)} \end{aligned}$$

Now, because $\alpha_n(\bar{p}(s))$ is strictly decreasing on $s \in [0, r]$,

$$k_n\alpha_n(\bar{p}(r)) - \beta(1-p^i(r)) < k_n\alpha_n(\bar{p}) - \beta(1-\bar{p}^i(r)) < V_{\bar{p}^i(r),n+1}$$

where the second inequality holds by the same reasoning presented in the explanation for (A.8). Thus we have

$$\bar{p}^i(r') < \bar{p}^i(r) < \frac{\beta - k_n\alpha_{n+1}(\bar{p}(r))}{\beta - k_{n+1}/(N-n)} < \frac{\beta - k_n\alpha_{n+1}(\bar{p}(r'))}{\beta - k_{n+1}/(N-n)}$$

which implies (A.9).

It follows from this that there exists an $r'' > r'$ such that for all $s \in [r', r'']$, $\alpha_n(\bar{p}(s))$ is weakly decreasing and $V_{\bar{p}^i(s),n+1} > k_n\alpha_n(\bar{p}(r')) - \beta(1-p^i(s))$. I now claim that

$$V_{\bar{p}(r'),n}(\delta_0) < V_{\bar{p}(r'),n}(\delta_{r''-r'}).$$

To see why, note that by definition,

$$V_{\bar{p}(r'),n}(\delta_{r''-r'}) - V_{\bar{p}(r'),n}(\delta_0) = \int_{r'}^{r''} k_n[\alpha_n(p(s)) - \alpha_n(p(r'))]d\Psi^i(s) + \int_{r'}^{r''} [V_{p^i(s),n+1} - (k_n\alpha_n(p(r')) - \beta(1-p^i(s)))]d\Psi^{-i}(s) + \sum_j (\Psi^j(r'') - \Psi^j(r'))k_n(\alpha_n(p(r'')) - k_n\alpha_n(p(r')))$$

Since $\alpha_n(p(s)) \geq \alpha_n(p(r'))$ and $V_{p^i(s),n+1} > k_n\alpha_n(p(r')) - \beta(1-p^i(s))$ $s \in [r', r'']$, it follows that $V_{\bar{p}(r'),n}(\delta_{r''-r'}) - V_{\bar{p}(r'),n}(\delta_0) > 0$. However, this contradicts Lemma 8. □

Proof of Proposition 2. Proof by induction. Fix an n, and assume that $\alpha_m(p)$ satisfies the above for all $m > n$ such that (p, m) is on-path.

We begin by showing that (ODE) must hold whenever $\alpha_n(p) < 1$. To this end, assume that $\alpha_n(p) < 1$. Then, by Lemma 8, there exists an $\epsilon > 0$ such that for all $\Delta \in (0, \epsilon)$,

$$\frac{V_{p,n}(\delta_\Delta) - V_{p,n}(\delta_0)}{\Delta} = 0 \tag{A.10}$$

Recall that by definition of V, that

$$V_{p,n}(\delta_0) = k_n\alpha_n(p) - \beta(1-p).$$

Meanwhile

$$V_{p,n}(\delta_\Delta) = \int_0^\Delta k_n\alpha_n(p(s))\Psi^i(s)ds + (N-n)\int_0^\Delta V_{p^{-i}(s),n+1}\Psi^{-i}(s)ds + (1 - \sum_j \lim_{s\to\Delta-} \Psi^j(s))[k_n\alpha_n(p(\Delta)) - \beta(1-p(\Delta))]$$

where Ψ is the first-report distribution associated with the strategy profile in which i plays δ_∞ and all $j \neq i$ play $F^j_{p,S}$. Specifically, for all $s > 0$,

$$\Psi^i(s) = p\lambda \int_0^s e^{-\lambda rn}(1 - F_{p,n}(r))^{N-n} dr$$

$$\Psi^{-i}(s) = p \int_0^s e^{-\lambda r(N-n)}(1-F_{p,n}(r))^{n-2} d(-e^{-\lambda r}(1-F_{p,n}(r))) + (1-p) \int_0^s (1-F_{p,n}(r))^{n-2} dF_{p,n}(r)$$

It follows from Lemma 1 that, for all j, Ψ^j is also absolutely continuous, I.e., there exists a function ψ_j such that:

$$\Psi^j(s) = \int_0^s \psi^j(r)dr.$$

Specifically, according to Lemma 1, one such ψ^i and ψ^{-i} are given by the following:

$$\psi^i(s) = p\lambda e^{-\lambda sn}(1 - F_{p,n}(s))^{N-n}$$

$$\psi^{-i}(s) = pe^{-\lambda sn}(\lambda + F'_{p,n}(s+) - \lambda F_{p,n}(s))(1 - F_{p,n}(s))^{n-2} + (1 - p)(1 - F_{p,n}(s))F'_{p,n}(s+)$$

Substituting these expressions for both $V_{p,n}(\delta_0)$ and $V_{p,n}(\delta_\Delta)$ into (A.10) and rearranging, we obtain that for all $\Delta \in (0, \epsilon)$,

$$K_1(\Delta) + K_2(\Delta) + K_3(\Delta) = 0 \tag{A.11}$$

where

$$K_1(\Delta) \equiv \frac{\int_0^\Delta k_n[(\alpha_n(p(s)) - \alpha_n(p)) + \beta(1 - p)]\psi^i(s)ds}{\Delta}$$

$$K_2(\Delta) \equiv \frac{(N - n) \int_0^\Delta [V_{p^{-i}(s),n+1} - k_n\alpha_n(p) + \beta(1 - p)]\psi^{-i}(s)ds}{\Delta}$$

$$K_3(\Delta) \equiv \frac{(1 - \sum_j \lim_{s\to\Delta-} \psi^j(\Delta))[k_n(\alpha_n(p(\Delta)) - \alpha_n(p)) + \beta(p(\Delta) - p)]}{\Delta}$$

Now, we consider $\lim_{\Delta\to 0+}$ of $K_1(\Delta)$, $K_2(\Delta)$, and $K_3(\Delta)$ separately.

For $K_1(\Delta)$, it follows from L'Hôpital's Rule, together with the continuity of $\alpha_n(p(\Delta))$ (i.e., Lemma 9) and $\psi^i(\Delta)$ in Δ that

$$\lim_{\Delta\to 0+} K_1(\Delta) = \lim_{\Delta\to 0+}[k_n(\alpha_n(p(\Delta)) - \alpha_n(p)) + \beta(1-p)]\psi^i(\Delta) = \beta(1-p)\psi^i(0) = \beta(1-p)p\lambda.$$

For $K_2(\Delta)$, it again follows from L'Hôpital's Rule, together with the right-continuity of $V_{p^{-i}(\Delta),n+1}$ in Δ that

$$\begin{aligned}\lim_{\Delta\to 0+} K_2(\Delta) &= (N-n)\lim_{\Delta\to 0+}[V_{p^{-i}(\Delta),n+1} - k_n\alpha_n(p) + \beta(1-p)]\psi^{-i}(\Delta)\\ &= (N-n)[V_{p^{-i},n+1} - k_n\alpha_n(p) + \beta(1-p)](\frac{\lambda p}{\alpha_n(p)})\end{aligned}$$

where the final inequality follows from the fact that at all (p,n) on-path, $\alpha_n(p) = \frac{\lambda p}{\lambda p + F'_{p,n}(0)}$.

For $K_2(\Delta)$, first note that by the continuous differentiability of $\Psi^j(s)$ that

$$\lim_{\Delta\to 0+}\sum_j \lim_{s\to\Delta-}\Psi^j(s) = 0.$$

Thus, it follows from the right-differentiability of $\alpha_n(p(\Delta))$ in Δ that

$$\begin{aligned}\lim_{\Delta\to 0+} K_3(\Delta) &= k_n\lim_{\Delta\to 0}\frac{\alpha_n(p(\Delta)) - \alpha_n(p)}{\Delta} + \beta\lim_{\Delta\to 0+}\frac{p(\Delta)-p}{\Delta} = k_n\frac{d}{d\Delta}\alpha_n(p(\Delta))\Big|_{\Delta=0+} + \beta p'(\Delta)\Big|_{\Delta=0+}\\ &= p'(\Delta)\Big|_{\Delta=0+}[k_n\alpha'_n(p) + \beta] = -\lambda pn(1-p)[k_n\alpha'_n(p) + \beta]\end{aligned}$$

Since we have shown that $\lim_{\Delta\to 0+} K_1(\Delta)$, $\lim_{\Delta\to 0+} K_2(\Delta)$, and $\lim_{\Delta\to 0+} K_3(\Delta)$ exist, and are given by the above expressions, it follows from (A.11) that

$$\lim_{\Delta\to 0+} K_1(\Delta) + \lim_{\Delta\to 0+} K_2(\Delta) + \lim_{\Delta\to 0+} K_3(\Delta) = 0.$$

Substituting in the above expressions for $K_1(\Delta)$, $K_2(\Delta)$ and $K_3(\Delta)$, we obtain (ODE).

Now, we wish to establish that (ODE) must hold whenever $k_n \geq \beta$ or $p > p_n^*$. It follows from Proposition 1 that $\alpha_n(p) < 1$, and thus by the above, (ODE) must hold.

Finally, we establish the two limit conditions presented in the proposition. We begin by establishing that when $k_n \geq \beta$, $\lim_{p\to 0+} \alpha_n(p) = \beta/k_n$. To this end, first note by Lemma 8 that for all $p > 0$, $V_{p,n}(\delta_0) = V_{p,n}(\delta_\infty)$. Note further that

$$\lim_{p\to 0+} V_{p,n}(\delta_\infty) = 0.$$

Thus,

$$\lim_{p\to 0+} V_{p,n}(\delta_0) = \lim_{p\to 0+} k_n\alpha_n(p) - \beta = 0,$$

and therefore, $\lim_{p\to 0+} \alpha_n(p) = \frac{\beta}{k_n}$. Next, let us consider the case where $k_n < \beta$. That $\lim_{p\to p_n^*+} \alpha_n(p) = 1$ follows from Lemma 9, since by Proposition 1, $\alpha_n(p_n^*) = 1$. □

Before proceeding with the rest of the characterization, I define a problem (P) on α. I then show that α consistutes an equlibrium if and only if it satisfies (P) (Lemma 10). Thus, existence and uniqueness of an equilibrium (Theorem 1) will reduce to establishing a unique solution to (P).

Definition 3. α is a solution to (P) if it satisfies the following for all $n \leq N$ and $p \in (0, 1]$:

- If $k_n < \beta$ and $p \leq p_n^* \equiv \frac{k_n-\beta}{k_n/n-\beta}$, then $\alpha_n(p) = 1$.
- If $k_n \geq \beta$ or $p < p_n^*$, then α satisfies (ODE), with limit condition $\lim_{p\to 0+} \alpha_n(p) = \beta/k_n$ if $k_n \geq \beta$ and $\lim_{p\to p_n*+} \alpha_n(p) = 1$ if $k_n < \beta$.
- $\alpha_n(1) = 0$.

Lemma 10. *(α, F) is an equilibrium if and only if at all (p, n) on-path, α is both consistent with F and a solution to (P).*

Proof of Lemma 10. Fix an (α, F). I begin by establishing the necessity of the three conditions specified in Definition 3 for (α, F) to be an equilibrium. First we establish the necessity of the first bullet of Definition 3. To this end, recall that by the selection assumption, $F_{1,n}(0) = 1$. Thus, it follows from (1.3) that $\alpha_n(1) = 0$ if $(p = 1, n)$ is on-path. Bullets two and three of Definition (3) follow immediately from Proposition 1 and Proposition 2, respectively.

Next, we establish the sufficiency of the above conditions for (α, F) to be an equilibrium. We begin by considering the case in which $k_n < \beta$ and $p \leq p_n^*$. It follows from (P) that $\alpha_n(q) = 1$ for all $q \leq p$. Thus, by (1.3), $F_{p,n} = \delta_\infty$. We wish to show that there exist no profitable deviations in this case, i.e., that $V_{p,n} = V_{p,n}(\delta_\infty)$. It suffices to show that

$$V_{p,n}(\delta_\infty) \geq V_{p,n}(\delta_s) \text{ for all } s \in [0, \infty). \tag{A.12}$$

First, note that for all $s \in (0, \infty)$,

$$V_{p,n}(\delta_s) = k_n(1 - p(1 - e^{-\lambda sn})(\frac{N-n}{N-n+1})) - \beta(1-p) \leq k_n - \beta(1-p) = V_{p,n}(\delta_0).$$

Further, $k_n \leq \beta$ and $p \leq p_n^*$ implies that

$$V_{p,n}(\delta_0) = k_n - \beta(1-p) \leq \frac{k_n}{n} = V_{p,n}(\delta_\infty)$$

Thus, $V_{p,n}(\delta_\infty) \geq V_{p,n}(\delta_s)$ for all $s \in [0, \infty)$

Next, we show that $F_{p,n}$ is optimal when $k_n \geq \beta$ or $p < p_n^*$. To this end, we begin by showing that

$$\frac{d}{d\Delta} V_{p,n}(\delta_\Delta) = 0 \text{ for all } \Delta \in [0, \infty) \text{ if } k_n \geq \beta \text{ and for all } \Delta \in [0, t^*) \text{ if } k_n < \beta \tag{A.13}$$

where t^* is the unique value such that $p(t^*) = p_n^*$. Note that

$$\begin{aligned} V_{p,n}(\delta_\Delta) = \int_0^\Delta k_n \alpha_n(p(s)) d\Psi^i(s) + \int_0^\Delta V_{p^i(s),n+1} d\Psi^{-i}(s) + \\ (1 - \sum_j \Psi^j(\Delta))(\alpha_n(p(\Delta)) - \beta(1 - p(\Delta))) \end{aligned} \tag{A.14}$$

where Ψ is the first-report distribution associated with the strategy profile in which i plays δ_∞ and all $j \neq i$ play $F_{p,n}$. Then, it follows that

$$\begin{aligned}
&\frac{d}{d\Delta}V_{p,n}(\delta_\Delta) \\
&= k_n\alpha_n(p(\Delta))\Psi^{i\prime}(\Delta) + (N-n)V_{p^i(\Delta),n+1}\Psi^{-i\prime}(\Delta) + (1-\sum_j \Psi^j(\Delta))p'(\Delta)[\alpha_n'(p(\Delta)) - \beta] \\
&- \sum_j \Psi^{j\prime}(\Delta)(k_n\alpha_n(p(\Delta)) - \beta(1-p(\Delta))) \\
&= (N-n)[V_{p^i(\Delta),n+1} - k_n\alpha_n(p(\Delta)) + \beta(1-p(\Delta))]\Psi^{-i\prime}(\Delta) - \beta(1-p(\Delta))\Psi^{i\prime}(\Delta) \\
&+ (1-\sum_j \Psi^j(\Delta))p'(\Delta)(k_n\alpha_n'(p(\Delta)) - \beta),
\end{aligned}$$

where $\Psi^{i\prime}(t) \equiv \frac{d}{dt}\Psi^i(t)$.

In the above, the existence of $\Psi^{j\prime}(\Delta)$ follows from the differentiability of α_n at $p(\Delta)$, and thus, the differentiability of $F_{p,n}$ at Δ. We wish to show that $\frac{d}{d\Delta}V_{p,n}(\delta_\Delta) = 0$. To this end, we begin by deriving expressions for $\Psi^{i\prime}(\Delta)$ and $\Psi^{-i\prime}(\Delta)$. First, it follows by definition of the first-report distribution that:

$$\Psi^i(\Delta) = p\lambda\int_0^\Delta (1-F_{p,n}(s))^{N-n}e^{-\lambda ns}ds.$$

Differentiating this, we obtain:

$$\Psi^{i\prime}(\Delta) = p\lambda(1-F_{p,n}(\Delta))^{N-n}e^{-\lambda n\Delta}$$

Meanwhile:

$$\Psi^{-i}(\Delta) = p\int_0^\Delta (1-F_{p,n}(s))^{n-2}e^{-\lambda(N-n)s}d((F_{p,n}(s)-1)e^{-\lambda s}) + (1-p)\int_0^\Delta (1-F_{p,n}(s))^{n-2}F_{p,n}'(s)ds$$

where the existence of $F_{p,n}'(s)$ again follows from the assumption that α_n is differentiable at $p(s)$.

Differentiating this, we obtain:

$$\Psi^{-i\prime}(\Delta) = p(1-F_{p,n}(\Delta))^{n-2}e^{-\lambda\Delta n}[F'_{p,n}(\Delta)+\lambda(1-F_{p,n}(\Delta))] + (1-p)(1-F_{p,n}(\Delta))^{n-2}f_{p,n}(\Delta)$$
$$= (1-F_{p,n}(\Delta))^{N-n}[\frac{f_{p,n(\Delta)}}{1-F_{p,n}(\Delta)}(pe^{-\lambda\Delta n}+(1-p))+pe^{-\lambda\Delta n}\lambda]$$

It follows from (1.3) and (1.2) that

$$\frac{F'_{p,n}(\Delta)}{1-F_{p,n}(\Delta)} = \lambda p(\Delta)(\frac{1}{\alpha_n(p(\Delta))}-1).$$

Substituting this, along with the definition of $p(\Delta)$ (A.1), we obtain:

$$\Psi^{-i\prime}(\Delta) = \lambda(1-F_{p,n}(\Delta))^{N-n}(pe^{-\lambda\Delta n}+(1-p))\frac{p(\Delta)}{\alpha_n(p(\Delta))}$$

Note further that

$$1-\sum_j \Psi^j(\Delta) = (1-F(\Delta))^{N-n}(pe^{-\lambda\Delta n}+(1-p)) \tag{A.15}$$

Substituting equations the expressions for $\Psi^{i\prime}(\Delta)$, $\Psi^{-i\prime}(\Delta)$, and $1-\sum_j \Psi^j(\Delta)$ into the above equation for $\frac{d}{d\Delta}V_{p,n}(\delta_\Delta)$, and simplifying, we obtain:

$$\frac{d}{d\Delta}V_{p,n}(\delta_\Delta) = K[\frac{(N-n)}{\alpha_n(p(\Delta))}(V^i_{p(\Delta),n+1}-k_n\alpha_n(p(\Delta))+\beta(1-p(\Delta))(1-\alpha_n(p(\Delta))))$$
$$-k_n\alpha'_n(p(\Delta))(1-p(\Delta))n]$$

where $K \equiv \lambda(1-F_{p,n}(\Delta))^{N-n}(pe^{-\lambda\Delta n}+(1-p))p(\Delta)$. Because (ODE) is satisfied at $(p(\Delta), n)$, using it to substitute in for $\alpha'_n(p(\Delta)$, we obtain (A.13).

Now, consider the case where $k_n \geq \beta$. To show $F_{p,n}$ is optimal, it suffices to show that all pure strategies δ_Δ yield the same payoff, i.e., that

$$V_{p,n}(\delta_0) = V_{p,n}(\delta_\Delta) \tag{A.16}$$

for all $\Delta \in [0, \infty]$. It follows directly from (A.13) that (A.16) holds for all $\Delta \in [0, \infty)$. It remains to show that (A.16) holds for $\Delta = \infty$. To this end, first note that by (A.13),

$$
\begin{aligned}
&V_{p,n}(\delta_0) = \lim_{\Delta\to\infty} V_{p,n}(\delta_\Delta) \\
&= \lim_{\Delta\to\infty} \int_0^{\Delta} k_n\alpha_n(p(s))d\Psi^i(s) + (N-n)\lim_{\Delta\to\infty}\int_0^{\Delta} V_{p^i(s),n+1}d\Psi^{-i}(s)+ \\
&\lim_{\Delta\to\infty}(1-\sum_j \Psi^j(\Delta))(k_n\alpha_n(p(\Delta)) - \beta(1-p(\Delta))) \\
&= \int_0^{\infty} k_n\alpha_n(p(\Delta))d\Psi^i(s) + (N-n)\int_0^{\infty} V_{p^i(\Delta),n+1}d\Psi^{-i}(s) = V_{p,n}(\delta_\infty)
\end{aligned}
$$

where the third equality follows from the limit condition $\lim_{p\to 0+}\alpha_n(p) = \beta/k_n$:

$$\lim_{\Delta\to\infty} k_n\alpha_n(p(\Delta)) - \beta(1-p(\Delta)) = \lim_{p\to 0+} k_n\alpha_n(0) - \beta = 0.$$

Finally, consider the case where $k_n < \beta$ and $p > p_n^*$. Because $\alpha_n(p(s)) = 1$ for all $s > t^*$, by (1.3), it follows that $F'_{p,n}(s) = 0$ for all such s. Then, the support of $F_{p,n}$ lies within $[0, t^*] \cup \infty$. Thus, to show $F_{p,n}$ is optimal, it suffices to show that δ_Δ is optimal for $\Delta \in [0, t^*] \cup \infty$. To this end, I first show that

$$V_{p,n}(\delta_\Delta) = V_{p,n}(0) \text{ for all } \Delta \in [0, t^*] \cup \infty \qquad \text{(A.17)}$$

and then show

$$V_{p,n}(\delta_{t^*}) \geq V_{p,n}(\delta_\Delta)\text{for all } \Delta \in (t^*, \infty). \qquad \text{(A.18)}$$

To show (A.17), first recall that it follows from (A.13) that

$$V_{p,n}(\delta_0) = V_{p,n}(\delta_\Delta) \text{ for all } \Delta \in [0, t^*).$$

It remains to show $V_{p,n}(\delta_0) = V_{p,n}(\delta_s)$ for $s \in \{t^*, \infty\}$. For $s = t^*$, it follows from the above that

$$V_{p,n}(\delta_0) = \lim_{\Delta\to t^*-} V_{p,n}(\delta_\Delta) = V_{p,n}(\delta_{t^*})$$

where the final inequality follows from (A.14), observing that α_n is continuous at p_n^* and Ψ^j is continuous at t^*. I will now show $V_{p,n}(\delta_{t^*}) = V_{p,n}(\delta_\infty)$. To this end, note that for all $\Delta \in [t^*, \infty]$:

$$V_{p,n}(\delta_\Delta) = \int_0^{t^*} k_n \alpha_n(p(s)) d\Psi^i(s) + (N-n) \int_0^{t^*} V_{p^i(s),n+1} d\Psi^{-i}(s) + (1 - \sum_j \Psi^j(t^*)) V_{p_n^*,n}(\delta_{\Delta - t^*})$$

Thus, to show $V_{p,n}(\delta_t^*) = V_{p,n}(\delta_\infty)$, it suffices to show that $V_{p_n^*,n}(\delta_0) = V_{p_n^*,n}(\delta_\infty)$. But it follows from the definition of p_n^* that:

$$V_{p_n^*,n}(\delta_0) = k_n - \beta(1 - p_n^*) = \frac{k_n p_n^*}{n} = V_{p_n^*,n}(\delta_\infty).$$

Similarly, to show (A.18), it suffices to show that $V_{p_n^*,n}(\delta_0) \geq V_{p_n^*,n}(\delta_\Delta)$ for all $\Delta \in (0, \infty)$, which we have established in (A.12). □

Proof of Theorem 1. Fix an n. Assume by induction that there exists a unique solution to (P) for all $m > n$. We wish to show that there exists a unique solution to (P) for n. To establish this, it suffices to show there exists a unique solution to the following two limit problems, when $\beta \leq k_n$ and $\beta > k_n$, respectively:

$$\text{(ODE) is satisfied on } (0,1), \text{ and } \lim_{p \to 0+} \alpha_n(p) = \beta / k_n \qquad \text{(LP: } \beta \leq k_n\text{)}$$

$$\text{(ODE) is satisfied on } (0, p^*), \text{ and } \lim_{p \to p_n^*+} \alpha_n(p) = 1. \qquad \text{(LP: } \beta > k_n\text{)}$$

To establish existence and uniqueness to the two above problems, we proceed by extending them to two boundary value problems. To this end, we begin by defining an extension of (ODE') of (ODE), which is identical to (ODE), except that it is well-defined when $p^i \geq 1$. Specifically, define:

$$\alpha_n'(p) = -\frac{1}{k_n(1-p)\alpha_n(p)} \frac{N-n}{N-n+1} [k_n \alpha_n(p) - \tilde{V}_{p^i,n+1} - \beta(1 - \alpha_n(p))(1-p)] \qquad \text{(ODE')}$$

where

$$\tilde{V}_{p^i,n+1} = \begin{cases} V_{p^i,n+1} & \text{if } p^i \in (0,1) \\ 0 & \text{if } p^i \geq 1 \end{cases}$$

Now let us define two boundary value problems on (ODE'):

$$\text{(ODE') is satisfied on } [0,1)\text{, and } \alpha_n(0) = \beta/k_n \qquad \text{(BVP: } \beta \leq k_n\text{)}$$

$$\text{(ODE') is satisfied on } (0,p_n^*]\text{, and } \alpha_n(p^*) = 1. \qquad \text{(BVP: } \beta \geq k_n\text{)}$$

Now we claim that the existence and uniqueness of a solution to (BVP: $\beta \leq k_n$) and (BVP: $\beta \geq k_n$) implies the existence and uniqueness of a solution to (LP: $\beta \leq k_n$) and (LP: $\beta > k_n$), respectively. Let us begin by considering the case where $k_n \geq \beta$. Assume that there exists a unique solution α_n to (BVP: $\beta \leq k_n$). Note that order for α_n to satisfy (BVP: $\beta \leq k_n$), it must be that $\lim_{p \to 0+} \alpha_n(p) = k_n/\beta$. Furthermore, (ODE) and (ODE') are equivalent on $(0,1)$. It follows that α_n is a solution to (LP: $\beta \leq k_n$), thus establishing existence. To establish uniqueness, assume by contradiction there exists some $\tilde{\alpha}_n$ defined on $p \in (0,1)$ that is a solution to (LP: $\beta \leq k_n$) where $\tilde{\alpha}_n(p) \neq \alpha_n(p)$. Now, define $\hat{\alpha}_n$, which extends the domain of $\tilde{\alpha}_n$, as follows:

$$\hat{\alpha}_n(p) = \begin{cases} \tilde{\alpha}_n(p) & \text{if } p \in (0,1) \\ k_n/\beta & \text{if } p = 0 \end{cases}$$

Because $\lim_{p \to 0+} \tilde{\alpha}_n(p) = k_n/\beta$, it follows that $\hat{\alpha}_n(p)$ satisfies (ODE') on $p \in [0,1]$ and is thus a solution to (BVP: $\beta \leq k_n$). Thus, (BVP: $\beta \leq k_n$) does not have a unique solution, a contradiction. Note that the argument in the case where $k_n < \beta$ is analogous.

It remains for us to establish that there exist unique solutions to both (BVP: $\beta \leq k_n$) and (BVP: $\beta \geq k_n$). We do this by invoking the Picard existence and uniqueness theorem, and thus begin by establishing that the right-hand side of (ODE') is Lipschitz continuous in $\alpha_n(p)$ and continuous in p for $p \in [-\epsilon, 1)$ and $\alpha_n(p) \in [c, 1+\epsilon]$ for any $c > 0$ and some $\epsilon > 0$. Since

$p^i \equiv \alpha_n(p) + (1 - \alpha_n(p))p$, it suffices to show that $\tilde{V}_{\cdot,n+1}$ is Lipschitz continuous in p^i for $p^i \geq 0$. In the case where $n = 1$, $\tilde{V}_{p^i,n+1} = 0$ for all p^i, and this is immediate. Next, suppose $n > 1$. First, consider the case where $k_{n+1} \geq \beta$. It follows from Lemma 8 that:

$$\tilde{V}_{p^i,n+1} = \begin{cases} k_n \alpha_{n+1}(p^i) - \beta(1 - p^i) & \text{if } p^i < 1 \\ 0 & \text{if } p^i > 1 \end{cases}$$

Because $\tilde{V}_{p^i,n+1}$ is continuously differentiable in p^i when $p^i \neq 1$, to establish that it is Lipschitz continuous it suffices to show that $\lim_{p^i \to 1-} V_{p^i,n+1} = 0$. Suppose this does not hold, by contradiction. Because $\alpha_{n+1}(\cdot)$ satisfies (ODE), this implies that $\lim_{p^i \to 1-} \alpha'_{n+1}(p^i) = \infty$. This in turn implies that $\lim_{p^i \to 1} \alpha_{n+1}(p^i) = \infty$, and thus that (ODE) is not satisfied at $p^i = 1$. Contradiction.

Next, consider the case where $k_{n+1} < \beta$. In this case:

$$\tilde{V}_{p^i,n+1} = \begin{cases} k_{n+1} p^i / n & \text{if } p^i < p^*_{n+1} \\ k_n \alpha_{n+1}(p^i) - \beta(1 - p^i) & \text{if } p^i \in (p^*_{n+1}, 1) \\ 0 & \text{if } p^i = 1 \end{cases}$$

By the reasoning from the case where $k_{n+1} \geq \beta$, $\tilde{V}_{p^i,n+1}$ is Lipschitz continuous for all $p^i > p^*_{n+1}$. Furthermore, Lipschitz continuity holds on $p^i < p^*_{n+1}$. To show that Lipschitz continuity holds across all p^i, it suffices to show that $\tilde{V}_{\cdot,n+1}$ is differentiable at p^*_{n+1}. To this end, we take the left- and right- derivative of $\tilde{V}_{\cdot,n+1}$ at p^*_{n+1} and show that they are equal:

$$\frac{d}{dp}\tilde{V}_{p^*-,n+1} = \frac{k_{n+1}}{N - n}$$

$$\frac{d}{dp}\tilde{V}_{p^*+,n+1} = -k_{n+1}\alpha'_{n+1}(p^*_{n+1}) + \beta = \frac{k_{n+1}}{1 - p^*_{n+1}} \frac{N - n}{N - n + 1} - \beta = \frac{k_{n+1}}{N - n}$$

Now, we show that there exists a unique solution for both (BVP: $\beta \leq k_n$) and (BVP: $\beta \geq k_n$) in some neighborhood of their respective boundary conditions. By the Picard Theorem, this follows immediately from our above-established result that the right-hand side of (ODE) is Lipschitz continuous in $\alpha_n(p)$ and continuous in p in some neighborhood of the boundary conditions $(\alpha_n(p) = 1, p = p^*)$ and $(\alpha_n(p) = \beta/k_n, p = 0)$.

Next, we seek to establish global existence and uniqueness of solutions to both (BVP: $\beta \leq k_n$) and (BVP: $\beta \geq k_n$). First, consider (BVP: $\beta \geq k_n$). The argument for (BVP: $\beta \leq k_n$) follows analogously. Let $[p^*, \bar{p})$ denote the largest right-open interval such that existence and uniqueness are both satisfied. Assume by contradiction that $\bar{p} < 1$. Let $\alpha_n(p)$ denote the solution along this interval.

We begin by showing that on this interval, $\alpha_n(p) \in (\underline{\alpha}, 1]$, where $\underline{\alpha} > 0$ is some constant. The upper bound is established as follows: suppose by contradiction that $\alpha_n(p) > 1$ somewhere on the interval. By the continuous differentiability of α_n along the interval, there must exist some $q < p$ such that $\alpha_n(q) = 1$ and $\alpha_n'(q) \geq 0$. However, it follows from (ODE') that

$$\alpha_n'(q) = -\frac{1}{k_n(1-q)} \frac{N-n}{N-n+1} [k_n - \tilde{V}_{p^i,n+1}] < 0$$

where the strict inequality follows from the fact that $\tilde{V}_{p^i,n+1} \leq k_{n+1} < k_n$. Contradiction. The lower bound is established as follows: suppose by contradiction that such a lower bound does not exist. Then, again by the continuous differentiability of α_n along the interval, there exists some $\hat{p} \in [p^*, \bar{p})$ such that

$$\lim_{p \to \hat{p}-} \alpha_n(p) = 0 \text{ and } \alpha_n(p) > 0 \text{ for all } p < \hat{p}$$

However, it then follows from (ODE) that $\lim_{p \to \hat{p}-} \alpha_n'(p) = \infty$. Thus, (ODE') is not satisfied on $[p^*, \bar{p})$. Contradiction.

Having established that on $[p^*, \bar{p})$, $1 \leq \alpha_n(p) > \underline{\alpha} > 0$, it follows from (ODE'), and the observation that $V_{p^i,n+1}$ is bounded, that α_n' is also bounded on this range. Thus, it follows that $\lim_{p \to \bar{p}-} \alpha_n(p) \equiv \bar{\alpha} > 0$ exists.

Now, consider the following modified boundary value problem, which is identical to (BVP: $\beta \geq k_n$), except with boundary condition $(\bar{p}, \bar{\alpha})$. by our prior-established result, we recall that (ODE') is

Lipschitz continuous in $\alpha_n(p)$ and continuous in p in some neighborhood of the boundary condition. Thus, we can again apply the Picard Theorem to obtain that there exists a unique solution to the modified boundary value problem in some neighborhood of $(\bar{p}, \bar{\alpha})$. Formally, there exists some $\epsilon > 0$ such that there is a unique solution $\tilde{\alpha}_n(p)$ on interval $(\bar{p} - \epsilon, \bar{p} + \epsilon)$. We can "paste" this solution $\tilde{\alpha}_n$, with our prior solution α_n. Formally, let

$$\hat{\alpha}_n(p) = \begin{cases} \alpha_n(p) & \text{if } p \in [p_n^*, \bar{p}) \\ \tilde{\alpha}_n(p) & \text{if } p \in [\bar{p}, \bar{p} + \epsilon) \end{cases}$$

Now, note that $\hat{\alpha}_n(p)$ is a unique solution to (BVP: $\beta \geq k_n$) on $[p_n^*, \bar{p} + \epsilon)$, which contradicts our earlier assumption that $[p^*, \bar{p})$ was the largest right-open interval such that existence and uniqueness are satisfied. Contradiction. □

Proof of Proposition 3. Let us begin by showing that $\alpha_n(p)$ is decreasing in p for all (p, n) on-path. By Lemma 10, it follows that when $k_N < \beta$, $\alpha_N(p) = 1$ for all p, and otherwise, $\alpha_N'(p) = 0$ for all p. Thus we have shown that $\alpha_N(p)$ is constant in p. Now, consider the case where $n < N$. Assume by induction that $\alpha_{n+1}(p)$ is weakly decreasing in p whenever $(p, n+1)$ is on path.

Assume by contradiction that there exists some $\bar{p}$ such that α_n is strictly increasing. Note that Lemma 10, $\alpha_n'(p) = 0$ whenever $\beta \geq k_n$ and $p < p_n^*$. Thus it must be that $\beta > k_n$ or $\bar{p} > p_n^*$. In this case, it again follows from Lemma 10 that (ODE) must be satisfied. Now define the function $X(p)$ as follows:

$$X(p) \equiv k_n \alpha_n(p) - \beta(1 - p^i) - V_{p^i, n+1} \tag{A.19}$$

Note that whenever (ODE) is satisfied, the following holds:

$$\alpha_n'(p) > (=) 0 \text{ if and only if } X(p) < (=) 0 \tag{A.20}$$

Thus, $X(\bar{p}) < 0$. Now, I claim that there exists $\underline{p} < \bar{p}$ such that $\lim_{p \to \underline{p}+} X(p) \geq 0$. To establish this, first consider the case where $k_n \geq \beta$. In this case,

$$\lim_{p\to 0+} X(p) = k_n \lim_{p\to 0+} \alpha_n(p) - \beta(1 - \lim_{p\to 0+}) - \lim_{p\to 0+} V_{\alpha_n(p),n+1} = (k_n + \beta) \lim_{p\to 0+} \alpha_n(p) - \beta - \lim_{p\to 0+} V_{\alpha_n(p),n+1}$$

When $\lim_{p\to 0+} \alpha_{n+1}(\alpha_n(p)) < 1$, it follows from Lemma 8 that

$$\begin{aligned}\lim_{p\to 0+} V_{\alpha_n(p),n+1} = \lim_{p\to 0+} V_{\alpha_n(p),n+1}(\delta_0) &= k_{n+1} \lim_{p\to 0+} \alpha_{n+1}(\alpha_n(p)) - \beta(1 - \lim_{p\to 0+} \alpha_n(p)) \\ &= k_{n+1}\alpha_{n+1}(\beta/k_n) - \beta(1 - \beta/k_n)\end{aligned}$$

Because $k_n \geq \beta$, the final equality follows from Lemma 10. Substituting this into our above expression for $\lim_{p\to 0+} X(p)$, we obtain

$$\lim_{p\to 0+} X(p) = \beta - k_{n+1}\alpha_{n+1}(\beta/k_n)$$

In the case where $k_{n+1} < \beta$, it follows directly that $\lim_{p\to 0+} X(p) > 0$. Otherwise, if $k_{n+1} \geq \beta$, then because $\lim_{p\to 0+} \alpha_{n+1}(p) = \beta/k_{n+1}$, it follows from the inductive assumption that $\alpha_{n+1}(p) \leq \beta/k_{n+1}$ for all p, and thus that $\lim_{p\to 0+} X(p) > 0$.

Meanwhile, when $\lim_{p\to 0+} \alpha_{n+1}(\alpha_n(p)) = 1$, it follows from the inductive assumption that $\alpha_{n+1}(q) = 1$ for all $q \geq \lim_{p\to 0+} \alpha_n(p)$. It thus follows that

$$\lim_{p\to 0+} V_{p^i,n+1} = \lim_{p\to 0+} V_{p^i,n+1}(\delta_\infty) = \frac{k_{n+1}}{N-n}\frac{\beta}{k_n}$$

Substituting into the above expression for $\lim_{p\to 0+} X(p)$ and simplifying, we obtain

$$\lim_{p\to 0+} X(p) = (\beta/k_n)(\beta - k_{n+1}/(N-n)) \geq 0,$$

where the inequality follows from the fact that $\alpha_{n+1}(\beta/k_n) = 1$, implying by Lemma 10 that $k_{n+1} \geq \beta$.

Next, consider the case where $k_n < \beta$. In this case,

$$\lim_{p \to p_n^* +} X(p) = k_n - \lim_{p^i \to 1-} V_{p^i, n+1}$$

If $\lim_{p^i \to 1-} \alpha_{n+1}(p^i) < 1$, then by Lemma 8,

$$\lim_{p^i \to 1-} V_{p^i,n+1} = \lim_{p^i \to 1-} V_{p^i,n+1}(\delta_0) = k_{n+1} \lim_{p^i \to 1-} \alpha_{n+1}(p^i) < k_n.$$

Thus, in this case, we obtain that $\lim_{p \to p_n^*+} X(p) > 0$. Meanwhile, if $\lim_{p^i \to 1-} \alpha_{n+1}(p^i) = 1$, by the inductive assumption, $\alpha_{n+1}(p) = 1$ for all p. Thus,

$$\lim_{p^i \to 1-} V_{p^i,n+1} = \lim_{p^i \to 1-} V_{p^i,n+1}(\delta_\infty) = \lim_{p^i \to 1-} \frac{k_{n+1} p^i}{N-n} = \frac{k_{n+1}}{N-n}.$$

We once again obtain $\lim_{p \to p_n^*+} X(p) > 0$. We have thus shown that there always exists $\underline{p} < \bar{p}$ such that $\lim_{p \to \underline{p}+} X(p) \geq 0$.

Because $X(\bar{p}) < 0$ by assumption, there must exist some $q \in [\underline{p}, \bar{p}]$ $X(q) < 0$ and $X'(q) < 0$. Note that differentiating our above expression for X, we have

$$X'(q) = k_n \alpha_n'(q) + \beta((1-q)\alpha_n'(q) + (1-\alpha_n(q))) - \frac{d}{dq} V_{\alpha_n(q)+(1-\alpha_n(q))q,n+1}. \tag{A.21}$$

First, consider the case where $\alpha_{n+1}(q^i) < 1$. By Lemma 8,

$$V_{q^i,n+1} = V_{q^i,n+1}(\delta_0) = k_{n+1}\alpha_{n+1}(q^i) - \beta(1-q^i).$$

Substituting this into (A.21), we obtain

$$X'(q) = k_n \alpha_n'(q) - k_{n+1}\alpha_{n+1}'(q^i)((1-q)\alpha_n'(q) + (1-\alpha_n(q))).$$

Note that because $X(q) < 0$ it follows from (A.20) that $\alpha_n'(q) > 0$. Furthermore, by the inductive assumption, $\alpha'_{n+1}(q^i) \leq 0$. Thus, in this case, $X'(q) > 0$. Contradiction.

Next, consider the case where $\alpha_{n+1}(q^i) = 1$. By the inductive assumption, $\alpha_{n+1}(p) = 1$ for all $p \leq q^i$. Thus,

$$V_{q^i,n+1} = V_{q^i,n+1}(\delta_\infty) = \frac{k_{n+1}q^i}{N-n}.$$

Substituting this into (A.21), we obtain

$$X'(q) = k_n\alpha'_n(q) + (\beta - \frac{k_{n+1}}{N-n})((1-q)\alpha'_n(q) + (1-\alpha_n(q)))$$

Because $\alpha_{n+1}(q^i) = 1$, by Proposition 1 (if $n < N-1$) and Lemma 10 (if $n = N-1$), it must be that $\beta \geq k_{n+1}$. Thus, $X'(q) > 0$. Contradiction.

Next, we will show that if $k_N \geq \beta$, then $\alpha_n(p) = \beta/k_n$. Assume that $k_N \geq \beta$. First consider the case where $n = N$. By Lemma 10, $\alpha'_n(p) = 0$ for all p on-path, and thus, $\alpha_N(p)$ is constant in p. Since Lemma 10 also asserts that $\lim_{p\to 0+} k_N\alpha_N(p) = \beta$, it must be that $\alpha_N(p) = \beta/k_N$ for all p. Now, consider $n < N$. Assume by induction that $\alpha_{n+1}(p) = \beta/k_{n+1}$ for all p. We begin by showing that $\alpha_n(p)$ is constant in p. Since $k_n > \beta$, by Lemma 10, (ODE) must hold at all p. By (A.20), showing $\alpha_n(p)$ is constant in p is equivalent to showing that $X(p) = 0$. To establish this, I begin by claiming that $V_{p^i,n+1} = V_{p^i,n+1}(\delta_0)$. In the case where $k_{n+1} > \beta$, it follows from Proposition 1 that $\alpha_{n+1}(p^i) < 1$, and thus this follows from Lemma 8. In the case where $k_{n+1} = \beta$, because $k_m > k_N \geq \beta$ for all $m < N$, it follows that $n+1 = N$. In this case, all pure strategies δ_s must yield the same value. In particular, for all $s \in [0,\infty]$, $V_{p,N}(\delta_s) = k_N p$. Thus, δ_0 is optimal. Having established that $V_{p^i,n+1} = V_{p^i,n+1}(\delta_0)$, we have:

$$V_{p^i,n+1} = k_{n+1}\alpha_{n+1}(p^i) - \beta(1-p^i) = \beta p^i$$

Substituting this into (A.19), we obtain $X(p) = k_n\alpha_n(p) - \beta$. Since we established above that $\alpha_n(p)$ is weakly decreasing, $\alpha_n(p) \leq k_n/\beta$ for all p, and thus $X(p) \leq 0$. Separately, by (A.20) $\alpha_n(p)$ weakly decreasing implies that $X(p) \geq 0$. Combining these inequalities, we have $X(p) = 0$.

Finally, I will show that $k_N < \beta$ implies that $\alpha_n'(p) < 0$ whenever $\alpha_n(p) < 1$. To this end, suppose $k_N < \beta$, and suppose by contradiction that at some q such that $\alpha_n(q) < 1$, $\alpha_n'(q) = 0$. It follows from (A.20) that $X(q) = 0$.

First, consider the case where $\alpha_{n+1}(q^i) = 1$. Recall from the above that in this case, we have

$$X'(q) = k_n\alpha_n'(q) + (\beta - \frac{k_{n+1}}{N-n})((1-q)\alpha_n'(q) + (1-\alpha_n(q)) = (\beta - \frac{k_{n+1}}{N-n})(1-\alpha_n(q)) \quad \text{(A.22)}$$

Now, I claim that $\beta > \frac{k_{n+1}}{N-n}$. In the case where $n = N-1$, this follows directly from our assumption that $k_N < \beta$. Meanwhile, in the case where $n < N-1$, because $\alpha_{n+1}(q^i) = 1$, this is a result of Proposition 1. It thus follows from (A.22) that $X'(q) > 0$. Since $X(q) = 0$, for some $p < q$, we must have $X(p) < 0$. By (A.21), $\alpha_n'(p) > 0$. This contradicts the above-established assertion that $\alpha_n(p)$ is weakly decreasing in p.

Next, consider the case where $\alpha_{n+1}(q^i) < 1$. As established above, in this case:

$$X'(q) = k_n\alpha_n'(q) - k_{n+1}\alpha_{n+1}'(q^i)[(1-q)\alpha_n'(q) + (1-\alpha_n(q))] = -k_{n+1}\alpha_{n+1}'(q)[1-\alpha_n(q)] > 0.$$

Again, this implies that there exists some $p < q$ such that $X(p) < 0$ and thus that $\alpha'(p) > 0$. Contradiction. □

A.3 Commitment solution

Here, we seek the optimal solution to the monopoly case of the baseline model in which the firm has the ability to commit to a reporting strategy. Formally, the only modification we introduce is that rather than F and α being determined simultaneously as they are in equilibrium, the firm chooses its strategy F before α is determined. Thus, in the commitment case, the credibility function is a function of the firm's strategy. We formalize this dependence by denoting the firm's credibility function as α_F. α_F is then given by (1.3) as in the non-commitment case, except that the strategy F upon which it is computed is the firm's choice of strategy, rather than the equilibrium strategy.

Because we are considering the monopoly case only, I will be dropping the n index from all functions and expressions. Furthermore, for convenience, I will be writing all functions as a function of calendar time t, rather than the common belief p as in the baseline model. Writing the functions in this way is without loss, since under a monopoly there is a one-to-one correspondence between the calendar time t and the common belief p.

The firm's objective is to choose a permissible strategy $F \in \mathcal{F}$ which maximizes its expected payoff over the course of the game. Specifically, its problem is given by the following:

$$\max_{F \in \mathcal{F}} \int_0^\infty [\alpha_F(t) - \beta(1 - p(t))(1 - \alpha_F(t))] d\Psi(t) \tag{A.23}$$

where, as in the baseline setup, $\Psi(t)$ denotes probability that the firm reports before time t under strategy F. Before proceeding, we highlight that the only difference between this problem and the problem of the monopoly case of the baseline model is that the credibility function is not taken as given, but is rather a function of the firm's choice of strategy F.

In the analysis that follows, it will be useful for us to cast this problem as a choice of an optimal credibility function α, rather than an optimal strategy F. To this end, I begin with a useful observation, which is analogous to Lemma 1, except under the commitment setting:

Lemma 11. *F must be continuous in equilibrium.*

We omit a proof for this claim, as it follows analogously to the proof for Lemma 1: if F exhibits a discontinuity at some time t, reporting at this time must yield a negative expected payoff. Thus, the firm can profitably deviate by shifting the mass that it had placed on reporting t to δ_∞.

It follows immediately from Lemma 11 that in equilibrium, both the firm's strategy F and the corresponding commitment function, α_F, are defined by the right-hazard rate $b(t)$ of the firm's strategy. That is,

$$\alpha_F(t) = \frac{\lambda p(t)}{\lambda p(t) + b(t)}$$

It further follows that Ψ is continuous and can thus be written as a function of α_F as follows:

$$\Psi(t) = 1 - e^{-\int_0^t (b(s)+p(s)\lambda)ds} = 1 - e^{-\int_0^t \frac{\lambda p(s)}{\alpha_F(s)}ds}$$

Having written Ψ in terms of α_F, and noting that at any given t $\alpha_F(t)$ is a one-to-one function of $b(t)$, we can cast the optimization problem given by (A.23) as one over α_F:

$$\max_{\alpha_F} \int_0^\infty \lambda p(t)[1 - \beta(1 - p(t))(\frac{1}{\alpha_F(t)} - 1)]e^{-\int_0^t \frac{\lambda p(s)}{\alpha_F(s)}ds}$$

In the following claim, I show that the optimal strategy for the firm consists of always truth telling (i.e., $\alpha_F(t) = 1$ for all t). In the proof that follows, I let $V(t, \alpha_F)$ denote the firm's value at time t given that it has chosen α_F.

Proposition 18. *In equilibrium, $\alpha_F(t) = 1$ for all t.*

Proof. Assume not, by contradiction. Then there exists a t^* such that $\alpha_F(t^*) < 1$. It follows from Lemma 11, and the assumption that F is right-continuously differentiable, that α_F must be right-continuous. Thus, there must exist a $\hat{\alpha} < 1$ and $\epsilon > 0$ such that $\alpha_F(t) < \hat{\alpha}$ for all $t \in [t^*, t^* + \epsilon]$.

Note that for any α_F, including the equilibrium α_F, we can write the time-0 value as follows:

$$V(0, \alpha_F) = \int_0^{t^*+\epsilon} \lambda p(t)[1 - \beta(1 - p(t))](\frac{1}{\alpha_F(t)} - 1)e^{-\int_0^t \frac{\lambda p(s)}{\alpha_F(s)}ds}dt + e^{-\int_0^{t^*+\epsilon} \frac{\lambda p(s)}{\alpha_F(s)}ds}V(t^* + \epsilon, \alpha_F) \tag{A.24}$$

Now, consider the following deviation $\tilde{\alpha}_F$, which is identical to α_F, except that it is 1 on the interval $[t^*, t^* + \epsilon]$:

$$\tilde{\alpha}_F(t) = \begin{cases} 1 & \text{if } t \in [t^*, t^* + \epsilon] \\ \alpha_F(t) & \text{otherwise} \end{cases}$$

Now, it follows from (A.24) that

$$V(0,\alpha_F) = V(0,\tilde{\alpha_F}) + \int_{t^*}^{t^*+\epsilon} \lambda p(t)[1-\beta(1-p(t))(\frac{1}{\alpha_F(t)}-1)]e^{-\int_0^t \frac{\lambda p(s)}{\alpha_F(s)}ds}dt - \int_{t^*}^{t^*+\epsilon} \lambda p(t)e^{-\int_0^t \lambda p(s)ds}dt$$
$$+(e^{-\int_{t^*}^{t^*+\epsilon} \frac{\lambda p(s)}{\alpha_F(s)}ds} - e^{-\int_{t^*}^{t^*+\epsilon} \lambda p(s)ds})V(t^*+\epsilon,\alpha_F) \tag{A.25}$$

Now, we will note the following two inequalities:

$$\int_{t^*}^{t^*+\epsilon} \lambda p(t)[1-\beta(1-p(t))(\frac{1}{\alpha_F(t)}-1)]e^{-\int_0^t \frac{\lambda p(s)}{\alpha_F(s)}ds}dt \leq \int_{t^*}^{t^*+\epsilon} \lambda p(t)[1-\beta(1-p(t))(\frac{1}{\bar{\alpha}}-1)]e^{-\int_0^t \frac{\lambda p(s)}{\bar{\alpha}}ds}dt$$
$$< \int_{t^*}^{t^*+\epsilon} \lambda p(t)e^{-\int_0^t \lambda p(s)ds}dt$$

$$e^{-\int_{t^*}^{t^*+\epsilon} \frac{\lambda p(s)}{\alpha_F(s)}ds} - e^{-\int_{t^*}^{t^*+\epsilon} \lambda p(s)ds} \leq e^{-\int_{t^*}^{t^*+\epsilon} \frac{\lambda p(s)}{\bar{\alpha}}ds} - e^{-\int_{t^*}^{t^*+\epsilon} \lambda p(s)ds} < 0$$

Applying these two inequalites to (A.25) we obtain

$$V(0,\alpha_F) < V(0,\tilde{\alpha}),$$

and thus, $\tilde{\alpha}_F$ serves as a profitable deviation. Contradiction. □

A.4 Proofs: comparative statics

Proof of Proposition 4. First, we establish part (a). Fix all other parameters and let $0 < \beta < \tilde{\beta}$. Let α and $\tilde{\alpha}$ denote the equilibrium credibility functions under β and $\tilde{\beta}$, respectively. Fix an n and assume inductively that the proposition holds for $n+1$ if $n < N$. Note that for any (p,n) and t, $p(t)$ will be the same under β and $\tilde{\beta}$. Thus to show the above claim, it suffices to show that for any p, $\alpha_n(p)$ is weakly increasing in β, and strictly so whenever $\alpha_n(p) < 1$.

We begin by showing that $\alpha_n(p) = 1$ implies that $\tilde{\alpha}_n(p) = 1$. First, consider the case where $n = N$. By Proposition 2, $\alpha_N(p) = 1$ implies that $k_N \leq \beta$. Thus, $k_N < \tilde{\beta}$, which by Proposition 1 implies that $\tilde{\alpha}_N(p) = 1$. Next, consider the case where $n < N$, and assume $\alpha_n(p) = 1$. By

Proposition 1, this implies that $k_n < \beta$ and $p \leq p_n^* \equiv \frac{\beta - k_n}{\beta - k_n/n}$. Further note that

$$\tilde{p}_n^* \equiv \frac{\tilde{\beta} - k_n}{\tilde{\beta} - k_n/n} > \frac{\beta - k_n}{\beta - k_n/n} \equiv p_n^*.$$

Thus, $k_n < \tilde{\beta}$ and $p < \tilde{p}_n^*$, which by Proposition 1 implies $\tilde{\alpha}_n(p) = 1$.

Now, suppose that $\alpha_n(p) < 1$. We wish to show that $\tilde{\alpha}_n(p) > \alpha_n(p)$. Suppose by contradiction that $\tilde{\alpha}_n(p) \leq \alpha_n(p)$. It follows from Proposition 2 that if $k_n > \tilde{\beta}$,

$$\lim_{q \to 0+} \alpha_n(q) = \beta/k_n < \tilde{\beta}/k_n = \lim_{q \to 0+} \tilde{\alpha}_n(q).$$

Meanwhile, if $k_n \leq \tilde{\beta}$.

$$\lim_{q \to \tilde{p}_n^*+} \alpha_n(q) < 1 = \lim_{q \to \tilde{p}_n^*+} \tilde{\alpha}_n(q)$$

To see why the latter must must hold, first consider the case where $n = 1$. It follows from Lemma 10 that $\tilde{\alpha}_n(q) = 1$ for all q. Meanwhile, it follows again from Proposition 2 that $\alpha_N(q)$ is constant in q, and because $\alpha_N(p) < 1$, $\lim_{q \to \tilde{p}_n^*+} \alpha_N(q) < 1$. In the case where $n < N$, because $p_n^* < \tilde{p}_n^*$, it follows from Proposition 1 that $\alpha_n(\tilde{p}_n^*) < 1$.

Thus, we have that both when $k_n > \tilde{\beta}$ and when $k_n \leq \tilde{\beta}$, there exists some $\hat{p} < p$ such that $\tilde{\alpha}_n(\hat{p}) > \alpha_n(\hat{p})$ and $\tilde{\alpha}_n$, α_n satisfy (ODE) on $[\hat{p}, p]$, for their respective value of β. Thus, there exists a $q \in [\hat{p}, p]$ such that $\alpha_n(q) = \tilde{\alpha}_n(q)$ and $\alpha_n'(q) \geq \tilde{\alpha}_n'(q)$. It follows from (ODE) that in order for the above two conditions to hold, it must be that

$$X \equiv (\beta - \tilde{\beta})(\frac{1 - \alpha_n(q)}{\alpha_n(q)})(1 - q) + \frac{V_{q^i,n+1} - \tilde{V}_{q^i,n+1}}{\alpha_n(q)} \geq 0 \tag{A.26}$$

where V and $\tilde{V}$ denote the value functions under β and $\tilde{\beta}$, respectively. First consider the case where $n = N$. Then $V_{q^i,n+1} = V_{\tilde{q}^i,n+1} = 0$, and thus $X < 0$, contradicting (A.26).

Next, consider the case where $n < N$. First suppose that $\alpha_{n+1}(q^i) = 1$. It follows from the inductive assumption that $\tilde{\alpha}_{n+1}(q^i) = 1$. Thus, by Lemma 10, $V_{q^i,n+1} = \frac{k_{n+1}q^i}{N-n} = \tilde{V}_{q^i,n+1}$. Again this implies that $X < 0$, contradicting (A.26). Now, suppose that $\alpha_{n+1}(q^i) < 1$. It then follows from

Lemma 8 that

$$V_{q^i,n+1} = V_{q^i,n+1}(\delta_0) = k_{n+1}\alpha_{n+1}(q^i) - \beta(1 - q^i)$$

Furthermore,

$$\tilde{V}_{q^i,n+1} = \tilde{V}_{q^i,n+1}(\delta_0) = k_{n+1}\tilde{\alpha}_{n+1}(q^i) - \tilde{\beta}(1 - q^i)$$

Thus, recalling from (A.2) that $q^i = \alpha_{n+1}(q) + (1 - \alpha_{n+1}(q))q$, we have

$$V_{q^i,n+1} - \tilde{V}_{q^i,n+1} \leq k_{n+1}(\alpha_{n+1}(q^i) - \tilde{\alpha}_{n+1}(q^i))$$

Substituting this into the above expression for X, we obtain

$$X \leq \frac{k_{n+1}(\alpha_{n+1}(q^i) - \tilde{\alpha}_{n+1}(q^i))}{\alpha_n(q)} < 0.$$

where the strict inequality follows from the inductive assumption that $\alpha_{n+1}(q^i) < \tilde{\alpha}_{n+1}(q^i))$. Again, this is a contradiction of (A.26).

Next, let us establish part (b). Let $\tilde{\lambda} > \lambda > 0$, and let α, $\tilde{\alpha}$ denote the equilibria under λ and $\tilde{\lambda}$, respectively, fixing all other parameters. We begin by showing that $\tilde{\alpha}_n(p) = \tilde{\alpha}_n(p)$ for any p and n. Fix an n and assume inductively that if $n < N$, $\alpha_{n+1}(p) = \tilde{\alpha}_{n+1}(p)$ for all p on-path.

Let V, $\tilde{V}$ denote the value functions under the equilibria associated with λ and $\tilde{\lambda}$, respectively. Note that $V_{p,n+1} = \tilde{V}_{p,n+1}$ for all p on-path. In the case where $n = N$, $V_{p,n+1} = \tilde{V}_{p,n+1} = 0$, and thus this holds trivially. In the case where $n < N$, this follows from the inductive assumption.

Now, note that by Lemma 10, α_n and $\tilde{\alpha}_n$ must both be a solution to (P) at all (p, n) on-path, which does not depend on λ. By Theorem 1, the solution to (P) is unique, and thus $\alpha_n(p) = \tilde{\alpha}_n(p)$ at all (p, n) on-path.

Now fixing any p and n, let $p(t)$ and $\tilde{p}(t)$ denote the common beliefs under λ and $\tilde{\lambda}$, respectively. It follows from (A.1) that $p(t) > \tilde{p}(t)$ for all $t > 0$. Thus, because $\alpha_n(p)$ and $\tilde{\alpha}_n(p)$ are both weakly decreasing in p (Proposition 3), it follows that $\alpha_n(p(t)) \leq \tilde{\alpha}_n(p(t))$. Furthermore, since $\tilde{\alpha}(p)$ is strictly decreasing in p (Proposition 3) whenever $\alpha_n(p) < 1$ and $k_N > \beta$, it follows that $\alpha_n(p(t)) < \alpha_n(\tilde{p(t)})$ in this case.

Finally, let us establish part (c). Let α and $\tilde{\alpha}$ denote the equilibria under N and $N+1$ firms, respectively, fixing all other parameters. We begin by showing that for all p, $\alpha_n(p) \geq \tilde{\alpha}_n(p)$, and $\alpha_n(p) > \tilde{\alpha}_n(p)$ when $\alpha_n(p) < 1$. To this end, fix an $n \in \{1, ..., N\}$ and assume inductively that the claim holds for $n+1$ whenever $n < N$.

We begin by showing that $\tilde{\alpha}_n(p) = 1$ implies that $\alpha_n(p) = 1$. Suppose that $\tilde{\alpha}_n(p) = 1$. By Proposition 1, $\beta > k_n$ and $p < \tilde{p}_n^* \equiv \frac{\beta - k_n}{\beta - k_n/(N+1-n)}$. Because $p_n^* \equiv \frac{\beta - k_n}{\beta - k_n/(N-n)} > \tilde{p}_n^*$, it follows from Proposition 1 that $\alpha_n(p) = 1$.

Now consider the case where $\tilde{\alpha}_n(p) < 1$. We wish tot show that $\tilde{\alpha}_n(p) < \alpha_n(p)$. To this end, we begin by making the following observation:

$$\text{If } \alpha_n \text{ and } \tilde{\alpha}_n \text{ both satisfy (ODE) at } q, \text{ and } \alpha_n(q) = \tilde{\alpha}_n(q), \text{ then } \alpha_n'(q) > \tilde{\alpha}_n'(q). \tag{A.27}$$

Let us now establish this. Note first that for α_n and $\tilde{\alpha}_n$ to both satisfy (ODE) at q, given that $\alpha_n(q) = \tilde{\alpha}_n(q)$, the following must hold:

$$\alpha_n'(q) = \frac{-1}{k_n(1-q)\alpha_n(q)} \frac{N-n}{N-n+1}(k_n\alpha_n(q) - V_{q^i,n+1} - \beta(1-\alpha_n(q))(1-q))$$

$$\tilde{\alpha}_n'(q) = \frac{-1}{k_n(1-q)\alpha_n(q)} \frac{N-n+1}{N-n+2}(k_n\alpha_n(q) - \tilde{V}_{q^i,n+1} - \beta(1-\alpha_n(q))(1-q)),$$

where V and $\tilde{V}$ denote the value functions under the equilibria with N and $N+1$ total firms, respectively. Note that if $n = N$, $\alpha_n'(q) = 0$. Meanwhile, by Proposition 3, $\tilde{\alpha}_n'(q) < 0$. Thus, $\tilde{\alpha}_n'(q) < \alpha_n(q)$ must hold. Next, consider the case where $n < N$. We begin by observting that $V_{q^i,n+1} > \tilde{V}_{q^i,n+1}$. To see why this must hold, first consider the case where $\tilde{\alpha}_{n+1}(q^i) = 1$. It then follows from the inductive assumption that $\alpha_n(q^i) = 1$. Then, by Lemma 10,

$$\tilde{V}_{q^i,n+1} = \tilde{V}_{q^i,n+1}(\delta_\infty) = \frac{k_{n+1}q^i}{N-n} < \frac{k_{n+1}q^i}{N-n-1} = V_{q^i,n+1}(\delta_\infty) = V_{q^i,n+1}.$$

Next, consider the case where $\tilde{\alpha}_n(q^i) < 1$. In this case, it follows from Lemma 8 that

$$\begin{aligned}\tilde{V}_{q^i,n+1} = \tilde{V}_{q^i,n+1}(\delta_0) = k_{n+1}\tilde{\alpha}_{n+1}(q^i) - \beta(1-q^i) < k_{n+1}\alpha_{n+1}(q^i) - \beta(1-q^i) \\ = V_{q^i,n+1}(\delta_0) \leq V_{q^i,n+1}\end{aligned}$$

where the strict inequality follows from the inductive assumption made above. Examining the two ODEs listed above, since by Proposition 3, $\alpha_n'(q) \leq 0$, it follows that $\tilde{\alpha}_n'(q) < \alpha_n'(q)$.

Now, assume by contradiction that $\alpha_n(p) \leq \tilde{\alpha}_n(p)$. We begin by showing that there exists a $q^* < p$ such that $\tilde{\alpha}_n(q^*) < \alpha_n(q^*)$. First consider the case where $k_n \geq \beta$. Then, by Proposition 2,

$$\lim_{q\to 0+} \alpha_n(q) = \lim_{q\to 0+} \tilde{\alpha}_n(q) = \frac{\beta}{k_n}$$

Then, by the continuous differentiability of α_n and $\tilde{\alpha}_n$ on $(0, p)$, it follows from Equation A.27 that for some $q^* < p$ sufficiently small $\alpha_n(q^*) > \tilde{\alpha}_n(q^*)$. Next, consider the case where $k_n < \beta$, and let $p_n^* \equiv \frac{\beta - k_n}{\beta/(N-n+1)-k_n}$. Note by Proposition 1 that $\alpha_n(p_n^*) = 1$. Meanwhile, because $p_n^* < \tilde{p}_n^* \equiv \frac{\beta-k_n}{\beta/(N-n+2)-k_n}$, it follows from Proposition 1 that $\tilde{\alpha}_n(p_n^*) < 1$, and thus, we have for $q^* = p_n^*$, $\tilde{\alpha}_n(q^*) < \alpha_n(q^*)$.

Since $\tilde{\alpha}_n(q^*) < \alpha_n(q^*)$ and $\tilde{\alpha}_n(p) \geq \alpha_n(p)$, by the continuous differentiability of α on $[q^*, p]$, there must exist some $q \in (q^*, p]$ such that $\alpha_n(q) = \tilde{\alpha}_n(q)$ and $\alpha_n'(q) \leq \tilde{\alpha}_n'(q)$. However, this is a contradiction of (A.27).

Now fixing any p and n, let $p(t)$ and $\tilde{p}(t)$ denote the common beliefs under N and $N+1$ firms, respectively. We wish to show that on some interval $[0, \bar{t}]$, where $\bar{t} > 0$, $\alpha_n(p(t)) \geq \tilde{\alpha}_n(\tilde{p}(t))$ is weakly increasing in t, and strictly so whenever $\alpha_n(p(t)) < 1$. First consider the case where $\alpha_n(p(t)) = 1$. In this case, the statement holds trivially. Next, consider the case where $\alpha_n(p) < 1$. It follows from the above that $\alpha_n(p) > \tilde{\alpha}_n(p)$. Now note that it follows from (A.1) that $\lim_{t\to 0+} p(t) - \tilde{p}(t) = 0$. Since $\alpha_n(p(t))$ and $\tilde{\alpha}_n(\tilde{p}(t))$ are both continuous in t (Lemma 9), it follows that for some $\bar{t} > 0$, $\alpha_n(p(t)) > \tilde{\alpha}_n(\tilde{p}(t))$ for all $t \in [0, \bar{t}]$. □

A.5 Proofs: heterogenous learning abilities

Here, we consider the extended model presented in Section 1.6. The objective is to establish Proposition 5. This proof will require extending certain results established in the baseline model to the extended model.

Regarding Lemmas 1-4, I will take for granted that these hold under the extended model. Formal proofs of this are omitted as all proofs presented under the baseline model will apply to the extended setting as well.

Next, I establish that Proposition 1 holds under the extended model. This claim is presented as Proposition 1'. In the analysis below, I let $V^i_{p,n}$ denote firm i's value.

Proposition 1'. *For all s, there exists a* $p^{i*} \in (0,1]$ *such that at any p on-path,* $\alpha^i_1(p) = 1$ *if and only if the following two conditions hold:*

1. $k_1 \leq \beta$
2. $p \leq p^{i*}$

Furthermore, $p^{j*} > p^{i*}$ *whenever* $\lambda^j > \lambda^i$ *and* $n > 1$.

Proof. Fix an i. Suppose that $k_1 \leq \beta$. By identical reasoning as Proposition 1, for all $q < \frac{\beta - k_1}{k_1}$, $\alpha^i_1(q) = 1$. Let

$$p^{i*} \equiv sup\{p|\alpha^i_1(p) = 1 \text{ for all } q < p\}$$

It follows by definition that $\alpha^i_1(p) = 1$ for all $p \leq p^{i*}_1$.

Next, we will show that $\alpha^i_1(q) < 1$ whenever $k_1 > \beta$ or $p > p^{i*}_1$. Suppose not by contradiction. First, consider the case where $k_1 > \beta$ and $\alpha^i_1(p) = 1$ for some p. Then we have that

$$V^i_{p,1}(\delta_0) = k_1 p + (k_1 - \beta)(1-p) > k_1 p \leq V^i_{p,1}(\delta_\infty)$$

Thus, i can profitably deviate at p. Contradiction. Next, consider the case where $q > p^{i*}_n$ and $\alpha^i_1(p) = 1$. In this case, a contradiction follows from identical reasoning to what is presented in

Proposition 1.

Finally, we show that $p^{j*} > p^{i*}$ whenever $\lambda^j > \lambda^i$. Suppose by contradiction that $p^{j*} \leq p^{i*}$. Note that because j is truth telling at $(p_S^{j*}, n = 1)$, $V^j_{p_1^{j*},1}(\delta_\infty) \geq V^j_{p^{j*},1}(\delta_0)$. Furthermore, because $p^{j*} \leq p^{i*}$, i is also truthful at $(p_n^{j^*}, n = 1)$. Thus,

$$V^j_{p_1^{j*},1}(\delta_0) = V^i_{p_1^{j*},1}(\delta_\infty) = k_1 - \beta(1-p).$$

Now, note that because $\lambda^j > \lambda^i$,

$$V^j_{p_1^{j*},1}(\delta_\infty) > V^i_{p_1^{j*},1}(\delta_\infty).$$

Combining these inequalities we have $V^i_{p_1^{j*},1}(\delta_\infty) < V^i_{p_1^{j*},1}(\delta_0)$. However, because $\alpha_1^i(p^{j*}) = 1$, $V^j_{p_n^{j*},1} = V^j_{p_n^{j*},1}(\delta_\infty)$. Contradiction. □

Next, we extend Proposition 2 to this setting. Note this entails deriving an ODE that applies to this extended model, (ODE').

Proposition 2'. *In equilibrium, for any p on-path, if $k_1 \geq \beta$ or $p > p^{i*}$, then the following must be satisfied:*

$$\alpha_1^{i\prime}(p) = -\beta - \frac{\sum_{j \neq i} \frac{\lambda^j}{\alpha_1^j(p)}}{\sum_j \lambda^j (1-p)} [\alpha^i(p) - \beta(1-p)] \qquad \text{(ODE')}$$

In addition, $\lim_{p \to 0+} \alpha_1^i(p) = \beta / k_1$ *must hold if $k_1 > \beta$, and* $\lim_{p \to p^{i*}+} \alpha_1^i(p) = 1$ *if $k_1 \leq \beta$.*

Proof. Let us first establish that (ODE') must hold under the conditions specified.

When $k_1 \geq \beta$ or $p > p^{i*}$, it follows from Proposition 1' that $\alpha_1^i(p(t)) < 1$. It then follows from Lemma 8 that there exsits an $\epsilon > 0$ such that for all $\Delta \in (0, \epsilon)$,

$$\frac{V^i_{p,1}(\delta_\Delta) - V^i_{p,1}(\delta_0)}{\Delta} = 0$$

Recall that $V^i_{p,1}(\delta_0) = k_1\alpha^i_1(p) - \beta(1-p)$. Meanwhile,

$$V^i_{p,1}(\delta_\Delta) = \int_0^\Delta k_1\alpha^i_1(p(s))\Psi^i(s)ds + (1 - \sum_j \lim_{s\to\Delta-} \Psi^j(s))[k_1\alpha_1(p(\Delta)) - \beta(1-p(\Delta))]$$

where Ψ is the first-report distribution associated with the strategy profile in which i plays δ_∞ and all $j \neq i$ play $F_{p,1}$. Specifically, for all $s > 0$,

$$\Psi^i(s) = p\lambda^i \int_0^s e^{-\prod_{j\in S}\lambda^j r} \prod_{j\neq i}(1 - F^i_{p,1}(r))dr$$

and for $j \neq i$,

$$\Psi^j(s) = p\int_0^s e^{-\prod_{k\neq j}\lambda^k r} \prod_{k\neq i\neq j}(1 - F^k_{p,1}(r))d(-e^{-\lambda^j r}(1 - F^j_{p,1}(r))) \\ +(1-p)\int_0^s \prod_{k\neq i\neq j}(1 - F^k_{p,1}(r))dF^j_{p,1}(r)$$

Substituting these two expressions into the above equation for $V^i_{p,1}(\delta_0)$ and following the same sequence of steps in Proposition 2 yields (ODE').

The two limit conditions are established by the same reasoning presented in Proposition 2. □

Proof of Proposition 5. Fix any (i, j) such that $\lambda^i > \lambda^j$. We want to show that $\alpha^i_1(p(t)) \leq \alpha^j_i(p(t))$ and that $\alpha^i_1(p(t)) < \alpha^j_i(p(t))$ whenever $\alpha^i_1(p(t)) < 1$. First suppose $\alpha^i_1(p) = 1$. In this case, $\alpha^i_1(p) \geq \alpha^j_1(p)$ is trivially satisfied.

Next, suppose $\alpha^i_1(p) < 1$. We want to show that $\alpha^i_1(p) > \alpha^j_1(p)$. Suppose by contradiction that $\alpha^i_1(p) \leq \alpha^j_1(p)$. First consider the case where $k_1 < \beta$. Then, let

$$q^* \equiv \inf\{q|\alpha^j_1(p) < 1 \text{ and } \alpha^j_1(p) < \alpha^i_1(p)\}.$$

Because the α^i_1 are continuous, it follows from Proposition 1', and the assumption that $\alpha^i_1(p) \leq \alpha^j_1(p)$, that $q^* < p$ exists. Again, by continuity, $\alpha^j_1(q^*) = \alpha^i_1(q^*)$. It then follows from (ODE') that $\alpha_1^{j\prime}(q^*) < \alpha_1^{i\prime}(q^*)$. But this implies that for some $q > q^*$, $\alpha_1^j(q^*) > \alpha_1^i(q^*)$. Contradiction.

Next, consider the case where $k_1 \geq \beta$. Recall by Proposition 2' that $\lim_{p \to 0+} \alpha_1^i(p) = \lim_{p \to 0+} \alpha_1^j(p)$. Thus, there exists some $q \in (0, p]$ such that $\alpha_1^i(p) \leq \alpha_1^j(p)$ and $\alpha_1^{i\prime}(p) \leq \alpha_1^{j\prime}(p)$. However, it again follows from (ODE') that $\alpha_1^{i\prime}(p) > \alpha_1^{j\prime}(p)$. Contradiction. □

Appendix B: Omitted Proofs from Chapter 2

B.1 Relevant notation and properties

Before proceeding, we will introduce some relevant notation and properties which will prove useful in the analysis which follows.

We begin with the sender's value function. Let $V_t^i(m, p)$ denote type i sender's value from sending message m at time t when she holds belief p. If $m = \emptyset$ this is equal to the sender's continuation value. If $m \in \{0, 1\}$, this value is a linear function of the sender's belief p:

$$V_t^i(m, p) = \tilde{p}R(m, 1) + (1 - \tilde{p})R(m, 0)$$

for $m \in \{0, 1\}$, where $\tilde{p}$ is the sender's belief $s = 1$ will realize, formally $\tilde{p} \equiv p\pi + (1 - p)(1 - \pi)$. We further let $V_t^i(p)$ denote the sender's value at time t under belief p.

$$V_t(p) \equiv \max_m V_t(m, p).$$

For brevity, in the analysis that follows, we will often drop the i superscript when referring to the bad sender.

Next, we define a function of the sender's strategy σ, which we call the *effective arrival rate*. This denotes the probability with which the sender will know the state is either 0 or 1 at time t, respectively. We now formally define this object:

Definition 4. The effective arrival rate of θ at time t is given by

$$\lambda_B^{t,\theta} = \begin{cases} \frac{\lambda_B}{2} & \text{if } t = 1 \\ (1 - \lambda_B^{t-1,\theta}\sigma_{t-1}(\emptyset, \theta))\lambda_B + \lambda_B^{t-1,\theta}\sigma_{t-1}(\emptyset, \theta) & \text{if } t > 1 \end{cases}$$

In the case where $\sigma_s(\emptyset, 0) = \sigma_s(\emptyset, 1)$ for all $s < t$, which holds in much of the analysis below, it follows from the above definition that $\lambda_B^{t,0} = \lambda_B^{t,1}$. Thus to simplify notation, we will in these cases remove the state superscript from the effective arrival rate, and let

$$\lambda_B^t \equiv \lambda_B^{t,0} + \lambda_B^{t,1}$$

Now, we will formally define the sender's interim reputation, which as noted in the main text, denotes the receiver's belief about the sender's type, given that she has not reported on or before t. While we provide a formula in the main text, that formula is only relevant under the equilibrium strategy which we derive. The general formula follows recursively from Bayes Rule and is given by the following for all $t \in \{1, ..., T\}$:

$$R_t = \frac{1}{1 + (\frac{1-R_{t-1}}{R_{t-1}} \frac{\lambda_B^{t,0}\sigma(\emptyset,0)+\lambda_B^{t,1}\sigma(\emptyset,1)+(1-\lambda_B^{t,0}-\lambda_B^{t,1})\sigma(\emptyset,\frac{1}{2})}{1-\lambda_G})}$$

Meanwhile, the reputation function $R_t(m, s)$ denotes the receiver's belief that the sender is the good type, given that she sends message m and time t, and the receiver observes private signal s. Let E_t denote the event that the sender does not report before time t, i.e., that $\tau \geq t$. Then Bayes Rule yields that for $m \in \{0, 1\}$:

$$R_t(m, s) = \frac{1}{1 + \frac{1-R_{t-1}}{R_{t-1}} L_t(m, s)}$$

where $L_t(m, s) \equiv \frac{Pr(m,s|B,E_t)}{Pr(m,s|G,E_t)}$, and $Pr(m, s|i, E_t)$ denotes the probability that (t, m, s) will realize, given that a sender of type i hadn't reported before time t.

B.2 Proofs for Static Model

As in the main text, we drop all time indices from functions here.

Proof of Lemma 2. Here we prove that the claim holds for $\theta = 1$. I a symmetric argument can be used to prove the claim for $\theta = 0$.

Suppose by contradiction that $\sigma(1,0) = \sigma(1,\frac{1}{2}) = 0$. Let $Pr(m,s|i)$ denote the probability that (m,s) is realized given the sender's type is i. Then by Bayes Rule:

$$Pr(1,1|B) = \lambda_B \sigma(1,1)\pi \text{ and } Pr(1,1|G) = \lambda_G \pi$$

It follows that

$$L(1,1) = L(1,0) = \frac{\lambda_B \sigma(1,1)}{\lambda_G} < 1$$

This in turn implies that $V(p) \geq V(1,p) > R_0$, i.e., that the bad sender's reputation will strictly improve with probability 1. Thus, R must be in violation of Bayes Rule: contradiction. □

Proof of Lemma 3. Here, we prove that $R(1,1) > R(1,0)$. That $R(0,0) > R(0,1)$ follows symmetrically.

By the expression for R, showing $R(1,1) > R(1,0)$ is equivalent to showing that $L(1,1) < L(1,0)$. It follows by Bayes Rule that

$$L(1,1) = \frac{1}{\lambda_G}[\frac{1-\pi}{\pi}P(m=1|\theta=0,B) + P(m=1|\theta=1,B)]$$

$$L(1,0) = \frac{1}{\lambda_G}[\frac{\pi}{1-\pi}P(m=1|\theta=0,B) + P(m=1|\theta=1,B)]$$

where $P(m|\theta,B)$ denotes the probability that message m is sent, given the state is θ and the sender is bad. It is given by:

$$Pr(m|\theta,B) = \sigma(m,\theta)\lambda_B + \sigma(m,\frac{1}{2})(1-\lambda_B)$$

It follows from Lemma 2 that $Pr(m=1|\theta=0,B) > 0$. Since by assumption $\pi > \frac{1}{2}$, it follows that

$L(1,1) < L(1,0)$. □

We will now prove a corollary of Lemma 3, which will be used in proving Proposition 6. This corollary establishes strict monotonicity of the value functions:

Corollary 2. *In any equilibrium, $V(1,p)$ ($V(0,p)$) is strictly increasing (decreasing).*

Proof of Corollary 2. We will prove that $V(1,p)$ is strictly increasing. That $V(0,p)$ is strictly decreasing follows symmetrically. Rerranging the above formula for $V(1,p)$, we have:

$$V(1,p) = p(2\pi - 1)[R(1,1) - R(1,0)] + [R(1,1)(1-\pi) + R(1,0)\pi]$$

By (3), it follows that $R(1,1) - R(1,0) > 0$, and thus $V(1,p)$ is strictly increasing in p. □

Proof of Proposition 6. We will prove that $\sigma(1,1) = 1$. That $\sigma(0,0) = 1$ follows symmetrically.

We will now show $\sigma(0,1) = \sigma(\emptyset,1) = 0$. First, suppose by contradiction that $\sigma(0,1) > 0$. It follows that message 0 is optimal under $p = 1$, and thus $V(0,1) \geq V(1,1)$. It follows from Corrolary 2 that

$$V(0,p) > V(0,p) \text{ for all } p < 1.$$

Thus, $\sigma(1,\frac{1}{2}) = \sigma(1,0) = 0$. This is a contradiction of Lemma 2.

Now, suppose by contradiction that $\sigma(\emptyset,1) > 0$. To this end we first show that $R(\emptyset,1) > R(\emptyset,0)$. By the above formula for V, this is equivalent to showing that $V(\emptyset,1) > V(\emptyset,\frac{1}{2})$, which follows from the following chain of inequalities:

$$V(\emptyset,1) \geq V(1,1) > V(1,\frac{1}{2}) \geq V(\emptyset,\frac{1}{2})$$

The first inequality in this chain follows from the assumption that message $\emptyset$ is optimal under $p = 1$. The second inequality follows from Corollary 2. The third inequality holds because, by the above reasoning $\sigma(1,0) = 0$, which implies by Lemma 2 that $\sigma(1,\frac{1}{2}) > 0$.

Next, we show $\sigma(\emptyset, 0) = 0$. It follows from Lemma 2 that $V(\emptyset, \frac{1}{2}) \leq V(0, \frac{1}{2})$. Because, $V(\emptyset, \cdot)$ is strictly increasing while $V(0, \cdot)$ is strictly decreasing, this implies $V(\emptyset, 0) < V(0, 0)$, thus showing that $\sigma(\emptyset, 0) = 0$. This implies:

$$P(m = \emptyset | \theta = 0, B) = (1 - \lambda_B)\sigma(\emptyset, \frac{1}{2}) < \lambda_B \sigma_B(\emptyset, 1) + (1 - \lambda_B)\sigma(\emptyset, \frac{1}{2}) = P(m = \emptyset | \theta = 1, B)$$

which implies:

$$L(\emptyset, 0) = \frac{(1-\pi)P(\emptyset|1, B) + \pi P(\emptyset|0, B)}{1 - \lambda_G} < \frac{\pi P(\emptyset|1, B) + (1-\pi)P(\emptyset|0, B)}{1 - \lambda_G} = L(\emptyset, 1).$$

This inequality on L is equivalent to $R(\emptyset, 0) > R(\emptyset, 1)$, which is a contradiction of the above. □

Proof of Proposition 7. We begin by proving the first part of the proposition, namely that $\sigma(m, \frac{1}{2}) > 0$ for all m. First, note that because by Proposition 6 $\sigma(1, 0) = \sigma(0, 1) = 0$, in order to satisfy Lemma 2, it must be that $\sigma(m, \frac{1}{2}) > 0$ for $m \in \{0, 1\}$. Next, suppose by contradiction that $\sigma(\emptyset, \frac{1}{2}) = 0$. It then follows by Lemma 2 that the bad sender never sends message $\emptyset$. Because the good type does so with strictly positive probability, it follows that a 1 message will reveal that the sender is good, i.e., $R(\emptyset, 1) = R(\emptyset, 0) = 1$. Thus a report of $\emptyset$ will guarantee a reputation of 1 for the bad sender, and thus must serve as a profitable deviation.

Next, we prove the second part of the proposition, namely, that $\sigma(1, \frac{1}{2}) = \sigma(0, \frac{1}{2})$. Assume by contadiction that this does not hold. Without loss of generality, let us assume that $\sigma(1, \frac{1}{2}) > \sigma(0, \frac{1}{2})$. By definition,

$$\begin{aligned} L(0,0) &= \frac{(1-\pi)(1-\lambda)\sigma(0, \frac{1}{2}) + \pi(\lambda_B + (1-\lambda_B)\sigma(0, \frac{1}{2})}{\pi \lambda_G} \\ &< \frac{(1-\pi)(1-\lambda)\sigma(1, \frac{1}{2}) + \pi(\lambda_B + (1-\lambda_B)\sigma(1, \frac{1}{2})}{\pi \lambda_G} = L(1,1) \end{aligned}$$

Furthermore,

$$\begin{aligned} L(0,1) &= \frac{\pi(1-\lambda_B)\sigma(0,\frac{1}{2}) + (1-\pi)(\lambda_B + (1-\lambda_B)\sigma(0,\frac{1}{2}))}{(1-\pi)\lambda_G} \\ &< \frac{\pi(1-\lambda_B)\sigma(1,\frac{1}{2}) + (1-\pi)(\lambda_B + (1-\lambda_B)\sigma(1,\frac{1}{2}))}{(1-\pi)\lambda_G} = L(1,0) \end{aligned}$$

It follows from these two inequalities that $R(0,0) > R(1,1)$ and $R(0,1) < R(1,0)$. Thus,

$$V(0,\frac{1}{2}) = \frac{1}{2}R(0,0) + \frac{1}{2}R(0,1) > \frac{1}{2}R(1,1) + \frac{1}{2}R(1,0) = V(1,\frac{1}{2})$$

However, this implies that $\sigma(1,\frac{1}{2}) = 0$, which is a contradiction of the first part of the claim. □

Before proceeding with the proof for Proposition 8, we introduce some relevant notation. Recalling the definition of σ^b above, let R^b, V^b, and L^b denote the reputation, value and likelihood functions, respectively that are consistent with σ^b. Furthermore, let

$$X(b) \equiv V^b(1,\frac{1}{2}) - V^b(\emptyset,\frac{1}{2})$$

Proof of Proposition 8. We wish to show that there exists a unique b such that σ^b constitutes an equilibrium. By Proposition 7, if σ^b constitutes an equilibrium, then $X(b) = 0$. So, we begin by showing that there exists a unique $b^* \in (0,1)$ such that $X(b^*) = 0$. To this end, we make two observations about $X(b)$:

1. $X(b)$ is continuous and strictly decreasing in b. It suffices to show that $V^b(1,\frac{1}{2})$ is continuous and strictly decreasing in b and $V^b(\emptyset,\frac{1}{2})$ is continuous and strictly increasing in b. First note that for all $s \in \{0,1\}$,

$$L^b(1,s) = \frac{(1-\lambda_B)b/2 + Pr(s|\theta=1)\lambda_B}{Pr(s|\theta=1)\lambda_G}$$

which implies that $R^b(1,s)$ is continuous and strictly decreasing in b for $s \in \{0,1\}$, and thus that $V^b(1,\frac{1}{2})$ is continuous and strictly decreasing in b.

Next, note that

$$L^b(\emptyset, s) = \frac{(1-\lambda_B)(1-b)}{(1-\lambda_G)} \text{ for } s \in \{0,1\}$$

which implies that $R^b(\emptyset, s)$, and consequently $V^b(\emptyset, \frac{1}{2})$ is continuous and strictly increasing in b for $s \in \{0,1\}$.

2. $X(0) > 0$ and $X(1) < 0$. To show $X(0) > 0$, note that by the above formulai, for $s \in \{0,1\}$,

$$L^0(1,s) = \frac{\lambda_B}{\lambda_G} < \frac{1-\lambda_B}{1-\lambda_G} = L^0(\emptyset, s)$$

Thus, $R^0(1,s) > R^0(\emptyset, s)$ for $s \in \{0,1\}$. Therefore, $X(0) > 0$. To show $X(1) < 0$, note that for $s \in \{0,1\}$, $L^1(1,s) > 0 = L^1(\emptyset, s)$. Thus, $R^1(1,s) < R^1(\emptyset, s)$ for all $s \in \{0,1\}$. Thus, $X(1) < 0$.

Combining the above two observations, it follows that there exists a unique $b^* \in (0,1)$ such that $X(b^*) = 0$. Thus we have shown that the only candidate equilibrium is $(m_B^{b^*}, R^{b^*})$. It remains to confirm that this is indeed an equilibrium, i.e., that the sender cannot profitably deviate at any possible belief.

Let us begin with the belief $p = \frac{1}{2}$. First, note that under $m_B^{b^*}$, $V^{b^*}(1, \frac{1}{2}) = V^{b^*}(0, \frac{1}{2})$. Thus, because $X(b^*) = 0$

$$V^{b^*}(1/2) = V^{b^*}(m, 1/2) \text{for all } m \tag{B.1}$$

Next, we will show there does not exist a profitable deviation when $p = 1$. That there does not exist a profitable deviation when $p = 0$ follows symmetrically. To show this, first note by definition of L:

- $L^{b^*}(\emptyset, 0) = L^{b^*}(\emptyset, 1)$
- $L^{b^*}(1, 0) > L^{b^*}(1, 1)$
- $L^{b^*}(0, 1) > L^{b^*}(0, 0)$

It follows from these three inequalities that

- $V^{b^*}(\emptyset, p)$ is constant in p
- $V^{b^*}(1, p)$ is strictly increasing in p
- $V^{b^*}(0, p)$ is strictly decreasing in p

Thus, it follows from (5) that

$$V^{b^*}(1,1) > V^{b^*}(\emptyset,1) > V^{b^*}(0,1)$$

Thus, $m = 1$ is the unique best response at $p = 1$, and there is no profitable deviation. □

B.3 Proofs for dynamic model

We begin by proving a Lemma that will be of use in our analysis below:

Lemma 12. *In any equilibrium, $V_t(1,p)$ ($V_t(0,p)$) is weakly increasing (decreasing) in p. Furthermore, $V_t(1,p)$ ($V_t(0,p)$) is strictly increasing (decreasing) in p whenever $\sigma_t(1,0)+\sigma_t(1,\frac{1}{2}) > 0$ ($\sigma_t(0,1)+\sigma_t(0,\frac{1}{2}) > 0$).*

Proof. We begin by proving that $V_t(1,p)$ is weakly increasing in p. That $V_t(0,p)$ is weakly decreasing in p follows symmetrically. First recall that

$$V_t(1,p) = p(2\pi-1)[R(t,1,1) - R_t(1,0)] + [R(t,1,1)(1-\pi) + R(t,1,0)\pi]$$

Thus, to show that $V_t(1,p)$ is weakly increasing in p, it suffices to show that $R(t,1,1) \geq R(t,1,0)$. To show this, by definition of R, it suffices to show that $L_t(1,1) \leq L_t(1,0)$. This holds by definition, since

$$L_t(1,1) = \frac{\lambda_B^{t,1}\sigma_t(1,1)/2 + \lambda_B^{t,0}\frac{1-\pi}{\pi}\sigma_t(1,0)/2 + (1-\lambda_B^{t,1}-\lambda_B^{t,0})\frac{\sigma_t(1,\frac{1}{2})}{2\pi}}{\lambda_G} \leq$$
$$\frac{\lambda_B^{t,1}\sigma_t(1,1)/2 + \lambda_B^{t,0}\frac{\pi}{1-\pi}\sigma_t(1,0)/2 + (1-\lambda_B^{t,1}-\lambda_B^{t,0})\frac{\sigma_t(1,\frac{1}{2})}{2(1-\pi)}}{\lambda_G} = L_t(1,0)$$

Next let's consider the case where $\sigma_t(1,0) + \sigma_t(1, \frac{1}{2}) > 0$. In this circumstance, the above weak inequality on $L_t(1,1)$ and $L_t(1,0)$ will become a strict inequality, thus yielding that $R(t,1,1) > R(t,1,0)$, and thus that $V_t(1,p)$ is strictly increasing in p. □

Proof of Lemma 4. Fix a t. Suppose by induction that all three claims hold for $s < t$ (this holds vacuously when $t = 1$).

We wish to show that $\sigma_t(1,0) = (0,1)$, and that $\sigma_t(m, \frac{1}{2}) > 0$ for $m \in \{0,1\}$.

We first consider the case where $\lambda_B^t < \lambda_G$. Let us begin by showing $\sigma_t(1,0) = \sigma_t(0,1) = 0$. Suppose by contradiction that $\sigma_t(0,1) > 0$. This implies by Lemma 12 that $V_t(0,p)$ is strictly decreasing in p, and that $V_t(1,p)$ is strictly increasing in p. It thus follows that $V_t(0,p) > V_t(1,p)$ for all $p < 1$. This, however, implies that the sender will only ever report 1 when $p = 1$, i.e.,

$$\sigma_t(1,0) = \sigma_t(1, \frac{1}{2}) = 0$$

Thus,

$$L_t(1,1) = L_t(1,0) = \frac{\lambda_B^t \sigma_t(1,1)}{\lambda_G} < 1$$

Since by assumption $\lambda_B^t < \lambda_G$, this implies that $R_t(1,1) = R_t(1,0) > R_{t-1}$, i.e., that the sender can guarantee an improved reputation by reporting 1. This implies that all senders' reputations would strictly improve at time t, implying that R_t must not be consistent with Bayes Rule. Next, we show that $\sigma_t(m, \frac{1}{2}) > 0$ for $m \in \{0,1\}$. Suppose by contradiction that $\sigma_t(1, \frac{1}{2}) = 0$. Since we showed above that $\sigma_t(1,0) = 0$, we once again obtain that $R_t(1,1) = R_t(1,0) > R_{t-1}$, implying that all senders would strictly improve their reputations at time t, implying that R_t must violate Bayes' Rule.

Next, we consider the case where $\lambda_B^t \geq \lambda_G$. Note it follows form the definition of λ_B^t that $\sigma_{t-1}(\emptyset, 1) = \sigma_{t-1}(\emptyset, 0) > 0$. Because the good sender must be acting optimally and good senders truthfully report arrivals, it follows that

$$V_{t-1}(1,1) = V_{t-1}(\emptyset, 1) = V_t(1,1)$$

where the first equality follows from the fact that $\sigma_{t-1}(\emptyset, 1) = \sigma_{t-1}(\emptyset 0) > 0$. Now suppose by contradiction that $\sigma_t(0,1) > 0$. Then by Lemma 12 $V_t(1,p)$ is weakly increasing in p and $V_t(0,p)$ is strictly decreasing in p. Implying that $\sigma_t(1,\frac{1}{2}) = \sigma_t(0,\frac{1}{2}) = 0$. Then, again applying Lemma 12, this implies that $V_t(1,p)$ is constant in p. Meanwhile, however, by the inductive assumption

$$V_{t-1}(1,\frac{1}{2}) \geq V_{t-1}(\emptyset,\frac{1}{2}) \geq \lambda_B^t V(1,\frac{1}{2}) + (1-\lambda_B^t)V_t(1,1) = V_t(1,1)$$

where the final inequality follows from the fact that $V_t(1,p)$ is constant in p. Combining this with the prior established fact that $V_{t-1}(1,1) = V_t(1,1)$, we obtain that $V_{t-1}(1,\frac{1}{2}) \geq V_{t-1}(1,1)$. However, by the inductive assumption, $\sigma_{t-1}(1,\frac{1}{2}) > 0$, which implies by *Lemma* 12 that $V_{t-1}(1,p)$ is strictly increasing in p. Contradiction. Next, we wish to show that $\sigma_t(m,\frac{1}{2}) > 0$ for $m \in \{0,1\}$. Suppose by contradiction that $\sigma_t(1,\frac{1}{2}) = 0$. By the previously established fact that $\sigma_t(1,0) = 0$, this implies by Lemma 12 that $V_t(1,p)$ is constant in p. We would once again obtain that $V_{t-1}(1,\frac{1}{2}) \geq V_t(1,1)$, implying that $V_{t-1}(1,\frac{1}{2}) \geq V_{t-1}(1,1)$, contradicting the fact that $V_{t-1}(1,p)$ must be strictly increasing in p.

Next, we show 1, namely, that $\sigma_t(\emptyset,1) = \sigma_t(\emptyset,0)$.Because we have established that $\sigma_t(0,1) = \sigma_t(1,0) = 0$, it suffices to show that $\sigma_t(1,1) = \sigma_t(0,0)$. Suppose not, i.e., suppose by contradiction that $\sigma_t(1,1) > \sigma_t(0,0)$.

First, I claim that $R_t(0,0) > R_t(1,1)$. Suppose not, i.e., that $R_t(0,0) \leq R_t(1,1)$. By definition of R, this implies that $L_t(1,1) \leq L_t(0,0)$, i.e.,

$$\begin{aligned}
\frac{\lambda_B^t \pi \sigma_t(0,0) + \frac{1}{2}(1-\lambda_B^t)\sigma_t(0,\frac{1}{2})}{\pi\lambda_G} &\geq \frac{\lambda_B^t \pi \sigma_t(1,1) + \frac{1}{2}(1-\lambda_B^t)\sigma_t(1,\frac{1}{2})}{\pi\lambda_G} \\
\Leftrightarrow \lambda_B^t \pi(\sigma_t(1,1) - \sigma_t(0,0)) &\leq \frac{1}{2}(1-\lambda_B^t)(\sigma_t(0,\frac{1}{2}) - \sigma_t(1,\frac{1}{2})) \\
\Leftrightarrow \lambda_B^t (1-\pi)(\sigma_t(1,1) - \sigma_t(0,0)) &< \frac{1}{2}(1-\lambda_B^t)(\sigma_t(0,\frac{1}{2}) - \sigma_t(1,\frac{1}{2})) \\
&\Leftrightarrow L_t(0,1) > L_t(1,0) \\
&\Leftrightarrow R_t(0,1) < R_t(1,0).
\end{aligned}$$

Thus, it follows that

$$V_t(0, \frac{1}{2}) = \frac{R_t(0,0) + R_t(0,1)}{2} < \frac{R_t(1,1) + R_t(1,0)}{2} = V_t(1, \frac{1}{2}).$$

However, this would violate the above established fact that $\sigma_t(0, \frac{1}{2}) > 0$. Contradiction.

Next I claim that $R_t(1,0) > R_t(0,1)$. To show this, assume by contradiction that that $R_t(0,1) \geq R_t(1,0)$. Given our above result that $R_t(1,1) < R_t(0,0)$, this would imply that

$$V_t(0, \frac{1}{2}) = \frac{R_t(0,0) + R_t(0,1)}{2} > \frac{R_t(1,1) + R_t(1,0)}{2} = V_t(1, \frac{1}{2}).$$

Again, contradicting the fact that $\sigma_t(0, \frac{1}{2}) > 0$.

Next, we show that $V_t(1,1) < V_t(0,0)$. To see this note that

$$V_t(1, \frac{1}{2}) = V_t(0, \frac{1}{2}) \Leftrightarrow [R_t(1,1) - R_t(0,0)] + [R_t(1,0) - R_t(0,1)] = 0$$
$$\Leftrightarrow \pi[R_t(1,1) - R_t(0,0)] + (1-\pi)[R_t(1,0) - R_t(0,1)] < 0 \Leftrightarrow V_t(1,1) < V_t(0,0)$$

which follows from our earlier observations that $R_t(1,1) < R_t(0,0)$ and $R_t(1,0) > R_t(0,1)$.

Now, note that because by assumption σ_s is report symmetric for all $s < t$, it follows that the sender's type is uncorrelated the state at time t. Formally, let $R_{t|\theta}$ denote the expected reputation of the sender (from the receiver's perspective), given that at time t she holds prior belief R_{t-1} about the sender's type. Formally,

$$R_{t|\theta} \equiv R_{t-1}[V_t(\theta,\theta)\lambda_G + R_t(1-\lambda_G)] + (1-R_{t-1})[V_t(\theta,\theta)\lambda_B + V_t(\theta, \frac{1}{2})(1-\lambda_B)]$$

If the sender's type is uncorrelated with the state at time t, it follows that $R_{t|\theta=1} = R_{t|\theta=0}$. However, examining the expression above, we see that $R_{t|\theta=0} > R_{t|\theta=1}$. This is a contradiction.

Finally, we show that $\sigma_t(1, \frac{1}{2}) = \sigma_t(0, \frac{1}{2})$. Suppose by contradiction that $\sigma_t(1, \frac{1}{2}) > \sigma_t(0, \frac{1}{2})$. Given that $\sigma_t(1,1) = \sigma_t(0,0)$, this would imply that both $L_t(1,1) > L_t(0,0)$ and $L_t(1,0) >$

$L_t(0,1)$, due to the following inequalities:

$$L_t(1,1) = \frac{\sigma_t(1,1)\pi\lambda_B + \sigma_t(1,\frac{1}{2})\frac{1}{2}(1-\lambda_B)}{\pi\lambda_G} > \frac{\sigma_t(0,0)\pi\lambda_B + \sigma_t(0,\frac{1}{2})\frac{1}{2}(1-\lambda_B)}{\pi\lambda_G} = L_t(0,0)$$

$$L_t(1,1) = \frac{\sigma_t(1,1)(1-\pi)\lambda_B + \sigma_t(1,\frac{1}{2})\frac{1}{2}(1-\lambda_B)}{(1-\pi)\lambda_G} > \frac{\sigma_t(0,0)(1-\pi)\lambda_B + \sigma_t(0,\frac{1}{2})\frac{1}{2}(1-\lambda_B)}{(1-\pi)\lambda_G} = L_t(0,0)$$

This in turn implies that, $R_t(1,1) < R_t(0,0)$ and $R_t(1,0) < R_t(0,1)$, thus

$$V_t(1,\frac{1}{2}) = \frac{R_t(1,1) + R_t(1,0)}{2} < \frac{R_t(0,0) + R_t(0,1)}{2} = V_t(0,\frac{1}{2})$$

Thus, $\sigma_t(1,\frac{1}{2}) = 0$, a violation of the above lemma. □

Lemma 13 (Boundedness of b_t)**.** *If* $\sigma_t(\emptyset, 1) = 0$ *for all* $s \geq t$*, then*

$$b_t < \frac{\lambda_G - \lambda_B^t}{1 - \lambda_B^t}$$

Proof. Suppose by contradiction that $b_t \geq \frac{\lambda_G - \lambda_B^t}{1-\lambda_B^t}$.

First consider the case where $t = T$. Then,

$$L_T(\emptyset,1) = L_T(\emptyset,1) = \frac{(1-\lambda_B^T)(1-b_T)}{1-\lambda_G} \leq 1$$

where the final inequality follows from our assumption above on b_T. It follows that

$$R_T(\emptyset,1) = R_T(\emptyset,0) \geq R_{T-1}.$$

Thus, $V_T(\frac{1}{2}) \geq V_T(\emptyset,\frac{1}{2}) \geq R_{T-1}$. However, we also now that

$$V_T(\emptyset,\frac{1}{2}) < V_T(\emptyset,1) \leq V_T(1)$$

Where the first inequality follows from part 3 of Lemma 4 combined with Lemma 12. But then it follows from the above inequalities that the bad sender's reputation at the end of period T must on average strictly exceed her reputation at the the beginning of T:

$$\lambda_B^T V_T(1) + (1 - \lambda_B^T) V_T(\frac{1}{2}) > R_{T-1}$$

meaning that the bad sender's reputation on average strictly increases in period T, violating Bayes Rule.

Next, consider the case where $t < T$. Assume by induction that $b_{t+1} < \frac{\lambda_G - \lambda_B^{t+1}}{1 - \lambda_B^{t+1}}$. Because by assumption $\sigma_t(\emptyset, 1) = 0$, $\lambda_B^{t+1} = \lambda_B$, and thus the above inequality becomes

$$b_{t+1} < \frac{\lambda_G - \lambda_B}{1 - \lambda_B}$$

Next, assume by contradiction that $b_t \geq \frac{\lambda_G - \lambda_B^t}{1 - \lambda_B^t}$. This immediately implies two things: $R_{t-1} < R_t$ and $b_t > b_{t+1}$. I claim this then implies that

$$R_t(1, s) < R_{t+1}(1, s) \text{ for } s \in \{0, 1\} \tag{B.2}$$

To see why this must hold, first note that

$$R_t(1, 1) = \frac{1}{1 + \frac{1 - R_{t-1}}{R_{t-1}} \frac{\lambda_B^t \pi + (1 - \lambda_B^t) b_t / 2}{\lambda_G \pi}}$$

$$R_{t+1}(1, 1) = \frac{1}{1 + \frac{1 - R_t}{R_t} \frac{\lambda_B^t \pi + (1 - \lambda_B^t) b_t / 2}{\lambda_G \pi}}$$

where because $R_t \geq R_{t-1}$, $b_t \geq \frac{\lambda_G - \lambda_B}{1 - \lambda_B} > b_{t+1}$, and the fact that $\lambda_B^t \geq \lambda_B$, it follows that $R_t(1, 1) < R_{t+1}(1, 1)$. One can analogously show that $R_t(1, 0) < R_{t+1}(1, 1)$.

Next, note that

$$V_t(\emptyset, \frac{1}{2}) = q R_{t+1}(1, 1) + (1 - q) R_{t+1}(1, 0)$$

where $q \equiv \pi\lambda_B + \frac{1}{2}(1-\lambda_B) > \frac{1}{2}$. Meanwhile,

$$V_t(1, \frac{1}{2}) = \frac{1}{2}R_{t+1}(1,1) + \frac{1}{2}R_{t+1}(1,0)$$

Then, applying Equation B.2 these expressions, we obtain that $V_t(\emptyset, \frac{1}{2}) > V_t(1, \frac{1}{2})$. However, this is a violation of part 2 of Lemma 4. Contradiction. □

Lemma 14 (Truthfully report arrivals)**.** *In any equilibrium, for every t*

$$\sigma_t(1,1) = \sigma_t(0,0) = 1$$

Proof. Recall that in Lemma 4, we showed $\sigma_t(1,0) = \sigma_t(0,1)$. Thus to prove the proposition, it suffices to show that $\sigma_t(\emptyset,1) = \sigma_t(\emptyset,0) = 0$ for all t. By part 1. of Lemma 4, it suffices to show $\sigma_t(\emptyset,1) = 0$. Suppose by contradiction this is not the case. Let t^* denote the last period such that this condition is not satisfied. Formally,

$$t^* = \max\{t \in \{1,...,T\} | \sigma_t(\emptyset,1) = 0\}.$$

First consider the case where $t^* = T$. In this case, because both messages $\emptyset$ and 1 must be optimal under belief p in period T,

$$V_T(\emptyset,1) = V_T(1,1)$$

Secondly, by part 3 of Lemma 4:

$$V_T(\emptyset, \frac{1}{2}) \leq V_T(1, \frac{1}{2}).$$

Note further that by Lemma 12 combined with Lemma 4, it follows that $V_T(1,p)$ is strictly increasing in p. Combining this with the two inequalities above, we obtain

$$V_T(\emptyset,1) > V_T(\emptyset, \frac{1}{2})$$

Since by Lemma 4 it follows that $\sigma_T(\emptyset, 0) > 0$ as well, by analogous reasoning as above, we also obtain that

$$V_T(\emptyset, 0) > V_T(\emptyset, \frac{1}{2}).$$

The above two inequalities imply that $V_T(\emptyset, p)$ is not monotonic in p. However, recall that

$$V_T(\emptyset, p) = p(2\pi - 1)[R_T(\emptyset, 1) - R_T(\emptyset, 0)] + [R_T(\emptyset, 1)(1 - \pi) + R_T(\emptyset, 0)\pi]$$

which is monotonic in p. Contradiction.

Next, consider the case in which $t^* < T$. Because it is optimal for time $t^* + 1$ senders who know the state to report it truthfully,

$$V_{t^*}(1, 1) = V_{t^*}(\emptyset, 1) = V_{t^*+1}(1, 1)$$

where the first equality follows from the fact that $\sigma_{t^*}(\emptyset, 1) > 0$. Furthermore, in order to ensure that faking at time t^* is optimal, which must hold by Lemma 4,

$$\frac{V_{t^*}(1, 1) + V_{t^*}(1, 0)}{2} \geq (\lambda_B + \frac{1 - \lambda_B}{2})V_{t^*+1}(1, 1) + \frac{1 - \lambda_B}{2}V_{t^*+1}(1, 0)$$

This it follows that $V_{t^*}(1, 0) > V_{t^*}(1, 0)$. This combined with our earlier equality $V_{t^*}(1, 1) = V_{t^*+1}(1, 1)$ implies that

$$R_{t^*}(1, 1) < R_{t^*+1}(1, 1) \text{ and } R_{t^*}(1, 0) > R_{t^*+1}(1, 0)$$

Recalling the above expression for the reputation function above, $R_{t^*}(1, 0) > R_{t^*+1}(1, 0)$ will hold only if

$$\frac{1 - R_{t^*-1}}{R_{t^*-1}}(\sigma_{t^*}(1, 1)\lambda_B^{t^*}(1-\pi)+\frac{1}{2}\sigma_{t^*}(1, \frac{1}{2})(1-\lambda_B^{t^*})) < \frac{1 - R_{t^*}}{R_{t^*}}(\sigma_{t^*+1}(1, 1)\lambda_B^{t^*+1}(1-\pi)+\frac{1}{2}\sigma_{t^*+1}(1, \frac{1}{2})(1-\lambda_B^{t^*+1}))$$

Rearranging, this is equivalent to:

$$[\frac{1-R_{t^*-1}}{R_{t^*-1}}\sigma_{t^*}(1,1)\lambda_B^{t^*} - \frac{1-R_{t^*}}{R_{t^*}}\sigma_{t^*+1}(1,1)\lambda_B^{t^*+1}](1-\pi) <$$
$$\frac{1}{2}[\frac{1-R_{t^*}}{R_{t^*}}\sigma_{t^*}(1,\frac{1}{2})(1-\lambda_B^{t^*+1}) - \frac{1-R_{t^*-1}}{R_{t^*-1}}\sigma_{t^*+1}(1,\frac{1}{2})(1-\lambda_B^{t^*})]$$

Next, I claim

$$\frac{1-R_{t^*-1}}{R_{t^*-1}}\sigma_{t^*}(1,1)\lambda_B^{t^*} - \frac{1-R_{t^*}}{R_{t^*}}\sigma_{t^*+1}(1,1)\lambda_B^{t^*+1} > 0. \tag{B.3}$$

Suppose by contradiction that the left-hand side is less than or equal to zero. By the above inequality, it would then follow that

$$[\frac{1-R_{t^*-1}}{R_{t^*-1}}(\sigma_{t^*}(1,1)\lambda_B^{t^*} - \frac{1-R_{t^*}}{R_{t^*}}(\sigma_{t^*+1}(1,1)\lambda_B^{t^*+1}]\pi < \frac{1}{2}\sigma_{t^*+1}(1,\frac{1}{2})(1-\lambda_B^{t^*+1}) - \frac{1}{2}\sigma_{t^*}(1,\frac{1}{2})(1-\lambda_B^{t^*})$$

However, this holds if and only if $R_{t^*}(1,1) > R_{t^*+1}(1,1)$, a contradiction of the above.

Now let us examine (B.3). First, note that since $\sigma_{t^*+1}(1,1) = 1$, and by definition of λ_B^t, it must be that $\lambda_B^{t^*}\sigma_{t^*} < \lambda_B^{t^*+1}$, it must be that $R_{t^*-1} < R_{t^*}$. From our earlier expression for R_t, this implies an upper bound on $\sigma_{t^*}(1,\frac{1}{2})$

$$1-\lambda_G > (1-\lambda_B^{t^*})(1-\sigma_{t^*}(1,\frac{1}{2})) + \lambda_B^{t^*}\sigma_{t^*}(\emptyset,1) \Leftrightarrow \sigma_{t^*}(1,\frac{1}{2}) > \frac{1-\lambda_B^{t^*+1}}{1-\lambda_B^{t^*}}\sigma_{t^*+1}(1,\frac{1}{2}) \tag{B.4}$$

Separately, recalling (B.3), the only way to ensure that both $R_{t^*}(1,1) < R_{t^*+1}(1,1)$ and $R_{t^*}(1,0) > R_{t^*+1}(0,1)$ is if

$$\sigma_{t^*}(1,\frac{1}{2}) < \frac{1-\lambda_B^{t^*+1}\sigma_{t^*+1}(1,\frac{1}{2})}{1-\lambda_B^{t^*}} \tag{B.5}$$

next, recall by Lemma 13 that $b_{t^*+1} < \frac{\lambda_G - \lambda_B^{t^*+1}}{1-\lambda_B^{t^*+1}}$ substituting this into (B.5), we have

$$\sigma_{t^*}(1,\frac{1}{2}) < \frac{\lambda_G - \lambda_B^t\sigma_t(\emptyset,1) - (1-\lambda_B^t\sigma_t(\emptyset,1))\lambda_B}{1-\lambda_B^t}$$

which contradicts (B.4). □

Proof of Proposition 9. First, we show that in equilibrium $\sigma_t(0,0) = \sigma_t(1,1) = 1$. This follows directly from part 2 of Lemma 4 and Lemma 14. Next, we must show that $b_t \equiv \sigma_t(1, \frac{1}{2}) = \sigma_t(0, \frac{1}{2}) > 0$ at all t, which is given directly by Lemma 4. It remains to show that $b_t < 1$. Suppose not by contradiction. Then it follows that the sender can guarantee a reputation of 1 by abstaining at t, regardless of her information. This is due to the fact that $b_t = 1$ implies that $R_t = 1$, and thus it follows from Bayes' Rule that $R_\tau(m, s) = 1$ for all m, s, and $\tau \geq t$. Thus, abstaining at t serves as a profitable deviation. Contradiction. □

Proof of Proposition 11. Combining Lemma 14 and Lemma 13, it follows that for all t,

$$b_t < \frac{\lambda_G - \lambda_B}{1 - \lambda_B}$$

Recall further that it follows from Lemma 4 that

$$R_t = \frac{1}{1 + \frac{1-R_{t-1}}{R_{t-1}} \frac{(1-\lambda_B)(1-b_t)}{1-\lambda_G}}$$

Combining this equality with the previous inequality on b_t implies the statement. □

Proof of Proposition 12. Suppose by contradiction that $b_t \leq b_{t+1}$. Under the equilibrium strategy, for any $\tau \in 1, ..., T$ and $s \in \{0, 1\}$,

$$L_\tau(1,1) = \frac{Pr(s|\theta = 1)\lambda_B + (1 - \lambda_B)b_\tau/2}{\lambda_G Pr(s|\theta = 1)}$$

It follows from the assumption that $L_t(1, s) \leq L_{t+1}(1, s)$ for all s. Recalling that

$$R_t(1,s) = \frac{1}{1 + \frac{1-R_{t-1}}{R_{t-1}} L_t(1,s)}$$

Proposition 11 above then implies that $R_t(1, s) < R_{t+1}(1, s)$ for $s \in \{0, 1\}$.

Next, recall the value of the uninformed sender at time t from reporting 1 is given by:

$$V_t(1, \frac{1}{2}) = \frac{R_t(1,1) + R_t(1,0)}{2}$$

Meanwhile, her value from abstaining at t is given by:

$$V_t(\emptyset, \frac{1}{2}) = (\lambda_B \pi + (1 - \frac{\lambda_B}{2})R_{t+1}(1,1) + \frac{1-\lambda_B}{2}R_{t+1}(1,0)$$

It follows from the above inequalities that there exists a $\bar{\lambda_B}$ such that if $\lambda_B < \bar{\lambda_B}$, $V_t(\emptyset, \frac{1}{2}) < V_t(1, \frac{1}{2})$. However, this is a contradiction of the equilibrium characterization above, which requires indifference between messages $\emptyset$ and 1 when $p = \frac{1}{2}$.

□

Proof of Proposition 10. First, we establish that for all t, $\alpha_t > 0$. Suppose not, by contradiction. Let us begin by establishing a set of inequalities. First, note that

$$R_t(1,1) = R_{t-1} + \alpha_t \leq R_{t-1}.$$

Separately, $R_t(1,1) > R_t(1,0)$. This holds by identical reasoning to what is presented in Lemma 3 (i.e., the static case). This implies that $R_t(1,0) < R_{t-1}$. Finally, by Proposition 11, $R_t < R_{t-1}$.

Now, assume that at time $t-1$, the sender has not yet reported. Then, the receiver's expected time-t belief about the sender's type at time $t-1$ is given by

$$\begin{aligned} E_{t-1}[Pr_t[\theta = G]] = 2Pr_{t-1}(m_t = 1 = s)R_t(1,1) + 2Pr_{t-1}(m_t = 1 \neq s)R_t(1,0) \\ +(1 - 2Pr_{t-1}(m_t = 1 = s) - 2Pr(m_t = 1 \neq s))R_t \end{aligned}$$

Since $\lambda_B < 1$ and $b_t < 1$ (by Proposition 9), $Pr_{t-1}(m_t = 1 \neq s) > 0$. Combining this with the above-established inequalities that $R_t(1,1) \leq R_{t-1}$, $R_t(1,0), R_{t-1}$, and $R_t < R_{t-1}$ yields

$$E_{t-1}[Pr_t[\theta = G]] < R_{t-1} = Pr_{t-1}[\theta = G].$$

This is the violation of the martingale property of the receiver's belief about the sender's type. Contradiction.

Now, we claim that for all t, $\alpha_t < -\beta_t$. Note that having established that $\alpha_t > 0$, this also implies that $\beta < 0$. To show this, suppose by contradiction that $\alpha_t \geq -\beta_t$. Then,

$$V_t^i(1) = V_t^i(1,1) = R_{t-1} + \pi\alpha_t - (1-\pi)\beta_t > R_{t-1} \text{ for all } i \in \{B, G\},$$

where the first equality follows from Proposition 9 and the strict inequality follows from the above-established fact that $\alpha_t > 0$. Separately,

$$V_t^B(\frac{1}{2}) = V_t^B(\frac{1}{2}, 1) = R_{t-1} + \frac{1}{2}\alpha_t - \frac{1}{2}\beta_t \geq R_{t-1}$$

Finally, I claim that $V_t^G(\frac{1}{2}) \geq V_t^B(\frac{1}{2})$. We prove this using backwards induction. Note that in the base, case we have

$$V_T^G(\frac{1}{2}) = V_T^B(\frac{1}{2}) = R_T(\emptyset)$$

where the final inequality again follows from Proposition 9. Now, fix any $t < T$, and suppose by induction that $V_{t+1}^G(\frac{1}{2}) \geq V_{t+1}^B(\frac{1}{2})$. We want to show that $V_t^G(\frac{1}{2}) \geq V_t^B(\frac{1}{2})$. To this end, note that

$$V_t^i(\frac{1}{2}) = V_t^i(\frac{1}{2}, \emptyset) = \lambda_i V_{t+1}^i(1,1) + (1-\lambda_i)V_{t+1}^i(\frac{1}{2})$$

Then

$$V_t^G(\frac{1}{2}) - V_t^B(\frac{1}{2}) = (\lambda_G - \lambda_B)[V_{t+1}^G(1,1) - V_{t+1}^B(1,1)] + (1-\lambda_G)[V_{t+1}^G(\frac{1}{2}) - V_{t+1}^B(\frac{1}{2})]$$

Since $V_{t+1}^G(\frac{1}{2}) \geq V_{t+1}^B(\frac{1}{2})$ by the inductive assumption, it suffice to show $V_{t+1}^G(1,1) \geq V_{t+1}^B(\frac{1}{2})$. This indeed holds, since

$$V_{t+1}^B(\frac{1}{2}) = V_{t+1}^B(\frac{1}{2}, 1) = \frac{1}{2}R_{t+1}(1,1) + \frac{1}{2}R_{t+1}(1,0) < \pi R_{t+1}(1,1) + (1-\pi)R_{t+1}(1,0) = V_{t+1}^G(1,1).$$

Now, once again assume that at time t, the sender has not yet reported. Then it follows from Bayes Rule that

$$\begin{aligned} E_{t-1}[Pr_T[\theta = G]] &= Pr_{t-1}[\theta = G] E_{t-1}[Pr_T[\theta = G] | \theta = G] + Pr_{t-1}[\theta = B] E_{t-1}[Pr_T[\theta = G] | \theta = B] \\ &= R_{t-1}(\lambda_G V_t(1,1) + (1-\lambda_G) V_t^G(\frac{1}{2})) + (1-R_{t-1})(\lambda_B V_t(1,1) + (1-\lambda_B) V_t^B(\frac{1}{2})) \end{aligned}$$

Since $\lambda_G > 0$ and $\lambda_B > 0$, it follow sfrom the above inequalities that

$$E_{t-1}[Pr_T[\theta = G]] > R_{t-1} = Pr_{t-1}[\theta = G].$$

This is a violation of the Martingale property of the receiver's belief about the sender's type. Contradiction.

□

Appendix C: Ommitted Proofs from Chapter 3

Proof of Lemma 5. First, we show that it is necessary that in any equilibrium, for all t and $\hat{h}$, $q^G(\hat{h}) = q^B(\hat{h})$. First, suppose by contradiction that there exists some t, $\hat{h}_t = (\hat{s}_1, ..., \hat{s}_{t-1})$ such that $q^G(\hat{h}_t) > q^B(\hat{h}_t)$. Further, let $\hat{h}_{t+1} \equiv (\hat{s}_1, \hat{s}_{t-1}, r)$ and $\hat{h}'_{t+1} \equiv (\hat{s}_1, \hat{s}_{t-1}, l)$. It follows from Bayes Rule that

$$Pr(i = G|\hat{h}_t, \hat{s}_t = s) = \frac{Pr(\hat{s}_t = s|i = G, \hat{h}_t)Pr(i = G|\hat{h}_t)}{Pr(\hat{s}_t = s|i = G, \hat{h}_t)Pr(i = G|\hat{h}_t) + Pr(\hat{s}_t = s|i = B, \hat{h}_t)Pr(i = B|\hat{h}_t)} \tag{C.1}$$

It follows from (C.1) that

$$\lambda_t(\hat{h}_{t+1}) = Pr(i = G|\hat{h}_t, \hat{s}_t = r) = \frac{1}{1 + \frac{q^B(\hat{h}_t)}{q^G(\hat{h}_t)}\frac{1-\lambda_t(\hat{h}_t)}{\lambda_t(\hat{h}_t)}} > \frac{1}{1 + \frac{1-q^B(\hat{h}_t)}{1-q^G(\hat{h}_t)}\frac{1-\lambda_t(\hat{h}_t)}{\lambda_t(\hat{h}_t)}}$$
$$= Pr(i = G|\hat{h}_t, \hat{s}_t = l) = \hat{\lambda}_t(\hat{h}'_{t+1}),$$

where the strict inequality follows from the above assumption that $q^G(\hat{h}_t) > q^B(\hat{h}_t)$. Because the firm is myopic, reporting $\hat{s}_t = \hat{r}$ is the firm's unique best response, and thus $m_t(s_t, \hat{h}_t) = 1$ for $s_t \in \{r, l\}$. But then by definition of q^i, $q^B(\hat{h}_t) = 1$, and thus $q^G(\hat{h}_t) \leq q^B(\hat{h}_t)$. Contradiction.

Next, we show sufficiency, i.e., that any $\{(m_t, \lambda_t)\}_{t=1,...,T}$ such that the λ_t are consistent with m_t and $q^B(\hat{h}) = q^G(\hat{h})$ at all report histories $\hat{h}$ is an equilibrium. Suppose not, by contradiction. Then at some t, $\hat{h}_t = (\hat{s}_1, ..., \hat{s}_{t-1})$ and s_t, $m_t(s_t, \hat{h}_t)$ is not a best response. But again, letting $\hat{h}_{t+1} \equiv (\hat{s}_1, \hat{s}_{t-1}, r)$ and $\hat{h}'_{t+1} \equiv (\hat{s}_1, \hat{s}_{t-1}, l)$, it follows from (C.1) that $\lambda_t(\hat{h}_{t+1}) = \lambda_t(\hat{h}'_{t+1})$. Thus, since the firm is myopic, all $m_t(s_t, \hat{h}_t)$ must yield the same time t payoff, and thus $m_t(s_t, \hat{h}_t)$ is trivially optimal. Contradiction. □

Proof of Proposition 13. It follows from Lemma 5 and equations (3.1) and (3.2) that a funciton m_t

comprises an equilibrium if and only if at every $\hat{h} \in \hat{\mathcal{H}}_t$,

$$\begin{aligned} p^G(\hat{h})\pi_G + (1 - p^G(\hat{h}))(1 - \pi_G) = [p^B(\hat{h})\pi_B + (1 - p^B(\hat{h})(1 - \pi_B)]m_t(r, \hat{h}) \\ +[p^B(\hat{h})(1 - \pi_B) + (1 - p^B(\hat{h}))\pi_B]m_t(l, \hat{h}). \end{aligned} \tag{C.2}$$

First, we see that $(m_t^*, \lambda_t^*)_{t=1,...,T}$ satisfies (C.2) at all t and $\hat{h}$ and is thus an equilibrium. It remains to show that it is the unique most informative equilibrium. To this end, fix a $\hat{h} \in \hat{\mathcal{H}}_t$, and suppose that $p(\hat{h}) \geq \frac{1}{2}$ (note that the proof for the case where $p(\hat{h}) < \frac{1}{2}$) follows analogously, and is this omitted). Now, fix some $m_t(\cdot, \hat{h}) \neq m_t^*(\cdot, \hat{h})$. Note that $m_t(r, \hat{h}) < 1$. If not, then $m_t(r, \hat{h}) = m_t^*(r, \hat{h})$, and then by (C.2), $m_t(l, \hat{h}) = m_t^*(l, \hat{h})$. Thus, we would have $m_t(\cdot, \hat{h}) = m_t^*(\cdot, \hat{h})$, a contradiction. Since $m_t(r, \hat{h}) < 1$, it follows from (C.2) that $m_t(l, \hat{h}) > m_t^*(l, \hat{h})$. Since this holds for all $m_t(\cdot, \hat{h}) \neq m_t^*(\cdot, \hat{h}) \in \mathcal{M}_t(\hat{h})$ and for all $\hat{h}$, it follows by definition (Criterion 2) that m_t^* is the unique most informative equilibrium. □

Proof of Corollary 1. Proof by induction. As the base case, we want to show $\lambda_1(\{r\}) = \lambda_1(\{l\}) = \lambda_0$. By Bayes Rule,

$$\lambda_1(\{h\}) = \frac{q^G(\emptyset)\lambda_0}{q^G(\emptyset)\lambda_0 + q^B(\emptyset)(1 - \lambda_0)}$$

By Lemma 5, $q^G(\emptyset) = q^B(\emptyset)$, and thus $\lambda_1(\{h\}) = \lambda_0$. We can analogously show that $\lambda_1(\{h\} = \lambda_0$.

Next, fix a $t > 1$ and suppose by induction that for all $\hat{h}_{t+1} \in \hat{\mathcal{H}}_t$, $\lambda_{t-1}(\hat{h}_t) = \lambda_0$. We want to show that for all $\hat{h}_{t+1} = (\hat{s}_1, ..., \hat{s}_t)$, $\lambda_t(\hat{h}_{t+1}) = \lambda_0$. To this end, let $\hat{h}_t = (\hat{s}_1, ..., \hat{s}_{t-1})$. First, suppose $\hat{s}_{t-1} = r$ (the case where $\hat{s}_{t-1} = l$ follows analogously). It then follows that

$$\lambda_t(\hat{h}_{t+1}) = \frac{q^G(\hat{h}_t)\lambda_{t-1}(\hat{h}_t)}{q^G(\hat{h}_t)\lambda_{t-1}(\hat{h}_t) + q^B(\hat{h}_t)(1 - \lambda_{t-1}(\hat{h}_t))} = \lambda_{t-1}(\hat{h}_t) = \lambda_0,$$

where the first equality follows from Bayes rule, the second equality follows from Lemma 5's assertion that $q^G(\hat{h}_t) = q^B(\hat{h}_t)$, and the final equality follows from the inductive assumption. □

Proof of Proposition 14. Let us first show that $b(\hat{h})$ is increasing in p_0. Note that it follows from the formula for $p^G(\hat{h})$ and $p^B(\hat{h})$ both are increasing in p_0. Separately, it follows from Proposition 13 that $b(\hat{h})$ is strictly increasing in $p^G(\hat{h})$, and decreasing in $p^G(\hat{h})$. Thus $b(\hat{h})$ must be strictly increasing in p_0. The comparative statics on π_G and π_B follow directly from $b(\hat{h})$. Finally, the comparative static on the r-surplus follows from the fact that both $p^G(\hat{h})$ and $p^B(\hat{h})$ is strictly decreasing in $2n - t + 1$, and that $b(\hat{h})$ is strictly increasing in $p^G(\hat{h})$ and decreasing in $p^B(\hat{h})$. □

Proof of Lemma 6. Fix a t and time-t report history $\hat{h}_t = (\hat{s}_1, ..., \hat{s}_{t-1}) \in \hat{\mathcal{H}}_t$. Now, suppose that $m_t(r, \hat{h}_t) < 1$. We want to show that $m_t(l, \hat{h}_t) = 0$. To this end, let $\hat{h}_{t+1} = (\hat{s}_1, ..., \hat{s}_{t-1}, r)$ and $\hat{h}'_{t+1} = (\hat{s}_1, ..., \hat{s}_{t-1}, l)$. Since $m_t(r, \hat{h}_t) < 1$, it follows that reporting l upon seeing signal r must be optimal at history $\hat{h}_t$, i.e., that

$$F_t(\lambda_t(\hat{h}_{t+1})) \leq F_t(\lambda_t(\hat{h}'_{t+1})) - \tau.$$

But this implies that

$$F_t(\lambda_t(\hat{h}_{t+1})) - \tau < F_t(\lambda_t(\hat{h}'_{t+1})).$$

The left-hand side is precisely the firm's payoff from reporting r upon observing signal l at history $\hat{h}_t$, whereas the right-hand side is the firm's payoff from reporting l upon observing signal l at history $\hat{h}_t$. It this follows that reporting l upon observing signal l is the unique vest response for the firm, i.e., that $m_t(l, \hat{h}_t) = 0$. The proof the claim that $m_t(l, \hat{h}_t) > 0$ implies that $m_t(r, \hat{h}_t) = 1$ follows analogously. □

Proof of Proposition 16. First, let us show that any equilibrium must satisfy the requirements of Proposition 16. To this end, fix a t, and let $(m_t, \lambda_t)_{t=1,...,T}$ denote an equilibrium. Assume inductively that (a) and (b) hold for all $s < t$, and $\hat{h}_s \in \hat{\mathcal{H}}_s$. Now, fix a $\hat{h}_t = (\hat{s}_1, ..., \hat{s}_t) \in \mathcal{H}_t$. First suppose that

$$\tau \geq |F_t(\bar{\lambda}_t(\hat{h}_{t+1})) - F_t(\bar{\lambda}_t(\hat{h}'_{t+1}))|,$$

where $\hat{h}_{t+1} = (\hat{s}_1, ..., \hat{s}_t, r)$ and $\hat{h}'_{t+1} = (\hat{s}_1, ..., \hat{s}_t, l)$. Now, suppose without loss of generality that $F_t(\bar{\lambda}_t(\hat{h}_{t+1})) - F_t(\bar{\lambda}_t(\hat{h}'_{t+1})) > 0$. Suppose by contradiction that $m_t(l, \hat{h}_t) > 0$ or $m_t(r, \hat{h}_t) < 1$.

Let us first consider the case where $m_t(l, \hat{h}_t) > 0$. It follows from Lemma 6 that $m_t(r, \hat{h}_t) = 1$.

Further let $\bar{q}^i(\hat{h}_t)$ denote the consumer's belief that type i will make a r report under $\hat{h}_t$ assuming that the bad firm is truthful. Meanwhile, let $q^i(\hat{h}_t)$ denote the same probability conditional on the bad firm playing the equilibrium strategy in period t. Now recall by Bayes Rule that

$$\lambda_t(\hat{h}_{t+1}) = \frac{q^G(\hat{h}_t)\lambda_{t-1}(\hat{h}_t)}{q^G(\hat{h}_t)\lambda_{t-1}(\hat{h}_t) + q^B(\hat{h}_t)(1-\lambda_{t-1}(\hat{h}_t))}$$

and analogously for $\bar{\lambda}_t(\hat{h}_{t+1})$. Since $q^B(\hat{h}_t) > \bar{q}^B(\hat{h}_t)$, it follows that $\lambda_t(\hat{h}_{t+1}) < \bar{\lambda}_t(\hat{h}_{t+1})$, and $\lambda_t(\hat{h}'_{t+1}) > \bar{\lambda}_t(\hat{h}_{t+1})$. It thus follows that

$$F_t(\lambda_t(\hat{h}_{t+1})) - \tau < F_t(\bar{\lambda}_t(\hat{h}_{t+1})) - \tau \leq F_t(\bar{\lambda}_t(\hat{h}'_{t+1})) < F_t(\lambda_t(\hat{h}'_{t+1})).$$

Note that the left-most expression is the firm's payoff from sending message r under signal l and the right-most term is the firm's payoff from sending message l under signal l. Since this inequality is strict, $m_t(l, \hat{h}_t) = 0$. Contradiction.

Next, suppose that $m_t(r, \hat{h}_t) < 1$. By Lemma 6, $m_t(l, \hat{h}_t) = 0$. It again follows from Bayes Rule that $\lambda_t(\hat{h}_{t+1}) > \bar{\lambda}_t(\hat{h}_{t+1})$ and $\lambda_t(\hat{h}'_{t+1}) < \bar{\lambda}_t(\hat{h}'_{t+1})$. Then:

$$F_t(\lambda_t(\hat{h}'_{t+1})) - \tau < F_t(\bar{\lambda}_t(\hat{h}'_{t+1}) - \tau < F_t(\bar{\lambda}_t(\hat{h}_{t+1})) < F_t(\lambda_t(\hat{h}_{t+1})),$$

Note that the left-most expression is the firm's payoff from sending message l under signal r and the right-most term is the firm's payoff from sending message r under signal r. Since the inequality is strict $m_t(r, \hat{h}_t) = 1$. Contradiction.

Next, suppose that $\tau < F_t(\bar{\lambda}_t(\hat{h}_{t+1})) - F_t(\bar{\lambda}_t(\hat{h}'_{t+1}))$. Suppose by contradiction that $\tau \neq F_t(\lambda_t(\hat{h}_{t+1})) - F_t(\lambda_t(\hat{h}'_{t+1}))$. First, suppose that $\tau > F_t(\lambda_t(\hat{h}_{t+1})) - F_t(\lambda_t(\hat{h}'_{t+1}))$. In this case, the firm's payoff from reporting r when receiving signal l is given by

$$F_t(\lambda_t(\hat{h}_{t+1})) - \tau < F_t(\lambda'_t(\hat{h}_{t+1}))$$

Thus, it must be that $m_t(l, \hat{h}_t) = 0$. Furthermore, since $F_t(\bar{\lambda}_t(\hat{h}_{t+1})) - F_t(\bar{\lambda}_t(\hat{h}'_{t+1})) > 0$, it must be that $m_t(r, \hat{h}_t) = 1$. Thus, the firm is truth telling, and thus $\lambda_t(\hat{h}_{t+1}) = \bar{\lambda}_t)(\hat{h}_{t+1})$, $\lambda_t(\hat{h}'_{t+1}) = \bar{\lambda}_t)(\hat{h}'_{t+1})$. Thus, $\tau > F_t(\bar{\lambda}_t(\hat{h}_{t+1})) - F_t(\bar{\lambda}_t(\hat{h}'_{t+1}))$. Contradiction.

Next, suppose that $\tau < F_t(\lambda_t(\hat{h}_{t+1})) - F_t(\lambda_t(\hat{h}'_{t+1}))$. By analogous reasoning, it follows that $m_t(l, \hat{h}_t) = 1$, and $m_t(r, \hat{h}_t) = 1$. But this implies that the bad firm never reports l, while the good firm does so with positive probability. Thus, it follows from Bayes Rule that

$$\lambda_t(\hat{h}'_{t+1}) = 1 < \lambda_t(\hat{h}_{t+1})$$

But this contradicts our earlier assumption that $\lambda_t(\hat{h}_{t+1}) - \lambda'_t(\hat{h}_{t+1}) > 0$.

Next, we want to show that a $(m^*_t, \lambda^*_t)_{t=1,...,T}$ satisfying (a) and (b) exists, is unique, and is indeed an equilibrium.

First, I claim that there exists a unique pair that both satisfies $(m^*_t, \lambda^*_t)_{t=1,...,T}$ and consistent with Bayes Rule. Fix a t and assume by induction that for all $s < t$, there a unique (m_s, λ_s) satisfying (1) and (2) and being consistent with Bayes Rule. Now, fix a $\hat{h}_t \in \hat{\mathcal{H}}_t$. First consider the case where $\tau \geq |F_t(\bar{\lambda}_t(\hat{h}_{t+1})) - F_t(\bar{\lambda}_t(\hat{h}'_{t+1}))|$. Existence and uniqueness in this case is trivial. Next suppose that $\tau < |F_t(\bar{\lambda}_t(\hat{h}_{t+1})) - F_t(\bar{\lambda}_t(\hat{h}'_{t+1}))|$. Next assume without loss that $F_t(\bar{\lambda}_t(\hat{h}_{t+1})) - F_t(\bar{\lambda}_t(\hat{h}'_{t+1})) > 0$. Now, note that we cannot have $m_t(r, h_t) < 1$. This would imply that $m_t(l, h) = 0$ (by Lemma 6), and it would follow from Bayes Rule that $\lambda_t(\hat{h}_{t+1}) - \lambda_t(\hat{h}'_{t+1}) > \tau$, and thus (b) would fail. Thus we can restrict attention to strategies where $m_t(l, \hat{h}_t) > 0$ and $m_t(r, \hat{h}_t) = 1$. Note further that $F_t(\lambda_t(\hat{h}_{t+1}))$ is continuous and strictly increasing in $m_t(l, \hat{h}_t)$, whereas $F_t(\lambda_t(\hat{h}'_{t+1}))$ is continuous and strictly decreasing in $m_t(l, \hat{h}_t)$. Furthermore, when $m_t(l, \hat{h}_t) = 0$ and $m_t(r, \hat{h}_t) = 1$, the firm is truthful, and thus $\lambda_t(\hat{h}_{t+1}) - \lambda_t(\hat{h}'_{t+1}) > \tau$. Meanwhile, when $m_t(l, \hat{h}_t) = 1$, $\lambda_t(\hat{h}_{t+1}) = 1$, and thus $\lambda_t(\hat{h}_{t+1}) - \lambda_t(\hat{h}'_{t+1}) \leq 0$. Thus by the intermediate value theorem, there exists a unique $m_t(l, \hat{h}_t) \in (0, 1)$ such that $\lambda_t(\hat{h}_{t+1}) - \lambda_t(\hat{h}'_{t+1}) = \tau$. Thus, there exists a unique (m_t, λ_t) satisfying part (b) as well.

Finally, it remains to show that the firm cannot profitably deviate from m^*_t under λ^*_t. Fix a t and $\hat{h}_t$, and first suppose that $\tau \geq |F_t(\bar{\lambda}_t(\hat{h}_{t+1})) - F_t(\bar{\lambda}_t(\hat{h}'_{t+1}))|$. The firm does not have an incentive to deviate from truth-telling under both signal l and r, as the potential reputational benefit $(|F_t(\bar{\lambda}_t(\hat{h}_{t+1})) - F_t(\bar{\lambda}_t(\hat{h}'_{t+1}))|)$ is outweighed by the cost of lying τ. Next, suppose that $\tau < |F_t(\bar{\lambda}_t(\hat{h}_{t+1})) - F_t(\bar{\lambda}_t(\hat{h}'_{t+1}))|$. Assume without loss that $F_t(\bar{\lambda}_t(\hat{h}_{t+1})) - F_t(\bar{\lambda}_t(\hat{h}'_{t+1})) > 0$. Note then that the firm is indifferent between sending message l and r under signal l, and thus cannot profitably deviate. Meanwhile, under signal r, sending message r is strictly optimal, and indeed we have $m^*_t(r, \hat{h}_t) = 1$. □

Proof of Proposition 17. Proof by counterexample. Suppose that $T = 2$. Let the firm's payoff in the first period be given by the weighted sum of the first-and second- period myopic payoffs:

$$u_1 = E[\lambda_1 + \tau\mathbb{I}(\hat{s}_1 = s_1) + \beta(\lambda_2 + \tau\mathbb{I}(\hat{s}_2 = s_2))]$$

where $\beta \in (0, 1)$. We know that

$$E[\lambda_2] = Pr^F(r)\lambda_2^{\hat{r}} + Pr^F(l)\lambda_2^l$$

where $Pr^F(s)$ denotes the firm's belief at time $t = 1$ that it will report s in the second period. Since the equilibrium in the second period is trivially myopic, assuming that τ is sufficiently small, $|\lambda_2(\{r, r\}) - \lambda_2(\{r, l\})| = \tau$. Let us assume without loss that $\lambda_2(\{r, r\}) - \lambda_2(\{r, l\}) = \tau$. Then,

$$E[\lambda_2] = \lambda_2(\{r, r\}) - Pr^F(l)\tau$$

Next, by Bayes consistency, we know

$$\lambda_1 = Pr(\hat{r})\lambda_2^{\hat{r}} + Pr(\hat{l})\lambda_2^l$$

where $Pr(\hat{r})$ and $Pr(\hat{l})$ are the probabilities from the *consumer's* perspective that it will report $\hat{r}$

and $\hat{l}$ in the second period, respectively. Combining these, we have

$$E[\lambda_2] = \lambda_1 + [Pr^F(\hat{r}) - Pr(\hat{r})]\tau.$$

Furthermore, since $E[\mathbb{I}(\hat{s}_2 = s_2)]$ is the probability the firm lies in the second period, we have

$$E[\mathbb{I}(\hat{s}_2 = s_2)] = Pr^F(l)\lambda_2^l$$

where λ_2^l is the firm's equilibrium level of bias in the second period. Thus,

$$E[\lambda_2 + \tau\mathbb{I}(\hat{s}_2 = s_2)] = \lambda_1 + \tau(Pr^F(r) - Pr(\hat{r})).$$

So we have simplified the firm's payoff in the first period to the following:

$$\lambda_1(1+\beta) + \tau\mathbb{I}(\hat{s}_1 = s_1) + \beta\tau(Pr^F(r) - Pr(\hat{r}))$$

Now we will show that the myopic equilibrium in the first period is not necessarily also a forward-looking equilibrium. Recall that for small enough τ, the myopic equilibrium is such that

$$\lambda_1^{\hat{r}} = \lambda_1^{\hat{l}} + \tau$$

Thus, the firm's net gain from playing $\hat{r}$ over $\hat{l}$ is given by

$$\beta(\lambda_1^{\hat{r}} - \lambda_1^{\hat{l}}) - \beta\tau(Pr(\hat{s}_2 = \hat{r}|\hat{s}_1 = \hat{r}) - Pr(\hat{s}_2 = \hat{r}|\hat{s}_1 = \hat{l}))$$

Since $\tau > 0$, we know $\lambda_1^{\hat{r}} - \lambda_1^{\hat{l}} > 0$. Pick any ϵ such that $0 < \epsilon < \lambda_1^{\hat{r}} - \lambda_1^{\hat{l}}$. We now that for a high enough p_0, we are guaranteed that

$$Pr(\hat{s}_2 = \hat{r}|\hat{s}_1 = \hat{r}) - Pr(\hat{s}_2 = \hat{r}|\hat{s}_1 = \hat{l}) < \epsilon.$$

This would mean that the firm has a strictly positive net gain from playing $\hat{r}$, and thus will deviate and always play $\hat{l}$. This proves that the myopic equilibrium does not hold in this case. $\square$