Spring 2021

Essays on Dynamic Games

Io Kuan Vong
Yale University Graduate School of Arts and Sciences, allenvongecon@gmail.com

Abstract

Essays on Dynamic Games

Io Kuan Vong

2021

This dissertation examines several economic questions related to dynamic incentives. In the first chapter, I study a dynamic game in which an expert advises a sequence of principals on their actions to match a hidden, randomly evolving state. The expert privately knows her competence. The principals learn about the state and the expert's competence from past advice and past action outcomes, both publicly observable. I find that the equilibrium can feature a "crisis of expertise," in which principals dismiss a competent expert's correct advice, and rely only on public information. Notably, the crisis happens precisely when the quality of public information is low, and thus when the competent expert's knowledge is much needed. Finally, I discuss policy implications for alleviating the crisis.

In the second chapter, I study the design of certification schemes for a firm who faces a sequence of consumers confronted with adverse selection and dynamic moral hazard problems. I show that a certifier extracts the full net market surplus and maximizes social welfare by using a low-standard, honors certification scheme: at each time, a firm who signed up for certification is either "certified" or "certified with honors;" moreover, an inept firm is always certified. I relate my findings to common concerns about the ability of such certification schemes to alleviate information asymmetries, despite the prevalence of these schemes in practice.

In the third chapter, I consider the strategic manipulation problem in multistage tournaments. In each stage, players are sorted into groups in which they play pairwise matches against each other. The match results induce a ranking over players in each group, and higher ranked players qualify to the next stage. Players prefer qualifying

to higher stages. In this setting, a player may potentially profit by exerting zero effort in some matches even when effort exertion is costless. Since such behavior manipulates the tournament, it is desired that full effort exertion is an equilibrium and any equilibrium ranking of qualifying players is immune to manipulation. To this end, I show that it is both necessary and sufficient to allow only the top-ranked player to qualify from each group. Otherwise, rankings can become a noisy indicator of players' strengths, while effort cost and heterogeneous prize spread can be of little relevance to players' effort choices.

Essays on Dynamic Games

A Dissertation
Presented to the Faculty of the Graduate School
of
Yale University
in Candidacy for the Degree of
Doctor of Philosophy

by
Io Kuan Vong

Dissertation Director: Johannes Hörner and Larry Samuelson

June 2021

Contents

List of Figures

List of Tables

Acknowledgements

I am immensely indebted to my advisors, Johannes Hörner, Larry Samuelson and Marina Halac. Their advice, encouragement, enthusiasm, time and patience shaped not only my completion of this dissertation, but also my personal and professional development. It has truly been a great privilege to be their student, and I thank them for all I have learned from them.

For helpful comments on earlier drafts of my papers, I thank Georgy Artemov, Ian Ball, Dirk Bergemann, Alessandro Bonatti, Tilman Börgers, Yi Chen, Joyee Deb, Eduardo Faingold, Florian Ederer, Péter Esö, Mira Frick, Tan Gan, John Geanakoplos, Carl Heese, Ryota Iijima, Patrick Lahr, Jan Knöpfle, Erik Madsen, George Mailath, Chiara Margaria, Meg Meyer, Weicheng Min, Xiaosheng Mu, Marco Ottaviani, Aniko Öry, Yujie Qian, Anna Sanktjohanser, Eran Shmaya, Philipp Strack, Juuso Välimäki, Stephan Waizmann, Xinyang Wang, Jun Xiao and Kai Hao Yang. I especially thank Claudio Mezzetti for his support in the early years of my graduate studies.

I thank my fellow PhD students, in particular Daisuke, Jinwook, Jeff, Jonathan, Lucas, Luís Fabiano, Marianne, Oscar, Soonwoo and Sukjoon, for making my years at Yale among the most memorable of my life.

To Lei Kei, thank you for putting up with me during my hardships.

Most importantly, I thank my parents, Agnes and Ken, and my sister, Kennes, for their love and support. And to Pinky, thank you for your love and inspiration, and for standing by me through thick and thin.

Chapter 1

The Crisis of Expertise

1.1 Introduction

Experts give advice to guide decisions. Popular discussions suggest that technology can lead to a "crisis of expertise" by making information abundant and publicly accessible to decision makers (see, *e.g.*, Nichols, 2017; Gurri, 2018; Eyal, 2019). A common description of this crisis is that decision makers dismiss the genuine advice that informed experts offer, and act solely on the basis of public information, for instance, from social media. Indeed, Nichols (2017, p. 105) writes: "Ask any professional or expert about the death of expertise, and most of them will immediately blame the same culprit: the Internet." Similarly, in a 2019 interview with *CNBC*, a financial planner explains that this is the reason why people often stay away from getting professional financial advice: "There is a lot more information online these days, [...] so people feel like they can do it themselves."[1]

How does the crisis of expertise affect the value of informed experts who offer genuine advice? Is the crisis due to the abundance of public information? And if so, is it simply a phenomenon where decision makers substitute high-quality public

1. "99% of Americans don't use a financial advisor—here's why," *CNBC*, November 11, 2019.

information for expert advice?

This paper addresses these questions in a simple model of expert advice. Decision makers take actions after assessing two relevant pieces of information—current expert advice and public information, the latter consisting of past advice and past action outcomes. My analysis proposes that the access to public information allows "fake" experts to benefit by pretending to give informed advice, despite them knowing no more than the decision makers. The presence of these fake experts causes skeptical decision makers to sometimes dismiss expert advice and act solely on the basis of public information. A striking prediction of my model is that decision makers dismiss expert advice only when the quality of public information is low, and thus when informed experts' knowledge is much needed. Hence, the crisis of expertise is a consequence of the lack, as opposed to the abundance, of high-quality public information.

I describe my model in Section 1.2. An expert faces a sequence of principals who arrive at random times. Each principal takes an action and wants his action to match a binary state. The principal faces uncertainty, as the state is hidden and evolves according to a Markov process. Before taking an action, the principal hires the expert by paying her upfront for advice, with the payment reflecting his perceived value of the advice. Once hired, the expert receives a private signal about the current state. The expert is privately informed about her signal quality: she is either an "informed" type who observes the state perfectly, or an "uninformed" type whose signal of the state is uninformative. The expert's advice and the state are publicly revealed after the principal takes an action. The goal of the expert is to maximize the payments she receives from the principals.

A central feature of the model is that the principals have noisy, evolving beliefs about both the state and the expert's type given the public information, namely, the history of past pieces of advice and what the state turned out to be at past times. The model is applicable to a variety of advising and consulting situations. To give a concrete example, suppose that the expert is a financial analyst and the principals are firms which choose to raise capital to fund either a risky project or a safe project. A project succeeds only if enough investors invest in it. The state of the market varies over time and can be "fear" or "greed;" if the market is fearful (resp., greedy), then only the safe (resp., risky) project receives enough funds. The firms are unsure about the state and therefore hire the analyst. They evaluate the advice, and then determine their optimal decisions, based on the past capital-raising outcomes by other firms as well as the past advice given by the financial analyst.

As the preceding discussion motivates, the goal of my analysis is to elucidate how the principals value genuine, informed advice given the public information. In Section 1.3, I thus characterize equilibria where the informed expert truthfully reports her private information relevant to the principal's decision-making, namely her type and current private signal, in each principal's arrival.

In any equilibrium, it is impossible for an informed type's report to convey information about her type to each principal, due to the uninformed type's incentive to appear informed. Hence, the expert's report in each arrival is effectively a report of the current state. The informed type reports her state observation truthfully. Instead, the uninformed type mixes between reporting either state in order to "gamble" for her reputation, namely the principals' belief that she is an informed type. To see how this mixture is determined, call the state that is most likely given the public information the *agreeing* state, and the other state is the *disagreeing* state. From the uninformed type's perspective, disagreeing is likely to be incorrect, and hence, to reveal her type. Yet, if it is correct, her reputation will jump up and given the equilibrium strategies, so will the principals' future payments. The uninformed type may thus want to gamble on the unlikely event that disagreeing is correct.

From each principal's perspective, if he knew that the expert is informed, then he would match his action with the experts report. In contrast, if the principal knew that the expert is uninformed, then he would follow the public state belief and match his action with the agreeing state, irrespective of the report. Because of the uninformed expert's gambling incentives, I show that there is a region of public beliefs about the state and the expert's type which corresponds to a crisis of expertise. Principals who arrive with beliefs in this region dismiss the expert's advice, that is, they match their actions with the agreeing state independently of the expert's report. Importantly, the crisis region excludes high enough, inconclusive state beliefs. This is because a principal who arrives with a high state belief, namely high-quality public information about the state, knows that the uninformed type shares the same state belief and has little incentive to gamble, as disagreeing is very likely to be incorrect.

In Section 1.4, I turn to the implications of my analysis, addressing the motivating questions in the opening paragraphs. First, I show that although the uninformed type's gambling undermines the value of expert advice, the damage it causes on the informed expert's payoff dissipates in the long run. In particular, the crisis of expertise is a short-run phenomenon, as the public information allows the informed expert to build her reputation. When the reputation grows sufficiently high, each arriving principal expects that the uninformed type is less likely to gamble, and thus matches his action with the expert's report. I relate my findings to concerns in the market for retail financial advice that the presence of fake experts undermine the value of informed advisors and their genuine advice.

Second, I discuss how my analysis speaks to the two faces of public information technologies, like social media, for decision-making. On the one hand, my findings predict that these technologies give decision makers access to resources that inform their decision-making and help them identify the expert's informativeness. On the other hand, my results show that these technologies also provide the uninformed expert with the same resources and endogenously lead to a convention that rewards successful contrarians. Such conventions are pervasive in practice and provide incentives for uninformed experts to gamble at the expense of decision makers. I relate these observations to recent empirical evidence on gambling by financial analysts.

Finally, I point to policies for alleviating the crisis of expertise, and relate my findings to recent advocates by political campaigns and consumer groups to equip consumers with high-quality public information. I argue that such policies are most needed in economic environments where state transitions occur neither very often nor very rarely. Specifically, my analysis predicts that such environments provide the uninformed type with the strongest incentives to gamble for her reputation.

My model belongs to the literature on expert reputation, and particularly relates to reputational cheap talk models. These are models where a careerist expert either herds with the public state belief (*e.g.*, Scharfstein and Stein, 1990; Prendergast, 1993; Zwiebel, 1995; Ottaviani and Sørensen, 2006a,b,c), or anti-herds with the public state belief (*e.g.*, Trueman, 1994; Avery and Chevalier, 1999; Levy, 2004).[2] The idea that experts might gamble for their reputations is familiar from these anti-herding models, albeit in different contexts and under different assumptions.

To characterize the crisis of expertise and its dynamics, my model has three central and distinguishing features relative to the existing reputational cheap talk models. First, the state evolves randomly over time.[3] Second, my model features receivers, *i.e.*, the principals, who take actions upon receiving advice. These two features allow me to characterize why a principal may act solely on the basis of public information and independently of expert advice.

2. The idea that experts strive to appear informed has motivated a large literature. See also Brandenburger and Polak (1996), Effinger and Polborn (2001), Prat (2005), Li (2007), Klein and Mylovanov (2017), Pavesi and Scotti (2019), Rüdiger and Vigier (2019), Smirnov and Starkov (2019) and Backus and Little (2020).

3. Indeed, my analysis takes a step towards the direction of developing a dynamic model of reputational cheap talk. When discussing directions for future research in a survey of the literature, Marinovic, Ottaviani, and Sørensen (2013) write that "a key challenge lies in finding a tractable and sufficiently general multi-period environment with learning about the precision [of the expert's information] as well as about the state."

Third, the expert's payoff in my model depends on the principals' perceived value of her advice, instead of being a reduced-form "reputational payoff," as is often assumed in the literature, that depends solely on the expert's reputation. This departure means that the expert's reputation concern is instrumental in my model: even if the principals believe that the expert is likely to be informed, they pay her a positive amount only if they expect the advice to convey useful information. As a result, the two expert types continuation payoffs typically differ from each other given any equilibrium history as they differ in the ability to give correct advice. Together, the three departures allow my analysis to examine the effect of uninformed experts' gambling on the value of informed experts and their genuine advice, to shed light on the formation of contrarian-rewarding conventions as an equilibrium outcome, to highlight how uninformed experts' gambling incentives depend on the uncertainty of the economic environment, as well as to capture the "fly-by-night" nature that uninformed experts reveal their identity in the short run. These phenomena are novel to the literature and are prominent in practice.

The basic incentive structure of my model is reminiscent of the bad reputation model a là Ely and Välimäki (2003). Nonetheless, my model does not feature bad reputation. In both models, the efficient outcome is attainable in a one-shot interaction as well as in repeated interactions with complete information on the expert's type. Moreover, in both models, the efficient outcome is unattainable in repeated interactions when there is incomplete information on the expert's type. In Ely and Välimäki (2003), the principals cannot perfectly verify the expert's signal at the end of each interaction. In turn, a "good" expert's desire to separate herself from a "bad" expert creates perverse incentives for the good expert to act against principals' interests. In contrast, in my model, such perverse incentives are absent, as the principals can perfectly verify whether the expert's report was correct following each action.

The broader literature on expert reputation examines models with different incentive structures. One strand of the literature consists of models where experts take costly actions to appear competent, starting with the seminal work by Holmström (1999).[4] My analysis is related to Prendergast and Stole (1996) in that a careerist agent's reaction to new information changes over time. In their model, this phenomenon is due to the agent's incentive to avoid revealing that she has made mistakes in the past. The agent's behavior remains inefficient in the long run. In my model, the expert's behavior changes over time for different reasons. The informed type's behavior changes in response to her evolving state belief. On the other hand, the uninformed type's behavior changes in response to both her evolving state belief and her evolving reputation. In particular, her behavior entails less efficiency losses as time progresses. Another strand examines experts who communicate to bias the principals' decisions, and strive to appear unbiased (*e.g.*, Sobel, 1985; Benabou and Laroque, 1992; Morris, 2001; Gentzkow and Shapiro, 2006).

There is also a literature on testing experts, starting with Foster and Vohra (1998), that aims to determining the existence of a test which cannot be passed by an expert without knowledge of the true data-generating process but can be passed almost surely by an expert who knows the process.[5] Recently, Deb and Stewart (2018), Deb, Pai, and Said (2018) and Deb, Mitchell, and Pai (2020) examine dynamic mechanisms that screen experts' ability. Different from these models, the expert in my model faces implicit incentives to appear informed, as her ability is evaluated by a "market," namely a sequence of short-lived principals, without commitment power.

4. See also Dewatripont, Jewitt, and Tirole (1999a,b), Majumdar and Mukand (2004), Bonatti and Hörner (2017), Cisternas (2018), and Halac and Kremer (2020).

5. For a survey, see Olszewski (2015).

1.2 Model

Setup. Time $t \in [0, \infty)$ is continuous, and the horizon is infinite. A long-lived expert faces a countable sequence of short-lived principals. The expert has a private type $\theta \in \Theta := \{\theta_I, \theta_U\}$. She is an informed type θ_I with probability $p_0 \in (0, 1)$, and is an uninformed type θ_U otherwise. The principals arrive according to a Poisson process $(N_t)_{t\geq 0}$ with intensity normalized to one. A hidden state $(s_t)_{t\geq 0}$, taking values on $S := \{\underline{s}, \bar{s}\}$ with initial value $s_0 = \bar{s}$, evolves according to a Markov process with switching rate (or volatility) $\gamma \in (0, \infty)$ at each state.[6] The expert's type, the state process and the arrival process are independent.

Actions and flow payoffs. When a principal arrives at time t, he pays the expert a wage $w_t \in \mathbb{R}_+$ upfront.[7] The wage equals the principal's marginal gain from interacting with the expert, as I specify below in (1.1). The expert sends a message $m_t \in M$ to the principal, where M is a finite set. The principal then takes an action $a_t \in S$, obtaining a benefit $b_t = 1$ if $a_t = s_t$ and $b_t = -1$ otherwise. The principal's flow payoff equals his benefit minus the wage, $b_t - w_t$. The expert's flow payoff equals the wage w_t.

The flow payoffs are discounted at rate $r > 0$. Given an infinite history of arrival times $(t_i)_{i=1}^{\infty}$ and wages $(w_{t_i})_{i=1}^{\infty}$ at those times, the expert's (normalized) payoff is

$$\sum_{i=1}^{\infty} r\, \mathrm{e}^{-rt_i}\, w_{t_i}.$$

6. The analysis readily extends to allowing for general initial state distributions, as well as asymmetric switching rates at the cost of extra notation. Section 1.5 discusses the extension to more states.

7. Although output-contingent contracts are theoretically attractive, my model assumes that wages are paid upfront, as is often the case in my motivating applications such as financial advice and business consulting (see, *e.g.*, McLachlin, 2000; Momparler, Carmona, and Lassala, 2015).

Information. At each principal's arrival time t, before sending a message, the expert receives a private signal $y_t \in S$. If the expert is an informed type, then the signal is equal to the current state. If the expert is an uninformed type, then the signal is equal to either $\bar{s}$ or $\underline{s}$ with equal probabilities, and is therefore pure noise. In addition, in each arrival, the principal's action and his realized benefit are public; the current state is thus publicly revealed once the benefit is realized.

Histories. A public history of length n is a list

$$h_n = (t_i, m_{t_i}, a_{t_i}, b_{t_i})_{i=1}^n,$$

collecting the times $(t_i)_{i=1}^n$ of the first n arrivals, and the message m_{t_i}, the action a_{t_i} and the benefit b_{t_i} in each such time-t_i arrival. A type θ's private history of length n is a list

$$h_n^\theta = (h_n, (y_{t_i})_{i=1}^n),$$

consisting of the length-n public history and her private signals in the first n arrivals. Let H_n be the set of length-n public histories, let $H := \cup_{n=0}^\infty H_n$, and let $h \in H$ denote a generic public history. Likewise, let H_n^θ be the set of type θ's length-n private histories, let $H^\theta := \cup_{n=0}^\infty H_n^\theta$, and let $h^\theta \in H^\theta$ denote a generic type θ's private history.

Strategies. Each expert type θ's strategy is a collection $\sigma^\theta = (\sigma_t^\theta)_{t\geq 0}$, where each

$$\sigma_t^\theta : H^\theta \times S \to \Delta M$$

is a measurable mapping that takes her history h^θ and the current signal y to a distribution over messages at each arrival time t, subject to the restriction that all arrival times in the history h^θ are strictly smaller than t. Let $\sigma := (\sigma^I, \sigma^U)$.

Each principal's strategy is to choose an action given a public history h and a current message m at his arrival time t. Since the principal is short-lived, he chooses the myopically optimal action and obtains a benefit that I denote by $b_t^*(h, m)$. Moreover, whenever a principal is indifferent between the two actions, without loss, I assume that he follows a pure strategy of matching his action to the last revealed state (hereafter, the last state); this follows from the fact that the principal's action does not affect the information of the expert or future principals as well as future wages. Finally, when specifying an equilibrium below, I omit the principals' strategies to ease the notation.

Beliefs. At each time t, the uninformed expert forms a public belief that the current state equals the last state. I denote this public belief by q_t. Because state transitions are Markov, q_t is plainly a function of the time-lapse l since the last state was revealed and is given by $q_t = \frac{1}{2}(1 + \mathrm{e}^{-2\gamma l}) =: \mu_l$. This public state belief is depicted in Figure 1.1. Crucially, $\mu_l > \frac{1}{2}$ for each $l \geq 0$. Thus, between the two possible states, the public state belief dictates that the current state is more likely to match the last state than to mismatch it. The belief falls as the time-lapse l grows, and it falls at a faster rate with l the higher is the state volatility γ.

I denote by q_t^I the informed expert's time-t belief that the current state equals the last state. At each time t absent a principal's arrival, q_t^I coincides with the public state belief q_t. At each time t when a principal arrives, the informed expert's state belief q_t^I jumps to either 1 or 0, depending on the private signal that she receives.

Given a strategy σ and a public history h at time t, an arriving principal forms a belief on the expert's type and the current state, *i.e.*, on the set $\Theta \times S$. This belief is a pair (p_t, q_t), where p_t denotes the public belief that the expert is informed and $q_t = \mu_l$ is the public belief that the state is equal to the last state given a time-lapse l. The belief p_t represents the expert's reputation at time t. Throughout, I refer to a

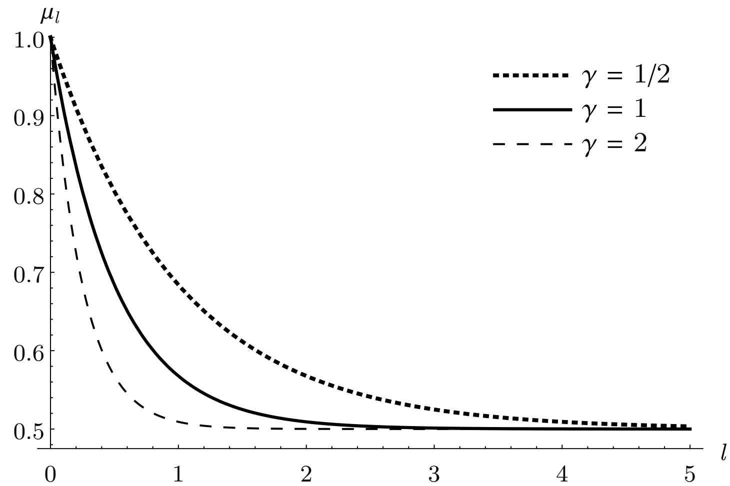

Figure 1.1: Public state belief

generic pair (p, q), or (p, μ), as the public type and state beliefs or simply the public beliefs.

Wage. In each arrival, the wage that a principal pays the expert equals his expected marginal benefit from receiving the expert's advice. If a principal with state belief μ had to choose an action without the expert's advice, he would have optimally matched his action with the last state, deriving an expected benefit of $2\mu - 1$. Thus, given a strategy σ and a history of play where a principal arrives at time t with associated public history h and state belief μ, the wage at the arrival is

$$w_t := \mathbf{E}[b_t^*(h, m)|h] - (2\mu - 1), \tag{1.1}$$

where the expectation is taken over messages m that the expert sends according to the strategy σ, conditional on the public history h.

Equilibrium. The solution concept that I use is perfect Bayesian equilibrium. In such an equilibrium, denoted by (σ, φ), each type's strategy σ^θ in the profile σ is optimal given the beliefs $\varphi := (p_t, q_t, q_t^I)_{t\geq 0}$ and the beliefs are consistent with the strategy profile σ wherever possible. The set of messages is an ingredient of the solution concept. In the spirit of Myerson (1986), I assume that $M = \Theta \times S$, allowing each type to disclose all her relevant private information in each arrival. A message is a pair $m = (\hat{\theta}, \hat{y})$, interpreted as an expert's report on her type and current signal.

I focus on perfect Bayesian equilibria that satisfy Properties 1.1 and 1.2 below, henceforth "equilibria." Property 1.1 concerns the informed type's behavior on path:

Property 1.1. *On path, in each arrival, the informed type reports her private information truthfully, i.e., upon observing a current signal $y = s$, she sends a message $m = (\theta_I, s)$.*

Property 1.1 can be interpreted as a social convention, where the informed type is expected to give advice on the basis of some code of ethics or fiduciary duty. Given Property 1.1, it is without loss to assume that an uninformed type never sends the messages $(\theta_U, \bar{s})$ or $(\theta_U, \underline{s})$, as these messages fully reveal her type and lead to zero future wages. Hence, I further simply the structure of the message set and maintain the assumption throughout that

$$M = S.$$

Each message is therefore interpreted as a report on the current signal. I say that an action, a signal or a report is *correct* if it matches the current state, and is *incorrect* otherwise. I also say that an action, a state, a signal or a report is *agreeing* with respect to the public state belief if it matches the last state, and is *disagreeing* otherwise.

In equilibrium, beliefs are updated by Bayes rule not only on path but also off path. Because Bayes' rule does not apply when beliefs are degenerate, an assumption is necessary to pin down what happens following degenerate beliefs off path. Property 1.2 takes care of this:

Property 1.2. *In each arrival with beliefs (p, μ) where either p or μ is degenerate, if the expert sends a message that is inconsistent with the strategy and the beliefs, then* *the expert's reputation jumps to zero.*

The analysis readily extends to allowing for other off-path beliefs. In Section 1.5, I discuss the consequences of doing so.

1.3 Equilibrium

In this section, I present my main result, characterizing the equilibrium. To streamline the presentation, I describe the result in a series of propositions. I begin by deriving the basic structure of any equilibrium. I then show that an equilibrium exists. Finally, I show that all equilibria exhibit qualitatively identical properties, and examine the properties in detail. Proofs are in the Appendix. In Section 1.4, I interpret the formal results in this section and discuss their implications.

1.3.1 Basic Structure

In any equilibrium, by Property 1.1, the informed type reports truthfully on path. Much difficulty in characterizing equilibria rests upon characterizing the uninformed type's strategy, because the set of her strategies is large. I first derive properties that the uninformed type's continuation payoffs and strategies must satisfy in any equilibrium. I then examine the principals' best responses given the uninformed type's equilibrium strategies.

Proposition 1.1. *In any equilibrium, the following properties hold:*

1. *Payoffs—there is a strictly increasing function* $V^U : [0,1] \to \mathbb{R}_+$ *such that at any uninformed type's history* h^U *where an arrival ends with reputation* p *after the principal's benefit is realized, the uninformed type's continuation payoff is* $V^U(p)$*. Moreover,* $V^U(0) = 0$*.*

2. *Behavior—in each arrival, the uninformed type's strategy is Markov in the public beliefs* (p, μ).

The monotone payoff structure reflects the fact that the uninformed expert is motivated by reputation concerns to appear informed. As the principals expect the informed type to truthfully report her signals, their willingness to pay strictly increases in the expert's reputation, holding fixed the uninformed type's continuation strategy. The uninformed type thus benefits from a higher reputation. When the reputation is zero, the principals do not find the expert's reports to be valuable, so the uninformed expert's continuation payoff equals zero.

In each arrival, given binary states, without loss, I take the uninformed expert's strategy to be a probability of sending an agreeing report. The uninformed expert chooses the agreeing probability to maximize her continuation payoff upon sending a report. This continuation payoff depends solely on the public state belief that the report is correct and the expert's current reputation; the latter determines the resulting reputation if the report turns out to be correct, given the chosen agreeing probability. (If the report turns out to be incorrect, the resulting reputation is zero.) Thus, her strategy is Markov in the public beliefs.

The Markovian structure simplifies the task of characterizing equilibria:

Corollary 1.1. *If an equilibrium exists, then it is characterized by some agreeing function* $\alpha : [0,1] \times (\frac{1}{2}, 1] \to [0,1]$ *that determines the uninformed expert's agreeing probability* $\alpha_{p,\mu}$ *in each arrival given public beliefs* $(p, \mu) \in [0,1] \times (\frac{1}{2}, 1]$.

Given that the principals expect the informed type to report truthfully, the informed type has no profitable deviation to misreport in an arrival. From her perspective, a deviation to misreport must result in a reputation of zero and thus zero future wages, giving her the lowest possible continuation payoff. Thus, without loss, I assume that upon any deviation by the informed type, she sends an agreeing report in every future arrival.

Equilibrium existence

By Corollary 1.1, any equilibrium is characterized by some uninformed expert's agreeing function given which the uninformed type finds no profitable deviation when principals the uninformed type's play follows such an agreeing function. Hence, to prove that an equilibrium exists, it suffices to show that there exists an agreeing function α that maximizes the uninformed type's equilibrium payoff, and then to verify that when the principals expect the uninformed type's play follows such an agreeing function, the uninformed type has no profitable deviation. The next result shows that an equilibrium exists.

Lemma 1.1. *An equilibrium exists.*

In general, there are multiple equilibria, each characterized by an agreeing function given which the uninformed type finds no profitable deviation when principals the uninformed type's play follows such an agreeing function. In the following, I show that all equilibria exhibit qualitatively identical properties. I also describe these properties in detail, as they are crucial for understanding the crisis of expertise.

Fix an arbitrary agreeing function $\bar{\alpha}$ that characterizes an equilibrium. Since the function $\bar{\alpha}$ is arbitrarily chosen, hereafter, I refer to the equilibrium that the function characterizes simply as the equilibrium.

In equilibrium, the function $\bar{\alpha}$ must satisfy the uninformed type's incentive constraint after receiving the upfront payment in each arrival on path with beliefs (p, μ),

$$\mu V^U\left(\frac{p}{p+(1-p)\bar{\alpha}_{p,\mu}}; \bar{\alpha}\right) - (1-\mu)V^U\left(\frac{p}{p+(1-p)(1-\bar{\alpha}_{p,\mu})}; \bar{\alpha}\right) \begin{cases} \geq 0, & \text{if } \bar{\alpha}_{p,\mu} = 1, \\ = 0, & \text{if } \bar{\alpha}_{p,\mu} \in (0,1), \\ \leq 0, & \text{if } \bar{\alpha}_{p,\mu} = 0. \end{cases} \tag{1.2}$$

The left side of the constraint is the difference between the uninformed expert's continuation payoff by agreeing and by disagreeing. From her perspective, an agreeing report is correct with probability μ. While an incorrect report leads to zero reputation (and hence zero continuation payoff), a correct agreeing report leads to reputation

$$\frac{p}{p + (1-p)\bar{\alpha}_{p,\mu}}. \tag{1.3}$$

Likewise, her continuation payoff by disagreeing, noting that such report is correct with probability $1 - \mu$ and a correct such report leads to reputation

$$\frac{p}{p + (1-p)(1-\bar{\alpha}_{p,\mu})}. \tag{1.4}$$

When the reputation p is positive, uniqueness of the agreeing probability arises from the uninformed type's reputation concerns. By choosing too high an agreeing probability, the reputation reward by sending a correct, disagreeing report is large, as principals expect that such report is more likely to be sent by an informed type. Likewise, by choosing too low an agreeing probability, the reputation reward by sending a correct, agreeing report is large, as principals expect that such report is more likely to be sent by an informed type.

Finally, if the reputation p is zero, then the uninformed expert's continuation payoff is zero. Thus, any strategy by the uninformed type is optimal. Without loss, I set $\bar{\alpha}_{0,\mu} = \frac{1}{2}$.

Principals' Best Responses

I next examine the principals' best responses given the uninformed type's strategy $\bar{\alpha}$. From each principal's perspective, if he knew that the expert is informed, he would match his action with the expert's report. In contrast, if he knew that the expert is uninformed, then he would optimally choose an agreeing action irrespective of the report. Because the principal is unsure of the expert's type, the equilibrium outcome is constrained inefficient.

In each arrival with beliefs (p, μ) , given that the informed type reports truthfully and the uninformed type reports according to the agreeing function $\bar{\alpha}$, the principal optimally matches his action with an agreeing report, as the agreeing report reinforces the public belief that the agreeing action is more likely to be correct. Specifically, the principal's belief that the report is correct given that it is disagreeing strictly exceeds one half irrespective of the value of $\bar{\alpha}_{p,\mu} \in [0, 1]$:

$$\underbrace{\frac{p\mu}{p\mu + (1-p)\bar{\alpha}_{p,\mu}}}_{\substack{\text{principal's belief that expert is} \\ \text{informed upon agreeing report}}} \times 1 + \left(\underbrace{1 - \frac{p\mu}{p\mu + (1-p)\bar{\alpha}_{p,\mu}}}_{\substack{\text{principal's belief that expert is} \\ \text{uninformed upon agreeing report}}} \right) \times \mu > \frac{1}{2}.$$

In contrast, the principal matches his action with a disagreeing report if and only if

$$\underbrace{\frac{p(1-\mu)}{p(1-\mu) + (1-p)(1-\bar{\alpha}_{p,\mu})}}_{\substack{\text{principal's belief that expert is} \\ \text{informed upon disagreeing report}}} \times 1 + \left(\underbrace{1 - \frac{p(1-\mu)}{p(1-\mu) + (1-p)(1-\bar{\alpha}_{p,\mu})}}_{\substack{\text{principal's belief that expert is} \\ \text{uninformed upon disagreeing report}}} \right) \times (1-\mu) > \frac{1}{2}.$$

The left side is the principal's belief that the report is correct given that it is disagreeing. The inequality is equivalent to writing

$$\bar{\alpha}_{p,\mu} > \kappa_{p,\mu} := \frac{\mu(2-p) - 1}{(2\mu - 1)(1-p)}. \tag{1.5}$$

Condition (1.5) admits a intuitive interpretation. The principal matches his action with a disagreeing report if and only if he is sufficiently confident that the disagreeing report is sent by an informed type. This is the case when the uninformed types agreeing probability is sufficiently high.

1.3.2 Characterization

In this section, I characterize the equilibrium. Given the agreeing function $\bar{\alpha}$, define a correspondence $C : [0,1] \rightrightarrows (\frac{1}{2}, 1]$ such that

$$C(p) := \{\mu \in (\tfrac{1}{2}, 1] : \bar{\alpha}_{p,\mu} \leq \kappa_{p,\mu}\}. \tag{1.6}$$

Given reputation p, the set $C(p)$ contains the state beliefs μ given which, in view of the preceding discussion, the agreeing probability $\bar{\alpha}_{p,\mu}$ is sufficiently low that principals prefer to ignore the expert's advice and choose the agreeing action. Summarizing the previous discussion:

Proposition 1.2. *In the equilibrium, in each arrival with beliefs (p, μ), if $\mu \in C(p)$, then the principal chooses an agreeing action independently of the expert's report. Otherwise, the principal matches his action with the expert's report.*

In view of Proposition 1.2, hereafter, I refer to graph(C) as the *crisis region.* I refer to a history of play at which the public beliefs (p, μ) are in the crisis region conditional on the expert being informed as a *crisis of expertise.* In the equilibrium, wages reflect the crises.

Corollary 1.2. *In the equilibrium, in each arrival with beliefs (p, μ), the wage is zero if (p, μ) is in the crisis region and is positive otherwise.*

In the equilibrium, in each arrival with beliefs (p, μ), the wage is

$$\max\Big(\underbrace{p + (1-p)\bar{\alpha}_{p,\mu}(2\mu - 1) - (2\mu - 1)}_{\substack{\text{marginal benefit of matching} \\ \text{action with expert's report}}}, 0\Big). \tag{1.7}$$

The arriving principal finds expert advice to be valuable if and only if he anticipates to match his action with the expert's report, that is, the principals' perceived marginal benefit of matching his action with the expert's report is positive.

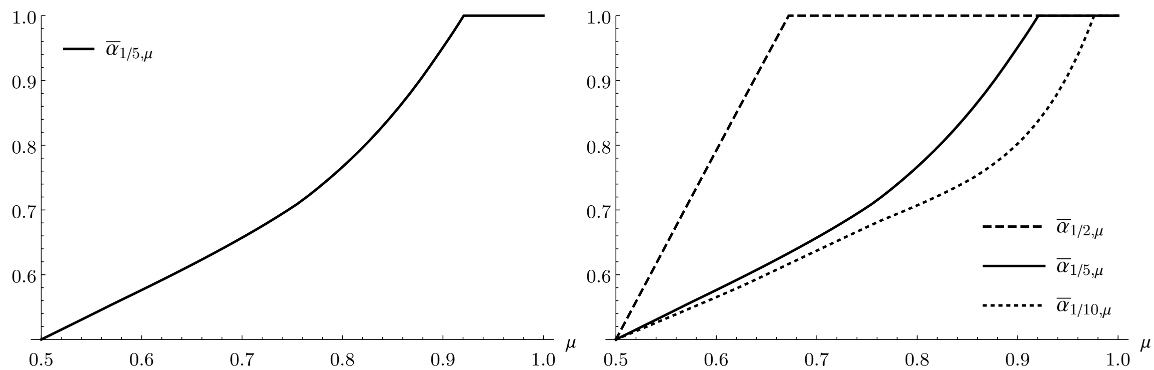

(a) $p = \frac{1}{5}$ **(b)** $p = \frac{1}{2}, \frac{1}{5}, \frac{1}{10}$

Figure 1.2: Agreeing probability ($\gamma = 1$, $r = 1$)

In the remainder of this section, I explore in detail the structure of the agreeing function $\bar{\alpha}$ and the crisis correspondence C. I begin with the agreeing function, elucidating the uninformed type's incentives to "gamble" for her reputation:

Proposition 1.3. *For every (p, μ), $\bar{\alpha}_{p,\mu} \geq \frac{1}{2}$. Moreover:*

1. *Lower state belief, more gambling—for every $p > 0$, there is $\mu^* \equiv \mu^*(p)$ such that $\bar{\alpha}_{p,\mu}$ strictly increases in μ on $(\frac{1}{2}, \mu^*]$ and equals 1 if $\mu \in (\mu^*, 1]$.*

2. *Lower reputation, more gambling—for every $\mu \in (\frac{1}{2}, 1]$, $\bar{\alpha}_{p,\mu}$ increases in p; further, $\mu^*(p)$ strictly decreases in p, satisfying $\mu^*(0) = 1$ and $\mu^*(1) = \frac{1}{2}$.*

The agreeing probability always exceeds $\frac{1}{2}$ due to the uninformed type's incentive to pool with the informed type. From the uninformed type's perspective, the informed type is more likely to agree than to disagree. In turn, in equilibrium, the reputation upon correctly disagreeing, as given by (1.4), must strictly exceed that upon correctly agreeing, as given by (1.3). By Proposition 1.1, the uninformed expert's reputation concerns provide her with incentives to disagree, gambling on the less likely event that disagreeing is correct for the larger reputation reward if disagreeing turns out to be correct.

Part 1 of Proposition 1.3 says that public information of high quality about the state disciplines the uninformed type's gambling. Given a higher state belief, the uninformed type's gambling incentive is lower as she believes that disagreeing is more likely to be incorrect. State beliefs above the cutoff μ^* "preempt" the uninformed type's gambling. On the other hand, state beliefs that are arbitrarily close to $\frac{1}{2}$ give the strongest gambling incentive to the uninformed type, who then reports both states with approximately equal probabilities. To illustrate, Figure 1.2a numerically plots the agreeing probability $\bar{\alpha}_{p,\mu}$ against the state belief μ, assuming that the reputation is $p = 1/5$, in a setting where the state volatility is $\gamma = 1$ and the discount rate is $r = 1$.

Part 2 of the proposition says that given a higher current reputation, the uninformed type agrees with a higher probability. This is because the reputation gain by correctly disagreeing is then smaller. To illustrate, Figure 1.2b numerically plots the agreeing probability against the state belief for several values of reputations, again taking $\gamma = 1$ and $r = 1$.

My next proposition concerns the structure of the crisis correspondence C and elucidate the principals' incentives.

Proposition 1.4. *For every p, the correspondence C satisfies $1 \in C(p)$. Moreover:*

1. *Given a positive reputation, high enough interior state beliefs preempt the crisis—for every $p > 0$, there exists $\bar{\mu} < 1$ such that $\mu \notin C(p)$ if $\mu \in [\bar{\mu}, 1)$.*

2. *Given a positive reputation, low enough state beliefs preempt the crisis—for every $p > 0$, there exists $\underline{\mu} < 1$ such that $\mu \notin C(p)$ if $\mu \in (\frac{1}{2}, \underline{\mu}]$.*

3. *Higher reputation p, smaller $C(p)$—there is $\bar{p} \in [0, 1)$ such that $C(p) \subsetneq C(p')$ if $p' < p < \bar{p}$. Further, $C(0) = (\frac{1}{2}, 1]$ and $C(p) = \{1\}$ if $p \geq \bar{p}$.*

The set $C(p)$ contains the unit state belief for every reputation p. When arriving with a unit state belief, each principal chooses an agreeing action as his state belief stays at 1 irrespective of the report he receives. Indeed, on path, both expert types

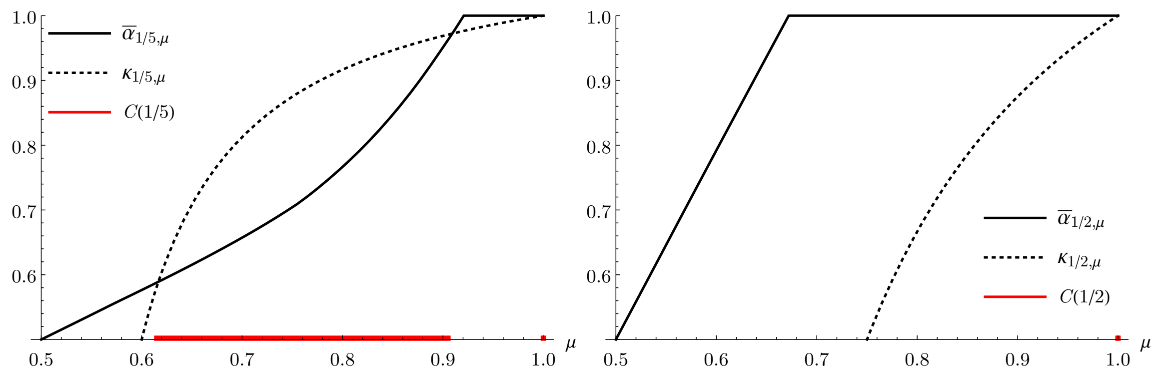

(a) $p = 1/5$ **(b)** $p = 1/2$

Figure 1.3: Agreeing probability and cutoff ($\gamma = 1$, $r = 1$)

agree in such arrival, inducing an agreeing action by the principal. On the other hand, if the expert deviates to disagree, then by Property 1.2, the principal attaches probability zero to the expert being informed and thus chooses the agreeing action.

Part 1 of Proposition 1.4 says that in any arrival where the reputation p is positive, the set $C(p)$ excludes high enough interior state beliefs. This shows that unless the public information about the state is conclusive, an arriving principal chooses an agreeing action on the basis of the public information irrespective of the expert's report only if the quality of public information about the state is low. A principal who arrives with a high enough state belief expects little gambling by the uninformed type, and thus matches his action with the report.

Part 2 says that the set $C(p)$ also excludes low enough state beliefs. A principal arriving with a low enough state belief matches his action with the expert's report, as he expects that the current state is sufficiently likely to have switched from the last state, and a disagreeing report is sent by an informed type.

Finally, Part 3 says that given a higher reputation, an arriving principal not only believes that the expert is more likely to be informed, but the principal also expects less gambling by the uninformed type in view of Proposition 1.3. To illustrate the result, consider Figure 1.3a. The set $C(\frac{1}{5})$ in the figure is the interval of state beliefs

μ given which $\bar{\alpha}_{1/5,\mu} \leq \kappa_{1/5,\mu}$, again taking $\gamma = 1$ and $r = 1$.[8] In contrast, in Figure 1.3b, the set $C(\frac{1}{2})$ contains only the unit state belief.

1.4 Implications

In this section, I leverage the equilibrium characterization to address the motivating questions in the opening paragraphs. I then discuss policy implications for alleviating the crisis.

1.4.1 Value of Informed Experts and Long-run Predictions

How does principals' access to public information affect the value of the informed expert and her genuine advice? In Section 1.3, I have been silent about the equilibrium payoff structure of the informed expert. Indeed, given that the informed expert reports the state truthfully in each arrival on path, it is *a priori* unclear how the uninformed expert's gambling affects the value of the informed expert to the principals. My next proposition shows that in the equilibrium, the informed type, like the uninformed type, also benefits from a higher reputation.

Proposition 1.5. *In the equilibrium, there is a strictly increasing function $V^I(\cdot;\bar{\alpha}) : [0,1] \to \mathbb{R}_+$ such that at any history h^I where an arrival ends with reputation p after the principal's benefit is realized, the informed type's continuation payoff is $V^I(p;\bar{\alpha})$. Moreover, $V^I(0;\bar{\alpha}) = 0$.*

The result follows because the equilibrium wage (1.7) strictly increases in the expert's reputation in view of Proposition 1.3. Given a higher reputation, principals

8. In general, a sharper characterization that addresses whether $C(p)$ defines a connected subset of state beliefs in $(\frac{1}{2}, 1]$ for any reputation p as in Figure 1.3a is an open question. The answer to this question hinges crucially upon the curvature of the agreeing probability $\bar{\alpha}_{p,\mu}$ in each state belief μ. As illustrated in Figure 1.2b, the curvature is in general non-monotone in the state beliefs, making an analytical characterization difficult to obtain due to the high dimensionality of the program $(*)$.

not only expect that the expert is more likely to be informed, but also that the uninformed type gambles less.

Further, Proposition 1.6 below shows that the informed expert's reputation increases over time, as her reports are always correct on path. Let $\mathbf{P}$ denote the probability measure over the set of outcomes of the game induced in the equilibrium.

Proposition 1.6. *In the equilibrium, on path, the reputation path* $(p_t)_{t\geq 0}$ *is increasing conditional on the expert being informed. Moreover, given any time* t*, for every* $\varepsilon > 0$*, there exists* $\underline{T} > 0$ *such that*

$$\mathbf{P}[p_t < p_{t+T}, T \geq \underline{T} | \theta = \theta_I] > 1 - \varepsilon.$$

The two propositions offer a perspective on why, despite concerns about the crisis of expertise, reputable experts who have consistently issued correct forecasts are competitively sought after in, for instance, financial markets (see, *e.g.*, Hong, Kubik, and Solomon, 2000; Hong and Kubik, 2003), business consulting (see, *e.g.*, Bourgoin and Harvey, 2018), or entrepreneurship (see, *e.g.*, Baddeley, 2018) in practice.

Importantly, because the informed expert's reputation increases over time, Proposition 1.7 below shows that the crisis of expertise disappears in the long run: principals act on the basis of expert advice *irrespective of* the quality of public information, even when they are skeptical of the expert's ability.

Proposition 1.7. *In the equilibrium, for every* $\varepsilon > 0$*, there exists* $T \geq 0$ *such that*

$$\mathbf{P}[q_t \notin C(p_t), t \geq T | \theta = \theta_I] > 1 - \varepsilon. \tag{1.8}$$

Once the informed type's reputation becomes sufficiently large, and hence all her future reputations are sufficiently large by Proposition 1.6, the public beliefs permanently exit the crisis region by by Proposition 1.4.

More generally, the constrained inefficiency in the equilibrium dissipates over time. If the expert is informed, then once the public beliefs permanently exit the crisis region, each arriving principal efficiently matches his action with the expert's correct report by Proposition 1.4.[9] On the other hand, if the expert is uninformed, then her reputation increases over time until she sends an incorrect report. As her reputation increases, she agrees and thus induces agreeing actions with higher probabilities. Once an incorrect report reveals her type, future principals only choose agreeing actions.

The results speak to the recent concerns in the retail financial advice industry. Take for example the case of Patricia Russell, who appeared in 2019 on *LinkedIn* as a certified financial planner and a graduate from prestigious universities. Her financial advice was quoted in major media outlets, allowing Patricia to benefit from an undeserving reputation as a financial expert before she was discovered to be fake.[10] The discovery of Russell and other fake financial advisors attracted considerable concerns that the presence of these fake experts diminishes the value of informed financial advisors and their genuine advice.[11] The present results suggest that although gambling by the uninformed type undermines the value of the informed expert's correct advice, the damage it causes on the informed expert's payoff dissipates in the long run.

Finally, in the equilibrium, the principals also learn about the uninformed expert over time.

Proposition 1.8. *In the equilibrium, on path, the reputation path* $(p_t)_{t\geq 0}$ *is increasing until it jumps down to zero and stays at zero thereafter, conditional on the expert being*

9. In particular, if the informed expert is arbitrarily patient, the informed type's equilibrium payoff tends to the equilibrium payoff that she would have obtained if her type is commonly known at the outset.

10. See "How This Fake Financial Expert Tricked Outlets Into Publishing Her Advice", *Huffpost*, August 1, 2019.

11. See "Fake 'Expert' Diminishes the Value of Genuine Financial Help", *The Seattle Times*, August 24, 2019.

uninformed. Moreover, for every $\varepsilon > 0$, there exists $T > 0$ such that

$$\mathbf{P}[p_t > 0, t \geq T | \theta = \theta_U] < \varepsilon.$$

Thus, the equilibrium outcome features a "fly-by-night" nature that many uninformed experts exhibit in practice, as in the case of Patricia Russell. In each arrival, the uninformed type is plainly betting on either the event that agreeing is correct or the event that disagreeing is correct. By luck, the uninformed type might send a number of correct reports, building an (undeserving) reputation in the short run. Given a large enough time, there are likely to be a large number of arrivals, and the uninformed type is likely to have sent an incorrect report and to have revealed her type. An example that illustrates such short-term reputation building is the case of Raoul Pal. While Pal successfully predicted the 2008 financial crisis, he also predicted that in 2012, the "biggest banking crisis in world history" and a series of sovereign defaults would take place.[12] These events, of course, did not happen. The financial advisory firm *Vestory* explicitly uses Pal as an example and writes: "Beware of the financial media spouting predictions. More than half the time, they are wrong, and when they're right, it's more likely only because they got lucky."[13]

1.4.2 The Implications of Public Information

Next, I turn to examine the equilibrium implications of public information. Is the crisis due to the abundance of public information? And if so, is it simply a phenomenon where decision makers substitute high-quality public information for expert advice? Briefly, the results in this section speak to the two faces of public information technologies like social media for decision-making: while the technologies allow decision

12. See, *e.g.*, "Former Hedge Funder Presents A Terrifying Vision Of THE END GAME," *Business Insider*, June 2, 2012.

13. See "Fake Financial News," *Vestory*, April 27, 2020.

makers to have easier access to resources that improve their decision-making, they also allow uninformed experts to spring to life at the expense of the decision makers.

On the one hand, public information is beneficial because it facilitates the principals' learning about the state and the expert's type, as shown in the preceding analysis. Their learning about the state improves their decision-making by increasing their reservation payoffs if they were to choose actions independently of expert advice. Moreover, the principals' learning about the expert's type not only improves the quality of their decisions concerning whether to match their actions with the expert's reports, but also improves the quality of expert advice by mitigating the uninformed type's gambling.

On the other hand, in the equilibrium, public information about the state also facilitates the uninformed type's learning and in turn harms the principals. Specifically, in each arrival, the uninformed expert learns how to better pool with the informed type. As a result, in view of Proposition 1.3, she is more likely to agree than to disagree in each arrival. In turn, a "contrarian-rewarding convention" endogenously arises, in the sense that a correctly disagreeing report gives both types of experts a higher continuation payoff than a correctly agreeing report: in each arrival with beliefs (p, μ), where $0 < p < 1$, for each type θ,

$$V^\theta\left(\frac{p}{p+(1-p)\bar{\alpha}_{p,\mu}}; \bar{\alpha}\right) < V^\theta\left(\frac{p}{p+(1-p)(1-\bar{\alpha}_{p,\mu})}; \bar{\alpha}\right).$$

In the face of such convention, the uninformed type faces incentives to gamble for her reputation at the expense of the principals. In particular, she uses the public information about the state to learn how to gamble.

In reality, conventions that reward successful contrarians are common and often appear in the form of implicit incentives. For example, the careers of financial analysts and entrepreneurs typically prosper upon accurate contrarian forecasts (see, *e.g.*, Baddeley, 2018; Bozanic, Chen, and Jung, 2019). My results rationalize such conventions.

My results also shed light on recent empirical evidence on experts' gambling patterns in the face of these conventions in the financial industries. For example, Bozanic et al. (2019) find that financial analysts who work at lower-tier brokerage houses are more likely to make calls contrary to the public beliefs for their career advancement.

1.4.3 The Case for Public Information

In recent decades, despite the concerns that the abundance of public information may have contributed to the crisis, political campaigns and consumer groups have consistently advocated the need to equip client advisees with higher-quality public information so that the clients can substitute public information for low-quality expert advice (see, *e.g.*, Howells, 2005). My results support these advocates to improve the quality of public information but for a different reason. My results suggest that higher quality of public information undermines gambling incentives by uninformed experts, and thus complementarily improves the quality of expert advice.[14]

The preceding discussions thus suggest that policies such that frequent auditing of experts, motivating the production of public information, and higher standards of expert certifications can be effective measures for alleviating the crisis of expertise and improving the quality of decision-making.

Consider first frequent auditing. Specifically, consider increasing the principals' arrival intensity from unity to some larger number. The interpretation is that some of these arriving principals are identical to those in the baseline model who care about matching their actions with the states, while the other arriving principals are

14. Interestingly, this observation resonates with the recent findings in the survey "Trust and Mistrust in Americans's Views of Scientific Experts" by the *Pew Research Center*. The survey documents that confidence in scientific experts is higher among survey participants who possess higher science knowledge (Funk, Hefferon, Kennedy, and Johnson, 2019).

inspectors who audit the expert. Given a high enough arrival rate, principals on average are likely to arrive with sufficiently high state beliefs, alleviating the crisis.

Similarly, policies that motivate frequent production of public information about the state can alleviate the crisis. Specifically, suppose that public news arrives as an independent Poisson process $(Z_t)_{t\geq 0}$ such that upon each news arrival, the current state is publicly revealed. In this setting, the agreeing function $\bar{\alpha}$ continues to characterize an equilibrium. This is because news arrival is exogenous and does not affect the expert's incentives in each arrival, conditional on the public beliefs. As in the case of frequent auditing, if news arrives sufficiently frequently, then principals on average arrive with sufficiently high state beliefs, alleviating the crisis.

Finally, higher certification standards that sufficiently improve the prior reputation can help to preempt the crisis. Indeed, as Proposition 1.4 highlights, if the prior belief p_0 is sufficiently high, and hence all reputations that arise on path are sufficiently high, then the crisis happens with probability zero.

My next result suggests that these policies to alleviate the crisis of expertise are "most needed" in environments where the state is neither too stable nor too volatile. Specifically, I show that the probability that some principal arrives in a crisis in the equilibrium, namely the probability that some principal arrives with beliefs in the crisis region conditional on the expert being informed, vanishes when the state volatility is too small or too large.

Given any history of play at which an arrival ends with resulting reputation p, the probability that some future principal arrives in a crisis is a function of p, denoted by $\lambda(p)$, and is recursively determined by:

$$\lambda(p) = \int_0^\infty \mathrm{e}^{-l} \Bigg[1 - \mathbf{1}_{\mu_l \notin C(p)} \Bigg(1 - \mu_l \lambda\Bigg(\frac{p}{p + (1-p)\bar{\alpha}_{p,\mu_l}}\Bigg) - (1-\mu_l)\lambda\Bigg(\frac{p}{p + (1-p)(1-\bar{\alpha}_{p,\mu_l})}\Bigg)\Bigg)\Bigg] \, \mathrm{d}l.$$

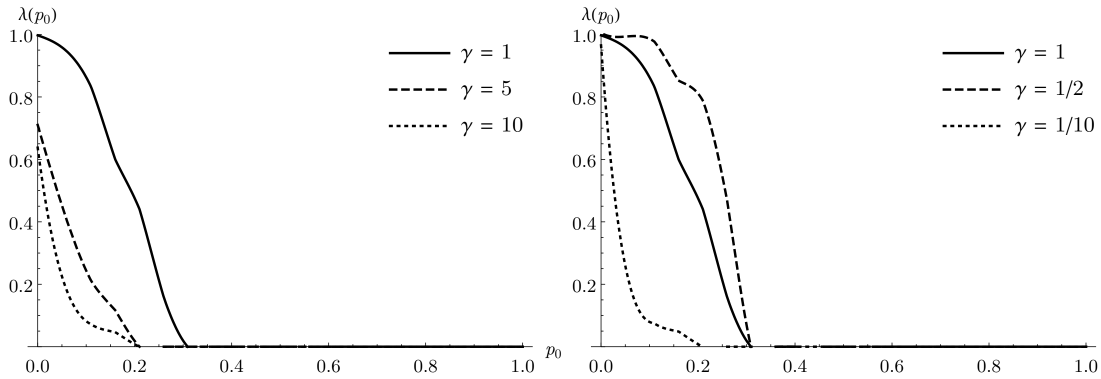

(a) Higher volatility **(b)** Lower volatility

Figure 1.4: Probability that some principal arrives in a crisis ($r = 1$)

The probability that some principal arrives in a crisis in the equilibrium is $\lambda(p_0)$.

Proposition 1.9. *For every $\varepsilon > 0$, there exists $\underline{\gamma}$ and $\bar{\gamma}$, where $0 < \underline{\gamma} < \bar{\gamma} < \infty$, such that for each $\gamma \notin (\underline{\gamma}, \bar{\gamma})$, in the equilibrium,*

$$\lambda(p_0) < \varepsilon. \tag{1.9}$$

In an environment where the state volatility is very low, principals on average arrive when the quality of public information is very high and disciplines the uninformed expert's gambling. In contrast, in an environment where the state volatility is very high, principals on average arrive when the quality of public information is very noisy, making the principals virtually indifferent between the two actions absent expert advice and thereby neglecting the uninformed expert's gambling.

Establishing comparative statics of the equilibrium probability that principals arrive in crises in relation to the volatility γ is less tractable, given the high dimensionality of the program $(*)$. Nonetheless, it is straightforward to obtain them numerically. Figure 1.4 plots the probability against the prior reputation p_0 conditional on the expert being informed, and Figure 1.5 plots the agreeing probability against the state belief, given several values of volatility γ. Figure 1.4 suggests that the probability is non-monotone in γ and is higher at intermediate values of γ, as one may expect from

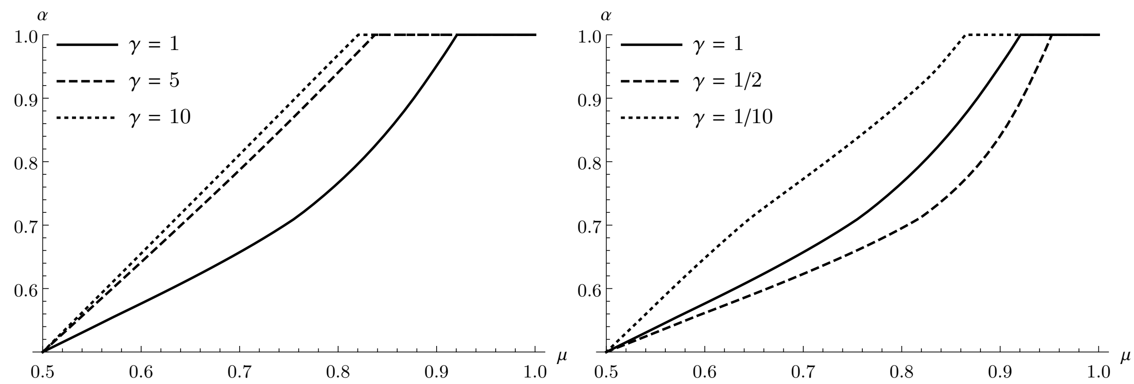

(a) Higher volatility **(b)** Lower volatility

Figure 1.5: Agreeing probabilities $(r = 1, p = 1/5)$

Proposition 1.9. At intermediate values of γ, principals are more likely to arrive with intermediate state beliefs that fall in the crisis region; moreover, the uninformed type has the strongest incentives to gamble at intermediate levels of volatility, as depicted in Figure 1.5.

To understand why the gambling incentives are strongest at intermediate levels of volatility, consider Figure 1.6 which plots the uninformed expert's continuation payoff in the equilibrium. At $\gamma = 1$, the convexity of the payoff at low reputation values gives the uninformed expert strong gambling incentives. In contrast, such convexity at low reputations vanishes for large γ, as the principals are likely to arrive with state beliefs of virtually 1/2 and therefore, according to (1.7), pay the expert wages approximately equal to the expert's reputations. Such convexity also vanishes for small γ, as the principals are likely to arrive with state beliefs of virtually one and pay the expert negligible wages.

1.5 Concluding Remarks

In this paper, I have developed a simple model of expert advice to examine the crisis of expertise. The crisis has attracted considerable attention amid recent technological advancement that makes information abundant and easily accessible. In my model,

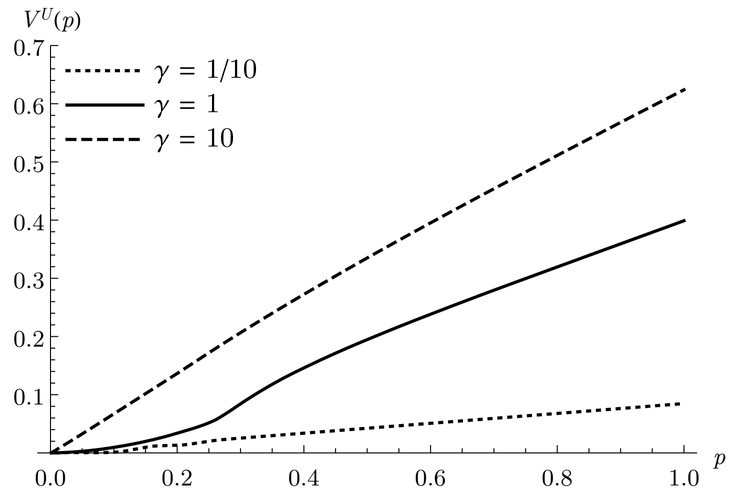

Figure 1.6: Uninformed type's continuation payoff

the crisis is an equilibrium phenomenon where decision makers dismiss the correct advice provided by informed experts and act solely on the basis of their noisy public information. I find that decision makers ignore valuable expert advice only when the quality of the public information is low, as it is then when uninformed experts have stronger incentives to gamble for their reputations at the expense of the decision makers. My analysis suggests that policies such as frequent auditing of experts, motivating the production of high-quality public information, and higher standards of expert certifications can be effective measures for alleviating the crisis of expertise.

I close the paper with a brief discussion of some of the assumptions that I made in my analysis. First, I have imposed Property 1.2 in my solution concept. Property 1.2 ensures that when an expert incorrectly reports the state given a reputation of one, then her reputation jumps down to zero. While the analysis readily extends to allowing for other off-path reputations, the property ensures that the principals' updating of the expert's reputation does not exhibit a discontinuity at reputation one. Indeed, if the uninformed expert reports incorrectly when her reputation is short of unity, then her reputation jumps down to zero.

Property 1.2 also ensures that when a principal arrives with a public state belief of one, then he chooses an agreeing action even if the expert sends a disagreeing report (off path), as the principal then believes that the expert is uninformed for sure. The analysis readily extends to allowing for other off-path beliefs by the principals. Such off-path beliefs would not affect the uninformed expert's incentives, as the event that principals arrive with unit state belief happens with probability zero.

Another assumption that I made is that the public information available to principals and experts consists only of past advice and past action outcomes. In reality, there may be public shocks, *e.g.*, a pandemic, that arrive at random times and produce information about the state. Suppose that high shocks and low shocks arrive as independent Poisson processes $(\bar{Z}_t)_{t\geq 0}$ and $(\underline{Z}_t)_{t\geq 0}$ such that upon each arrival of a high (resp., low) shock, the current state publicly jumps to $\bar{\theta}$ (resp., $\underline{\theta}$). In this setting, like in the public-news setting discussed in Section 1.4.3, the agreeing function $\bar{\alpha}$ continues to characterize an equilibrium.

Finally, I have assumed that the state is binary. The analysis readily extends to more than two states if, for each principal, the benefit that he receives from not matching the state is independent of what non-matching action he chooses, while his benefit from matching the state is higher. If the payoff structure is different, for instance, if the state is a real number and each principal has a quadratic loss function such that he strictly prefers an action that is closer to the state, then the fundamental notion of agreeing and disagreeing, and more importantly the notion of a crisis of expertise, do not carry over. The nature of the problem is then different and is beyond the scope of this paper. I view this as a promising venue for future research.

1.6 Appendix A: Proofs

1.6.1 Proof of Proposition 1.1

Given any equilibrium, let $\mathcal{V}^U(h^U)$ denote the uninformed type's continuation payoff at the end of an arrival when the principal's benefit is realized such that her history is updated to h^U.

I first show that in any equilibrium, given any history h^U, the uninformed type's continuation payoff $V^U(h^U)$ strictly increases in the associated reputation p. Fix an equilibrium profile (σ, φ). Given any play history at which an arrival ends with reputation p and history h^U, consider the uninformed expert's continuation payoff $\mathcal{V}^U(h^U)$. To emphasize the associated reputation $p \equiv p(h^U)$, in the following, I rewrite $\mathcal{V}^U(h^U)$ as $\mathcal{V}^U(h^U; p)$. Let $\sigma^U|_{h^U}$ denote the uninformed type's continuation strategy. The expected wage in the next principal's arrival is

$$\int_0^\infty \mathrm{e}^{-(r+1)l}\, w(p, \mu_l; \alpha)\, \mathrm{d}l,$$

where

$$w(p, \mu_l; \alpha) := \max(p + (1-p)(2\mu - 1)(2\alpha_l - 1) - (2\mu_l - 1), 0)$$

represents the flow wage given that the uninformed expert sends an agreeing message with probability α_l when the next principal arrives given a time-lapse l in the future. Suppose that the reputation p increases to some $p' > p$ and consider the continuation game. By choosing the continuation strategy $\sigma^U|_{h^U}$ (which is feasible), the uninformed expert's continuation payoff is $\hat{\mathcal{V}}^U(h^U; p') > \mathcal{V}^U(h^U; p)$, as $w(\cdot; p', \cdot) > w(\cdot; p, \cdot)$ and the probabilities that her messages match the states in each future arrival remain the same. Since the equilibrium payoff $\mathcal{V}^U(h^U; p')$ in the continuation game must be at least $\hat{\mathcal{V}}^U(h^U; p')$, the claim follows.

Next, I show that in any equilibrium, there exists a strictly increasing function $V^U : [0,1] \to \mathbb{R}_+$ such that given any history h^U at the end of an arrival with associated reputation $p \in [0,1]$, $\mathcal{V}^U(h^U) = V^U(p)$. In equilibrium, the informed expert plays a strategy $\bar{\sigma}^I$ given which she reports truthfully in each arrival. I establish several key properties that the uninformed expert's continuation payoffs at each history must satisfy in any equilibrium. I then use these properties to characterize the uninformed expert's equilibrium strategy.

Fix an uninformed expert's equilibrium strategy σ^U. First, given any two histories $h^U, \tilde{h}^U$ at which an arrival ends with an identical reputation p, $\mathcal{V}^U(h^U) = \mathcal{V}^U(\tilde{h}^U) =: V^U(p)$ for some function V^U, for otherwise the uninformed expert has a profitable deviation at either h^U or $\tilde{h}^U$. Clearly, $V^U(0) = 0$ because all future wages are zero.

Now, in any equilibrium, at any history of play where a principal arrives with public beliefs (p, μ) and the uninformed expert faces a history h^U, the uninformed expert's agreeing probability, denoted by α_{h^U}, must satisfy the incentive constraint

$$\mu V^U\left(\frac{p}{p+(1-p)\alpha_{h^U}}; \bar{\alpha}\right) - (1-\mu)V^U\left(\frac{p}{p+(1-p)(1-\alpha_{h^U})}; \bar{\alpha}\right) \begin{cases} \geq 0, & \text{if } \alpha_{h^U} = 1, \\ = 0, & \text{if } \alpha_{h^U} \in (0,1), \\ \leq 0, & \text{if } \alpha_{h^U} = 0. \end{cases}$$

Since V^U is strictly increasing, if $p > 0$, then α_{h^U} must be such that for some cutoff $\mu^*(p)$,

$$\alpha_{h^U} = \begin{cases} 1, & \text{if } \mu \geq \mu^*(p), \\ \bar{\alpha}_{p,\mu}, & \text{otherwise,} \end{cases}$$

where $\alpha^*_{p,\mu}$ uniquely solves

$$\mu V^U\left(\frac{p}{p+(1-p)\bar{\alpha}_{p,\mu}};\bar{\alpha}\right) = (1-\mu)V^U\left(\frac{p}{p+(1-p)(1-\bar{\alpha}_{p,\mu})};\bar{\alpha}\right).$$

Because V^U is strictly increasing and $\mu > \frac{1}{2}$, $\alpha^*_{p,\mu} \in (\frac{1}{2}, 1]$. In contrast, if $p = 0$, then any α_{h^U} is optimal because $V^U(0;\bar{\alpha}) = 0$ independently of the uninformed type's agreeing probability.

Observe that the uninformed expert's equilibrium strategy in the arrival must therefore be Markov in the public beliefs (p, μ). Moreover, it is unique if $p > 0$. Finally, because $V^U(\cdot;\bar{\alpha})$ is strictly increasing, the agreeing probability increases in p by construction.

1.6.2 Proof of Lemma 1.1

Let $L \equiv L^\infty([0,1] \times (\frac{1}{2}, 1]; [0,1])$ be the set of Lebesgue measurable functions that map from $[0,1] \times (\frac{1}{2}, 1]$ to $[0,1]$. Since each agreeing function $\alpha \in L$ characterizes a candidate equilibrium, it induces a set $P(\alpha) \subseteq [0,1]$ of reputations that arise at the end of each arrival on path, as well as a function $V^U(\cdot;\alpha) : [0,1] \to \mathbb{R}_+$ such that $V^U(p;\alpha)$ represents the uninformed type's continuation payoff at the end of each arrival that results in reputation p after the principal's benefit is realized.

The uninformed expert's maximum equilibrium payoff is given by:

$$\sup_{\alpha \in L} V^U(p_0;\alpha). \tag{$*$}$$

Clearly, the objective in the problem $(*)$ is continuous in α when L is endowed with the weak-* topology. Moreover, L is weak-* compact, as $L^\infty([0,1] \times (\frac{1}{2}, 1]; \mathbb{R})$ is the dual of $L^1([0,1] \times (\frac{1}{2}, 1]; \mathbb{R})$ and hence the unit ball $L = \{\alpha \in L^\infty([0,1] \times (\frac{1}{2}, 1]; \mathbb{R}) : \|\alpha\|_\infty \leq 1\}$ is weak-* compact by Alaoglu's theorem (*e.g.*, Aliprantis and Border, 2006, Theorem 6.21). Thus, a solution to the program exists and the supremum is attained as a maximum. Moreover, the function $V^U(p, s; \alpha)$, holding fixed the agreeing probability $\alpha_{p,\mu,s}$ induced by the agreeing function as constant for each (p, μ, s), is continuous in $p \in [0, 1]$. For each $(p, \mu) \in (0, 1] \times (\frac{1}{2}, 1]$, the equilibrium agreeing probability $\bar{\alpha}_{p,\mu} \in [0, 1]$ is characterized by (1.2). Such $\bar{\alpha}_{p,\mu} \in [0, 1]$ exists, because $V^U(p; \bar{\alpha})$ is continuous in p and is strictly increasing.

1.6.3 Proof of Proposition 1.2

The result is clear from the discussion in the main text.

1.6.4 Proof of Corollary 1.2

It is straightforward to verify that

$$w(p, \mu; \alpha) = (p + (1 - p)(2\mu - 1)(2\bar{\alpha}_{p,\mu} - 1) - (2\mu - 1))^+ > 0 \quad \iff \quad \mu \notin C(p).$$

1.6.5 Proof of Proposition 1.3

I have shown that the agreeing probability exceeds 1/2 given any beliefs (p, μ), as well as part 1 in the proof of Proposition 1.1. It remains to show part 2. The cutoff is

$$\mu^*(p) = \min\{\mu \in [\tfrac{1}{2}, 1] : \mu V^U(p; \bar{\alpha}) \geq (1 - \mu)V^U(1; \bar{\alpha})\}.$$

Since V^U is strictly increasing, μ^* is strictly decreasing. Since $V^U(0) = 0$, $\mu^*(0) = 1$.

1.6.6 Proof of Proposition 1.4

That $1 \in C(p)$ for any reputation $p \in [0, 1]$ follows by definition of C. To show part 1 and part 2, note that by definition of $\bar{\alpha}$ and κ, there exists $\mu^* < 1$ such that for every $\mu \in [\mu^*, 1]$, $\bar{\alpha}_{p,\mu} = 1$ but $\kappa_{p,\mu} < 1$. Thus, the set $C(p)$ excludes high enough interior state beliefs μ. On the other hand, for small enough state beliefs μ, $\kappa_{p,\mu} < \frac{1}{2}$. Because $\bar{\alpha}_{p,\mu} > \frac{1}{2}$ for each $\mu \in (\frac{1}{2}, 1]$, the set $C(p)$ also excludes low enough state beliefs μ.

Finally, the monotonicity of the crisis correspondence C follows directly from the fact that $\bar{\alpha}_{p,\mu}$ increases in p and $\kappa_{p,\mu}$ strictly decreases in p for each $\mu \in (\frac{1}{2}, 1]$. To see that $\bar{p} < 1$, observe that given any p, $\kappa_{p,\mu}$ strictly increases in μ, and $\kappa_{p,\mu} \leq \frac{1}{2}$ for $\mu \leq \frac{1}{2}(1+p)$; moreover, $\alpha_{p,\mu} > \frac{1}{2}$ given any beliefs (p, μ). Given a large enough p, the gambling cutoff $\mu^*(p)$ is close to $\frac{1}{2}$ such that $\kappa_{p,\mu} = \frac{1}{2}$ at some $\mu > \mu^*(p)$ and $\bar{\alpha}_{p,\mu} = 1$ for every $\mu \geq \mu^*(p)$.

1.6.7 Proof of Proposition 1.5

In equilibrium, given that the uninformed expert's strategy is Markov in the public beliefs in each arrival, the informed expert's continuation payoff at the end of an arrival resulting in reputation p is characterized by the recursive equation

$$V^I(p;\bar{\alpha}) = \int_0^\infty \mathrm{e}^{-(r+1)l} \Bigg\{ rw(\bar{\alpha}_{p,\mu_l}; p, \mu_l) + \mu_l V^I\left(\frac{p}{p+(1-p)\bar{\alpha}_{p,\mu_l}}; \bar{\alpha}\right) \\ + (1-\mu_l)(1-\bar{\alpha}_{p,\mu_l}) V^I\left(\frac{p}{p+(1-p)(1-\bar{\alpha}_{p,\mu_l})}; \bar{\alpha}\right)\Bigg\} \,\mathrm{d}l. \tag{1.10}$$

Clearly, by definition, $V^I(0;\bar{\alpha}) = 0$. I argue that V^I is strictly increasing. It is useful to observe that by construction of $\bar{\alpha}_{p,\mu_l}$ in the proof of Proposition 1.1, the two induced beliefs when the state is revealed at the end of the arrival, namely

$$\frac{p}{p+(1-p)\bar{\alpha}_{p,\mu_l}}, \quad \frac{p}{p+(1-p)(1-\bar{\alpha}_{p,\mu_l})},$$

are increasing in p.

Let $B([0,1])$ be the set of bounded functions $f : [0,1] \to \mathbb{R}_+$ equipped with the sup norm $\|\cdot\|$. Consider the operator T^I given by

$$(T^I f)(p) = \int_0^\infty \mathrm{e}^{-(r+1)l} \left\{ rw(\bar{\alpha}_{p,\mu_l}; p, \mu_l) + \mu_l f\left(\frac{p}{p + (1-p)\bar{\alpha}_{p,\mu_l}}\right) \right. \\ \left. + (1-\mu_l)(1-\bar{\alpha}_{p,\mu_l}) f\left(\frac{p}{p + (1-p)(1-\bar{\alpha}_{p,\mu_l})}\right) \right\} \mathrm{d}l. \tag{1.11}$$

I show that V^I is well-defined as the unique fixed point of T^I, which is a contraction mapping. I verify Blackwell's sufficient conditions for a contraction, namely monotonicity and discounting. Monotonicity follows immediately since, given f and $\hat{f}$ such that $f \leq \hat{f}$, $T^I f \leq T^I \hat{f}$. Discounting follows since, given any $\phi > 0$, $T^I \circ (\phi + f(p)) < T^I \circ f(p) + \phi$.

Next, let $B'([0,1])$ be the set of bounded and increasing functions on $[0,1]$, and let $B''([0,1])$ be the set of bounded and strictly increasing functions on $[0,1]$. Since $B'([0,1])$ is a closed subset of $B([0,1])$, to prove that V^I is strictly increasing, it suffices to show that $T^I(B'([0,1])) \subseteq B''([0,1])$ by a standard property of contraction mappings (see, *e.g.*, Stokey, Lucas, and Prescott, 1989, p. 52, Corollary 1). The latter follows because f is increasing, $p/(p + (1-p)\bar{\alpha}_{p,l})$ and $p/(p + (1-p)(1-\bar{\alpha}_{p,l}))$ are increasing in p for each μ_l, and

$$\int_0^\infty \mathrm{e}^{-(r+1)l}\, w(\bar{\alpha}_{p,\mu_l}; p, \mu_l)\, \mathrm{d}l$$

is strictly increasing in p.

1.6.8 Proof of Proposition 1.6

Let $\mathbf{P}^I$ denote the probability measure over the set of outcomes of the game induced in the equilibrium by conditioning on an informed type. Conditional on the expert being an informed type, given any play history at which a principal arrives with beliefs

(p, μ), the reputation once the state is revealed is one of

$$\frac{p}{p + (1-p)\bar{\alpha}_{p,\mu}} \geq p, \quad \frac{p}{p + (1-p)(1-\bar{\alpha}_{p,\mu})} > p.$$

Thus $(p_t)_{t \geq 0}$ is increasing. Now, fix an arbitrary play history and a sequence of arrival times $(t_i)_{i=0}^{\infty}$. Given any arrival time t_i, for every $\eta > 0$, there exists another arrival time $t_{i'}$, where $i' - i$ is sufficiently large, such that, $\mathbf{P}^I[p_{t_i} = p_{t_{i'}}] < \eta$. This is because for $p_{t_i} = p_{t_{i'}}$, it must hold that in all arrivals in between t_i and $t_{i'}$, the agreeing states must be correct and the time-lapses are sufficiently small such that the uninformed type agrees with probability one. This event happens with negligible probability if $i' - i$ is large. Thus, for every $\varepsilon > 0$, there exist $T > 0$ and $\bar{n} > 0$ such that for every $t \geq \underline{T}$ and for every $n \geq \bar{n}$,

$$\begin{aligned} \mathbf{P}[N_t < n] &< \frac{\varepsilon}{2}, \\ \mathbf{P}^I[p_t = p_{t+T}, T \geq \underline{T}, N_t \geq n] &< \frac{\varepsilon}{2}. \end{aligned}$$

Then

$$\begin{aligned} \mathbf{P}^I[p_t = p_{t+T}, T \geq \underline{T}] &= \mathbf{P}^I[p_t = p_{t+T}, T \geq \underline{T}, N_t < n] + \mathbf{P}^I[p_t = p_{t+T}, T \geq \underline{T}, N_t \geq n] \\ &< \mathbf{P}[N_t < n] + \mathbf{P}^I[p_t = p_{t+T}, T \geq \underline{T}, N_t \geq n] \\ &< \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \varepsilon. \end{aligned}$$

Thus

$$\mathbf{P}^I[p_t < p_{t+T}, T \geq \underline{T}] = 1 - \mathbf{P}^I[p_t = p_{t+T}, T \geq \underline{T}] \geq 1 - \varepsilon.$$

1.6.9 Proof of Proposition 1.7

Let $\mathbf{P}^I$ denote the probability measure over the set of outcomes of the game induced in the equilibrium by conditioning on an informed type. Fix $\bar{p}$ such that for every $p \geq \bar{p}$, $C(p) = \{1\}$. By Proposition 1.7, for every $\varepsilon > 0$, there exist $T > 0$ and $\bar{n} > 0$ such that for every $t \geq T$ and for every $n \geq \bar{n}$,

$$\mathbf{P}[N_t < n] < \frac{\varepsilon}{2},$$
$$\mathbf{P}^I[p_t < \bar{p}, t \geq T, N_t \geq n] < \frac{\varepsilon}{2}.$$

Then

$$\begin{aligned}\mathbf{P}^I[q_t \in C(p_t), t \geq T] &< \mathbf{P}^I[p_t < \bar{p}, t \geq T] \\ &= \mathbf{P}^I[p_t < \bar{p}, t \geq T, N_t < n] + \mathbf{P}^I[p_t < \bar{p}, t \geq T, N_t \geq n] \\ &< \mathbf{P}^I[N_t < n] + \mathbf{P}^I[p_t < \bar{p}, t \geq T, N_t \geq n] \\ &= \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \varepsilon.\end{aligned}$$

1.6.10 Proof of Proposition 1.8

Let $\mathbf{P}^U$ denote the probability measure over the set of outcomes of the game induced in the equilibrium by conditioning on an uninformed type. Because $(N_t)_{t \geq 0}$ is a Poisson process with unit intensity and because the uninformed expert sends an incorrect message in each arrival with positive probability (except in the zero-probability event that the state belief is one in the arrival), given $\varepsilon > 0$, there exist $T > 0$ and $\bar{n} > 0$ such that for every $t \geq T$ and for every $n \geq \bar{n}$,

$$\mathbf{P}[N_t < n] < \frac{\varepsilon}{2},$$
$$\mathbf{P}^U[p_t > 0, t \geq T, N_t \geq n] < \frac{\varepsilon}{2}.$$

Then

$$\begin{aligned}\mathbf{P}^U[p_t > 0, t \geq T] =& \mathbf{P}^U[p_t > 0, t \geq T, N_t < n] + \mathbf{P}^U[p_t > 0, t \geq T, N_t \geq n] \\ <& \mathbf{P}^U[N_t < n] + +\mathbf{P}^U[p_t > 0, t \geq T, N_t \geq n] \\ <& \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \varepsilon.\end{aligned}$$

1.6.11 Proof of Proposition 1.9

Fix $\varepsilon > 0$. Let $\mathbf{P}^I$ denote the probability measure over the set of outcomes of the game induced in the equilibrium by conditioning on an informed type. By definition, $\lambda(p_0) < \mathbf{P}^I[q_t \in C(p_t),\ t \geq 0]$. Thus, it suffices to show that there exists $\underline{\gamma} \equiv \bar{\gamma}(p_0, r, \varepsilon), \bar{\gamma} \equiv \bar{\gamma}(p_0, r, \varepsilon)$, where $0 < \underline{\gamma} < \bar{\gamma} < \infty$ such that for each $\gamma \notin (\underline{\gamma}, \bar{\gamma})$, in the equilibrium,

$$\mathbf{P}^I[q_t \in C(p_t),\ t \geq 0] < \varepsilon.$$

By Proposition 1.7, there exists $T \geq 0$ such that

$$\mathbf{P}^I[q_t \in C(p_t), t \geq T] < \frac{\varepsilon}{2}. \tag{1.12}$$

Fix such T. By Proposition 1.4 and that $p_t \geq p_0$ for any $t \geq 0$ conditional on the expert being informed,

$$\mathbf{P}^I[q_t \in C(p_t), t \in [0, T)] \leq \mathbf{P}^I[q_t \in C(p_0), t \in [0, T)].$$

Next, let $\underline{l}(\gamma) := \min\{l > 0 : \mu_l \in C(p_0)\}$ and $\bar{l}(\gamma) := \max\{l > 0 : \mu_l \in C(p_0)\}$. By definition of C in Proposition 1.4, $\underline{l}(\gamma), \bar{l}(\gamma)$ exists and are well-defined. Moreover,

$$\underline{l}(\gamma), \bar{l}(\gamma) \to \begin{cases} \infty, & \text{as } \gamma \to 0, \\ 0, & \text{as } \gamma \to \infty. \end{cases}$$

Then,

$$\mathbf{P}^I[\mu_l \in C(p_0), l \geq 0] \leq \int_{\underline{l}(\gamma)}^{\bar{l}(\gamma)} \mathrm{e}^{-l}\,\mathrm{d}l \to \begin{cases} 0, & \text{as } \gamma \to 0, \\ 0, & \text{as } \gamma \to \infty. \end{cases}.$$

Thus, there exists $\underline{\gamma}$, $\bar{\gamma}$ such that for every $\gamma \notin (\underline{\gamma}, \bar{\gamma})$, in the equilibrium,

$$\mathbf{P}^I[\mu_l \in C(p_0), l \geq 0] < \frac{\varepsilon}{2T}. \tag{1.13}$$

Now,

$$\begin{aligned} \mathbf{P}^I[q_t \in C(p_t), t \in [0,T]] &= \sum_{n=0}^{\infty} \mathbf{P}[N_T = n]\mathbf{P}^I[q_t \in C(p_t), t \geq [0,T] \mid N_T = n] \\ &\leq \sum_{n=0}^{\infty} \mathbf{P}[N_T = n] n \mathbf{P}^I[\mu_l \in C(p_0), l \geq 0] \\ &\leq \frac{\varepsilon}{2T} \sum_{n=0}^{\infty} \mathbf{P}[N_T = n] n \\ &= \frac{\varepsilon}{2}, \end{aligned} \tag{1.14}$$

where the second last line follows from (1.13), and the last line follows from the property of Poisson process (with unit intensity) that the expected number of arrivals up to time T is plainly T. Thus,

$$\begin{aligned}\mathbf{P}^I[q_t \in C(p_t), t \geq 0] &= \mathbf{P}^I[q_t \in C(p_t), t \in [0, T)] + \mathbf{P}^I[q_t \in C(p_t), t \geq T] \\ &\leq \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \varepsilon.\end{aligned}$$

where the last line follows from (1.12) and (1.14).

Chapter 2

Firm Certification

2.1 Introduction

A group of tourists contemplates whether to stay at a particular hotel. They are unsure whether the hotel is capable of providing high quality. Moreover, even if the hotel is capable of providing high quality, they are unsure whether it will exert costly effort to do so. The resulting adverse selection and moral hazard limit the tourists' willingness to pay the hotel. The hotel, in response, pays *Intertek* to certify its quality provisions. *Intertek* monitors the hotel's quality provisions over time, and periodically updates the hotel's certification status. At any time, the hotel is either certified or not certified, and might also receive "the Best Overall Hotel Award" from *Intertek* if it has consistently delivered high quality.[1]

Third-party certification of firms, like in the *Intertek* example, is pervasive in many markets. The example makes clear that a key role of firm certification is to address adverse selection and dynamic moral hazard problems. For instance, when explaining why firms should consider signing up for its certification service, *Intertek* writes: "We can strengthen your brand by ensuring consistency of quality, standards and risk

1. See, for instance, "Intertek Cristal Announces Winners of 2020 Cristal Global Awards for Outstanding Performance in the Travel & Tourism Industry," *Intertek*, November 09, 2020.

management in all things you do..."[2] Similarly, the ISO 9001 certification website writes: "Using ISO 9001 helps ensure that customers get consistent, good-quality products and services, which in turn brings many business benefits."[3]

"Honors" certification schemes as in the *Intertek* example, which either certify or do not certify participating firms, and occasionally provide public awards to highlight the top-performing firms, are common. It is also common that certification intermediaries award the honors implicitly by promoting the top-performing firms on their websites or press releases, instead of providing explicit awards. The honors certification structure, however, largely censors participating firms' evolving track records, and may therefore destroy the firms' incentives to exert costly effort to provide high quality. It thus appears to conflict with the goal to address dynamic moral hazard.

At the same time, many certification schemes allegedly feature low certification standards. While a certified firm is generally perceived as a firm that can be relied upon to provide high quality, the report "The False Promise of Certification" by the NGO *Changing Markets Foundation* (Brad, Delemare, Hurley, Lenikus, Mulrenan, Nemes, Trunk, and Urbancic, 2018) argues otherwise by documenting that many firms get certified simply by participating in a scheme. Low standards limit the schemes' effectiveness to inform consumers about firms' ability to provide high quality. In turn, low standards might disrupt participating firms' incentives to exert costly effort to provide high quality, limiting the role of certification to address dynamic moral hazard. Low standards also appear to be be at odds with the role of certification to address adverse selection.

If the two common properties of firm certification, namely an honors certification structure and a low certification standard, limit the value of certification in alleviating

2. See `https://cristalstandards.com/` (accessed on February 16, 2021).

3. See `https://www.iso.org/iso-9001-quality-management.html` (accessed on February 16, 2021).

information asymmetries and in turn the certifiers' revenue, why are they common?

To address this question, I develop a certification model featuring both adverse selection and dynamic moral hazard in a firm's quality provision. A firm (she) trades with a succession of consumers over time. The firm privately knows her type (either competent or inept) and her effort towards quality provision, and is tempted to shirk against consumers. The firm may pay a certifier upfront to participate in a certification scheme. The scheme repeatedly rates the firm given her track record. Ratings link the firm's past behavior to market expectations of her future behavior, affecting her future revenue. The certifier chooses a scheme *ex ante* to maximize his revenue.

I characterize schemes that allow the certifier to extract the full net market surplus, which I call full-surplus-optimal (FSO) schemes. These schemes must attract participation by both competent and inept firms. They must display "coarse" ratings that censor a firm's track record, as well as "transfer" ratings that transfer some positive surpluses produced by a competent firm to an inept firm. I show that FSO schemes can be implemented with an honors certification structure and a lower certification standard: in each period, a participating firm receives one of the three ratings, namely "not certified," "certified," or "certified with honors," depending on the firm's track record; moreover, the standard is low in the sense an inept firm is frequently certified. The "certified" rating is a transfer rating.

Low-standard, honors certification schemes convey a limited amount of information regarding a firm's type and evolving track record to consumers. The competent firm then faces persistent reputational incentives to exert high effort in order to separate from an inept type, generating the maximum surplus in each period. As a result, the scheme maximizes social welfare. Moreover, by certifying the inept firm, the certifier transfers some surpluses that are produced by the competent firm from the competent firm to the inept firm. The certifier then extracts the resulting surpluses from both types, which sum to the full net market surplus.

My construction of FSO schemes contains novel elements. In my model, the folk theorem does not hold. Thus, schemes such as those based on review strategies (*e.g.*, Radner, 1985; Varas, Marinovic, and Skrzypacz, forthcoming) are in general not FSO even in the limit of no discounting. Under a review strategy, the certifier divides the horizon into review phases and rates the firm at the end of each phase given the realized quality provisions. A short review phase cannot motivate high effort by a competent firm when the effort cost is high, as it fails to sufficiently separate a competent firm from an inept one. A longer review phase inevitably leads to shirking, since a firm that has produced several bad outcomes expects to fail the review and thus shirks for the remainder of the phase. The punishment phase upon failing a review might also call for further shirking. The probability of shirking does not vanish for large discount factors.

The features of low certification standard and the censoring of track records have attracted ample concerns about the ability of certification schemes to alleviate information asymmetries in experience-good and credence-good markets,[4] in environmental certifications,[5] and in financial markets.[6] These features have also attracted considerable attention in the academic literature of certification. Existing certification models (*e.g.*, Lizzeri, 1999; Stahl and Strausz, 2017; Harbaugh and Rasmusen, 2018; Hopenhayn and Saeedi, 2019) are often cast in a static, pure adverse-selection setting,[7] focusing on certifying the quality of finished products and abstracting from

4. For example, see "Grading New York Restaurants: What's in an 'A'?" *The New York Times*, January 19, 2011.

5. For instance, *Changing Market Foundation* calls on schemes in different industries to be more transparent and more selective about who to certify (Brad et al., 2018).

6. Many regulators criticize that corporate credit ratings are too coarse to accurately inform the public about the firms' ability to meet their debt obligations (see, *e.g.*, Partnoy, 2006; Pagano and Volpin, 2010; Partnoy, 2017). In turn, they might fail to discipline firms to avoid excessive risk-taking and to manage their investments diligently.

7. Some models have an *ex ante* moral-hazard stage where the product quality is chosen (*e.g.*, Albano and Lizzeri, 2001; Zapechelnyuk, 2020) but abstract from dynamic moral hazard considerations.

the role of certification in incentivizing dynamic quality provisions. This role is key to certifying *firms*, as opposed to finished products, in applications: quality certifications incentivize hotels, restaurants, and other experience-good firms to put in costly effort to best serve consumers; eco-labels incentivize firms to continually invest in new environmental technologies (*e.g.*, OECD, 1997; Goodland, 2002); corporate credit ratings discipline firms to continually manage their investments diligently.

To the best of my knowledge, my model provides the first analysis of certification that features both adverse selection and dynamic moral hazard, and delivers new insights concerning the popularity of low certification standards and honors certification structures. The important insight from existing pure adverse-selection models is that, these features maximize participation by firms and thus the certifiers' revenue.[8] But as the opening paragraphs motivate, when the dynamic moral hazard problem is present, these properties might backfire, disrupting firms' incentives to provide high quality and in turn destroying the certifiers' revenue. My analysis thus elucidates how a profit-maximizing certifier trades off screening and incentives to provide high quality.[9] Moreover, contrary to the pure adverse-selection models, simply attracting participation by inept firms does not suffice for the certifier's profit maximization. Specifically, in my model, the certifier must employ transfer ratings upon attracting participation by inept firms *and* induce high effort by competent firms at these ratings to maximize his profit.

My results also contribute to the literature on repeated games. My model builds on a classic firm-reputation game; see, *e.g.*, Mailath and Samuelson (2001), Hörner

8. In particular, firms participate in certification to signal their quality. In this regard, the literature that studies endogeneous entry into markets by firms who face reputational incentives is also related (see, *e.g.*, Atkeson, Hellwig, and Ordoñez, 2015; Vellodi, 2019).

9. To focus on how low-standard, honors certification addresses adverse selection and dynamic moral hazard, my analysis abstracts from other important reasons for censoring information, such as biased certifiers or difficulty to interpret exact information by consumers, that are examined in the certification literature. See also Dranove and Jin (2010) for a survey.

(2002), and also Bar-Isaac and Tadelis (2008) for a survey. As discussed above, the folk theorem does not hold, and schemes with familiar structures such as those based on review strategies are in general not FSO. Much difficulty in constructing FSO schemes in my model stems from the absence of a Stackelberg type of firm who always exerts high effort. In FSO schemes, "top" ratings must give a firm high profits to motivate her to exert high effort at lower ratings. When a Stackelberg type is present, top ratings give a firm high profits by allowing her to shirk while consumers' payments remain high (*e.g.*, Ekmekci, 2011). In my model, the absence of a Stackelberg type rules out such possibility and yields a new trade-off between rating transparency and efficiency. Top ratings induce high payments in equilibrium only if, at these ratings, consumers believe that the firm is likely competent and the firm exerts high effort. The FSO schemes must therefore be sufficiently transparent for consumers to learn. However, they cannot be too transparent: to motivate high effort at top ratings, a firm that performs badly must be punished with low ratings that induce low payments. Given consistent high effort, low payments can be induced only if consumers are unsure of the firms type.

More broadly, the complementarity between adverse selection and moral hazard has been highlighted extensively in the literatures of information design and limited records for dynamic incentive provision (*e.g.*, Lizzeri, Meyer, and Persico, 2002; Liu, 2011; Ekmekci, 2011; Liu and Skrzypacz, 2014; Che and Hörner, 2018; Kovbasyuk and Spagnolo, 2018; Bhaskar and Thomas, 2019; Hörner and Lambert, 2020; Lorecchio and Monte, 2021).[10] The goal to rationalize low-standard honors scheme, as well as the certifier's profit-maximizing motive, contrast with these literatures. My analysis also highlights the role of transfer ratings in maximizing the certifier's profits. In particular, as illustrated in Section 2.5, these transfer ratings constrain the certifier's

10. Relatedly, Marinovic, Skrzypacz, and Varas (2018) take a certification technology as exogenously given and study a firm's investment behavior.

ability to provide incentives for quality provision.

2.2 Model

2.2.1 The Market

Time is discrete and the horizon is infinite. A long-lived firm (she) faces a sequence of short-lived consumers (it). A certifier (he) designs and sells a certification scheme.

The horizon is divided into two stages. In the first stage, the certifier offers the firm a (certification) scheme, defined below in Definition 2.1.[11] The firm privately learns her type θ, which is competent (C) with probability $\mu \in (0, 1)$ and is inept (I) otherwise.[12] The firm observes the scheme, and either accepts or rejects the scheme in that stage. Upon acceptance, the firm reports a type to the certifier and pays a certification fee, which may depend on her report as specified by the scheme. Upon rejection, without loss of generality,[13] the firm receives a null rating $\varnothing$ forever. The firm can mix over her participation decisions $d \in \{C, I, \varnothing\}$, where $d = \theta$ refers to acceptance with a report θ, and $d = \varnothing$ stands for rejection (identified with the null rating). Let $\pi(d|\theta)$ denote the probability that a firm chooses decision d given her type θ. This ends the first stage. Observe that, for ease of exposition, I have modeled the scheme as a direct mechanism. The certifier can equivalently offer the firm a menu of schemes, each intended for a type, and charge the firm a fee tailored for each

11. Because the scheme is chosen once and for all, an interpretation is that the certifier commits to the scheme. Commitment may arise from his reputational concerns (Coffee, 1997). A famous example of a certifier's desire to signal such commitment is the expulsion by Better Business Bureaus of its Los Angeles affiliate in 2009, after several eateries ware discovered to pay for high ratings. See "Better Business Bureau Expels Los Angeles Area Chapter," *Los Angeles Times*, March 12, 2013.

12. Footnote 26 discusses how the results extend when there are more than two types.

13. As Section 2.2.4 explains, this assumption gives the firm the lowest possible equilibrium payoff upon rejection in the constructions in later sections. If the certifier were able to rate a firm who rejected the scheme, giving the firm a null rating if and only if a firm rejected would be optimal, since this minimizes the firms rejection payoff.

scheme.[14]

The scheme, alongside the firm's participation decision, induces a repeated game in the second stage. In this stage, with periods $t = 0, 1, \ldots$, the firm interacts with the consumers. At the beginning of each period, the firm's rating is updated before a new consumer enters. The consumer pays the firm its expected payoff from the trade outcome upfront. The firm then chooses effort, which generates a noisy outcome. Finally, the consumer leaves. Figure 2.1 summarizes the timing of moves.

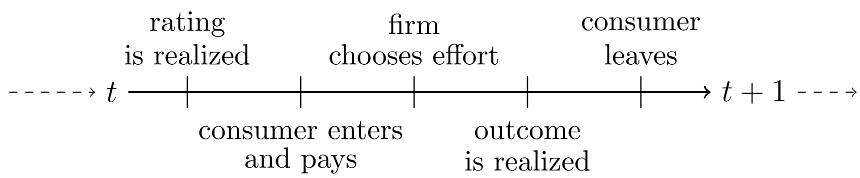

Figure 2.1: Timing of moves in each period $t = 0, 1, \ldots$

There is a unit continuum of effort levels. A competent firm chooses effort from the continuum. An inept firm can only exert zero effort. Effort level e yields a good outcome $\bar{y}$ with probability $e(1-\rho)+(1-e)\rho$ and a bad outcome $\underline{y}$ with complementary probability, where $\rho \in (0, \frac{1}{2})$.[15] Thus, higher effort is more likely to lead to a good outcome.[16]

Consumers receive a payoff of 1 from a good outcome and a payoff of 0 from a bad outcome. A firm's profit in each period equals a consumer's payment minus her effort cost. An effort level e costs the firm ce, where $c > 0$. I assume that high effort, *i.e.*, effort 1, is efficient, so that $e = 1$ maximizes the expected stage surplus $e(1 - \rho - c) + (1 - e)\rho$. In turn, this implies that the social surplus associated with

14. In practice, it is often the case that certifiers set fees with two components—one that is common to all participating firms, and the other one that depends on the individual participating firms' characteristics such as the firms' sizes and operational structures.

15. The assumption that the probability ρ is independent of effort e is for notational convenience.

16. The results extend when there are more outcomes associated with the effort levels, so long as each consumer's expected payoff given high effort is $1 - \rho$ and that given zero effort is ρ. As will be evident, ratings in the constructions that follow plainly induce payments that are linear combinations of the expected payoff given high effort and that given zero effort.

high effort exceeds that with zero effort:

$$1 - \rho - c > \rho. \tag{2.1}$$

In an induced game, the time, the scheme and the participation rule are commonly known. The firm's type and effort choices are her private information. The firm observes the rating and outcome in each period. Each consumer sees the current rating and a finite number K of most recent past ratings.[17]

I now define a (certification) scheme.

Definition 2.1. *A certification scheme is a pair (f, S), where*

$$f : \{C, I, \varnothing\} \to \mathbb{R}_+$$

defines a fee given a firm's participation decision, and $S = (S_t)_{t=0}^{\infty}$ is a rating system where each

$$S_t : \{C, I, \varnothing\} \times R^t \times \{\bar{y}, \underline{y}\}^t \to \Delta(R)$$

maps the firm's participation decision and her previous ratings and outcomes up to and including time $t-1$ to a distribution over a set of ratings R, from which the period-t rating is drawn. The mapping f has the restriction that $f(\varnothing) = 0$, and the mapping S has the restriction that $S_t(\varnothing, \cdot) = 1 \circ \{\varnothing\}$ for each period t.

The fees and the system (in particular, the range R) are freely chosen by the certifier. The restriction $f(\varnothing) = 0$ says that rejection is free. The restriction

17. Finite past rating observations capture the feature of limited records that are prominent in practice. I show in Section 2.9 that if consumers observe infinitely many past ratings, then FSO schemes, as I formally define in Definition 2.2, fail to exist, not even virtually in the sense defined in Section 2.5. Moreover, it does not matter whether consumers also observe outcomes. One interpretation is that consumers observe their outcomes and report to the certifier before exiting. Another interpretation is that the firm sells a credence good; consumers do not observe the outcomes but the certifier possesses the expertise to do so.

$S_t(\varnothing, \cdot) = 1 \circ \{\varnothing\}$ for each period t says that the system displays a null rating in each period upon a firm's rejection. Let $G(f, S, d)$ denote a game induced by a scheme (f, S) and a decision $d \in \{C, I, \varnothing\}$.

I next describe strategies, payoffs and equilibrium in an induced game in the second stage. I then define the certifier's problem in the first stage. Finally, I discuss my modeling choices.

2.2.2 Induced Game

Consider an induced game $G(f, S, d)$ in the second stage. This is a repeated game with incomplete information that unfolds as described in Section 2.2.1. Given a prior belief μ that the firm is competent, a current rating r and K recent past ratings $q^K := (r^{-K}, \ldots, r^{-1})$, where r^{-k} denotes the rating published k periods ago, a period-t consumer forms a belief that the firm is competent via Bayes' rule given its knowledge of the scheme (f, S) and the participation rule π. This belief, denoted by $\varphi_t(q^K, r)$, is interpreted as the firm's reputation.

Let $H_t := (R \times [0, 1] \times \{\bar{y}, \underline{y}\})^t \times R$ be the set of the firm's histories in period t before she chooses effort, consisting of previous ratings, efforts and outcomes up to and including period $t - 1$ and the period-t rating. The firm's strategy is a collection $\sigma^d = (\sigma^d_t)^\infty_{t=0}$, where each

$$\sigma^d_t : H_t \times \{C, I\} \to [0, 1]$$

specifies an effort level in period t after each firm's history $h_t \in H_t$ given her type θ, with the restriction that $\sigma^d_t(\cdot, I) = 0$. When there is no risk of ambiguity, I write the competent firm's strategy $\sigma^d_t(\cdot, C)$ simply as $\sigma^d_t(\cdot)$. The strategy σ^d induces a probability measure P^d_θ over the set of infinite histories in the induced game, conditional on the firm's type θ.

A consumer's payment, which is its expected payoff from an outcome, depends on its belief on the firm's type and on the competent firm's effort. Specifically, a period-t consumer who observes a current rating r and the K recent past ratings q^K pays the firm

$$p_t(q^K, r) = \rho + (1 - 2\rho)\varphi_t(q^K, r)\mathbf{E}^{\pi}\left[\mathbf{E}^{P^d_C}[\sigma^d_t(h_t)|q^K, r]\middle|q^K, r\right]. \tag{2.2}$$

The conditional expectations $\mathbf{E}^{\pi}[\cdot|q^K, r]$ and $\mathbf{E}^{P^d_C}[\cdot|q^K, r]$ capture that given ratings (q^K, r), a consumer forms beliefs on the firm's participation decision (and hence the induced game being played) given the participation rule, and on the competent firm's histories (which determine her effort) given her strategy. Given the payoff normalization on outcomes (payoff of 1 from a good outcome and 0 from a bad outcome), the payment is plainly the consumer's belief of receiving a good outcome.

A type-θ firm, with discount factor $\delta \in (0, 1)$, chooses a strategy to maximize profits:

$$(1 - \delta)\mathbf{E}^{P^d_\theta}\left[\sum_{t=0}^{\infty} \delta^t(p_t(q^K, r) - ce_t)\right], \tag{2.3}$$

where e_t denotes the effort level chosen in period t.

The solution concept that I use is *weak perfect Bayesian equilibrium* (Mas-Colell, Whinston, and Green, 1995, p. 285). In such an equilibrium, denoted by (σ^d, φ), σ^d is maximizing after each history of the firm of each type given consumers' beliefs $\varphi = (\varphi_t)_{t=0}^{\infty}$, and the beliefs are consistent with Bayes' rule given σ^d whenever possible.[18]

Observe that the sole role of each rating is to influence a consumer's belief of receiving a good outcome and, in turn, the payment as given by (2.2). Hence, one can, without loss, take the set of non-null ratings to be the set of beliefs. Hereafter,

18. An equilibrium exists, in which the competent firm always exerts zero effort. In this equilibrium, on path, each consumer pays the firm ρ irrespective of the realized rating, according to (2.2).

I assume that $R = [0,1] \cup \{\varnothing\}$.

Finally, in an induced game upon rejection, a competent firm optimally exerts zero effort as null ratings convey no information regarding her past play to consumers. Thus, if a scheme displays a null rating with zero probability upon a firm's acceptance, then, in equilibrium, each consumer who sees a null rating infers that the firm has rejected the scheme. This means that, upon seeing a null rating, the consumer expects zero effort and pays the firm ρ, according to (2.2), irrespective of its belief on the firm's type.

2.2.3 The Certifier's Problem

To formulate the certifier's problem in the first stage, I allow the certifier to not only choose a scheme but also a participation rule for the firm, subject to constraints that the firm optimally follows the rule given the scheme and some collection of equilibria, one in each of the three induced games.[19] I denote a collection of equilibria by $(\sigma, \varphi) := \{(\sigma^d, \varphi)\}_{d \in \{C,I,\varnothing\}}$. Moreover, by the revelation principle, I assume without loss that the certifier chooses a scheme and a participation rule to incentivize truthful reporting by a firm upon acceptance at the optimum.

Given a scheme (f, S), a participation rule π, a decision d, and an equilibrium (σ^d, φ) in the induced game, define

$$U(d, \sigma^d; \theta) := (1-\delta)\mathbf{E}^{P_\theta^d}\left[\sum_{t=0}^{\infty} \delta^t (p_t(q^K, r) - ce_t)\right] - f(d) \tag{2.4}$$

as a type-θ firm's payoff, that is, her profits from the interaction with consumers

19. That is, the certifier picks his preferred participation rule and collection of equilibria. This aligns with the mechanism design literature in which the principal picks his favorite equilibrium. An alternative formulation of the certifier's problem could be that he faces his "worst-case" collection of equilibria as in the robust mechanism design literature. The solution to such problem is trivial in my setting, as the worst-case collection of equilibria calls for consistent zero effort by the competent firm and the certifier obtains zero payoff.

minus the fee payment to the certifier. Then the certifier's problem is as follows. The certifier chooses a scheme (f, S) and a participation rule π to maximize his expected payoff

$$\mu\pi(C|C)f(C) + (1-\mu)\pi(I|I)f(I), \tag{2.5}$$

subject to incentive compatibility that each type-θ firm truthfully reports upon acceptance:

$$U(\theta, \sigma^{\theta}; \theta) \geq U(\theta', \sigma^{\theta'}; \theta), \quad \theta' \neq \theta,$$

as well as individual rationality that each type-θ firm follows the participation rule:

$$\pi(d|\theta) \geq \pi(d'|\theta) \quad \Longrightarrow \quad U(d, \sigma^{d}; \theta) \geq U(d', \sigma^{d'}; \theta),$$

in some collection of equilibria (σ, φ).[20]

2.2.4 Discussion of Modeling Choices

In the model, adverse selection arises as the firm's type is her private information. Moral hazard arises as the competent firm faces a myopic temptation to shirk against consumers. Absent a scheme, the competent firm optimally exerts zero effort, as consumers do not observe any past play. The firm's equilibrium payoff is then ρ. The firm would obtain higher profits if she could convince consumers that she is competent and that she will exert high effort. By adequately revealing information to consumers, schemes may act as a *screening* device to address adverse selection and as a *commitment* device to address moral hazard. Expression (2.2) highlights the role

20. I show that a solution to the certifier's problem exists by construction for the parameter values under consideration in later sections.

ratings play in coordinating consumers' beliefs of the firm's type and effort.

I pause for a brief discussion of my modeling choices. First, I have assumed that consumers pay their expected payoffs. This captures the competitive demand for the firm's service and is familiar from existing reputation models.[21]

Second, observe that by definition, I allow a scheme to display a null rating upon a firm's acceptance. For example, in the *Intertek* example, the null rating is interpreted as "not certified." Upon seeing (only) the null rating, consumers are unsure whether the firm did not sign up for certification or she signed up but failed to be certified.

Third, the fee in my model captures the source of a certifier's revenue and is familiar from the certification literature. In practice, the monetary transfer from a firm to a certifier often involves a stream of payments that are fixed *ex ante*. The fee in my model can without loss be interpreted as a discounted sum of a stream of these payments. While output-contingent fees are theoretically attractive, they are ruled out in my model, as is often the case in practice.

Finally, note also that even if the firm had the choice to terminate the relationship with the certifier in each period, she would have no strict incentive to do so in equilibrium, as continuing allows her to be at least as well off.

2.3 Full-Surplus Optimality

My analysis focuses on *full-surplus-optimal (FSO) schemes*, which are schemes that allow the certifier to seize the full net market surplus. In this section, I formally define an FSO scheme. I derive properties of the payoffs and behavior of both types of firms in games induced by any FSO scheme, as well as properties of any FSO scheme, when an FSO scheme exists. This section is, however, silent on when such a scheme exists. Instead, the properties help to motivate the constructions of FSO schemes in later

21. See, *e.g.*, Holmström (1999), Mailath and Samuelson (2001) and Board and Meyer-ter Vehn (2013).

sections. Proofs are in Section 2.10.

To define full-surplus optimality, I first derive the certifier's payoff when he extracts the full net market surplus, *i.e.*, an upper bound on his payoff. *Ex ante*, the probability of a firm being competent (resp., inept) in any period is μ (resp., $1-\mu$), by the martingale property of beliefs. If a firm is competent (resp., inept), then the maximum market surplus is $1-\rho-c$ (resp., ρ). A firm's payoff upon rejection is at least ρ, which is the lower bound on her equilibrium payoff and is guaranteed by a firm who rejects and then always exerts zero effort. The certifier's payoff when he extracts the full net market surplus is therefore

$$\bar{W} := \mu(1-\rho-c-\rho) + (1-\mu)(\rho-\rho) = \mu(1-2\rho-c). \tag{2.6}$$

Definition 2.2. *A scheme (f, S) is FSO if the certifier attains the payoff $\bar{W}$ in the certifier's problem by using this scheme given some participation rule and some collection of equilibria.*

Fix an FSO scheme (f, S), and its associated participation rule π and collection of equilibria. The derivation of $\bar{W}$ in (2.6) reveals several key properties of the participation rule, the scheme and the collection of equilibria. First, the collection of equilibria must be *efficient*, namely, that the competent firm exerts high effort consistently (*i.e.*, after every history that occurs with positive probability) in every induced game $G(f, S, d)$ in which she plays with positive probability, that is, $\pi(d|C) > 0$. This ensures that the competent firm produces the maximum surplus $1-\rho-c$ upon acceptance. Efficiency implies that FSO schemes make a firm as valuable to consumers as she can be. Moreover, FSO schemes maximize social surplus because consumers pay the firm their expected payoffs, and the firm's payment to the certifier is simply a transfer.

The second property is that the participation rule satisfies *full competence accep-*

tance, *i.e.*, $\pi(C|C) = 1$. (Truthful reporting follows from incentive compatibility.) This ensures that the certifier extracts the competent type's surplus with probability one. Indeed, efficiency requires full competent acceptance: if a competent firm rejects with positive probability, then she optimally exerts zero effort in the game upon rejection (Section 2.2.2), disrupting efficiency.

These two properties imply that *positive inept acceptance*, namely, $\pi(I|I) > 0$, is necessary, although only the competent type can produce a positive net surplus for the certifier by (2.6). Suppose on the contrary that the participation rule specifies that only the competent firm accepts. If the certifier's payoff equals $\bar{W}$, the competent firm must then receive a payment $1-\rho$ in each period upon acceptance. But then she would deviate to exert zero effort and save the cost after each history upon acceptance (and truthfully reporting), disrupting efficiency.[22]

Moreover, the collection of equilibria satisfies *full extraction*, *i.e.*, a firm's payoff (2.4) equals ρ independent of her type θ given any participation decision d she chooses with positive probability, that is, $\pi(d|\theta) > 0$. This ensures that the certifier seizes all the available net surplus in the market.

Given a participation rule and a collection of equilibria, I say that a scheme has a *low certification standard* if the rating system displays a rating that does not fully reveal to a consumer the firm's type with positive probability upon a firm's acceptance; a scheme is *coarse* if there is a rating that is displayed with positive probability upon a firm's acceptance and does not fully reveal to a consumer the entire history of outcomes.

Any FSO scheme must have a low certification standard and must be coarse. Ab-

22. An inept firm need not accept with probability one given an FSO scheme. On the one hand, if an inept type accepts with a higher probability, the certifier extracts revenues from the inept firm with a higher probability. Yet, a higher inept firm's acceptance probability lowers consumers' beliefs that a firm is competent and in turn the revenues that can be extracted by the certifier. The derivation of (2.6) shows that these two countervailing forces on the certifier's payoff balance out. Footnote 23 identifies an explicit construction where an inept firm mixes in her acceptance given an FSO scheme.

sent a low standard, ratings upon a competent report fully reveal that the firm is competent. Given efficiency, payments always equal $1-\rho$ by (2.2), again disrupting the firm's incentives for high effort upon acceptance. If a scheme has a low standard but is not coarse, then there is a history after which a competent firm, who has accepted the scheme and consistently exerted high effort in the past, optimally deviates to exert zero effort, as consumers' beliefs (and thus payments by (2.2)) are then virtually identical irrespective of future outcomes. This again disrupts efficiency.

I collect these observations in Lemma 2.1 below, together with a converse result and a condition on the parameters that is necessary for efficiency.

Lemma 2.1. *If a scheme is FSO given a participation rule and a collection of equilibria, then the participation rule satisfies full competence acceptance and positive inept acceptance, and the collection of equilibria satisfies efficiency and full extraction. Moreover, the scheme must have a low certification standard and must be coarse. Conversely, if a scheme and a participation rule solve the certifier's problem given a collection of equilibria that satisfies efficiency and full extraction, then the scheme is FSO. Finally, an efficient collection of equilibria exists only if*

$$c < \bar{c} \equiv \bar{c}(\mu, \rho) := \delta(1-2\rho)^2. \tag{2.7}$$

Condition (2.7) can be interpreted as follows. Given that the effort cost is linear in the effort level, to motivate consistent high effort, the effort cost must be smaller than the largest difference between the discounted expected benefit of high effort versus that of zero effort:

$$c < \delta\big\{(1-\rho)\overbrace{\big((1-\rho)-\rho\big)}^{\substack{\text{largest positive}\\ \text{payment difference}\\ \text{upon good outcome}}} + \rho\overbrace{\big(\rho-(1-\rho)\big)}^{\substack{\text{largest negative}\\ \text{payment difference}\\ \text{upon bad outcome}}}\big\} = \delta(1-2\rho)^2.$$

This largest difference is such that, upon a good (resp., bad) outcome, which is more likely conditional on high (resp., low) effort, the competent firm achieves the largest positive (resp., negative) difference in payments. The above inequality then follows because payments, given by (2.2), are bounded above by $1 - \rho$ and below by ρ. The lower bound is strict since payments strictly exceed ρ given efficiency.

Condition (2.7) refines the static efficiency condition (2.1). When a competent firm exerts high effort in an induced game, she incurs the cost immediately. However, she captures the benefit of high effort, namely higher continuation profits upon a good outcome relative to a bad outcome, only partially. This is because future profits are discounted and the competent firm may unluckily produce a bad outcome despite having exerted high effort, thereby failing to capture higher continuation profits.

Lemma 2.1 holds for *any* discount factor. One might conjecture that as the discount factor tends to 1, the certifier's optimal payoff tends to $\bar{W}$ even if the inept type rejects for sure. This is not true. Indeed, if the participation rule specifies that the inept firm rejects with probability one, then ratings at best induce *frequent* (but not consistent) high effort by the competent firm upon her acceptance, and the certifier's optimal payoff falls short of $\bar{W}$. That is, the folk theorem does not hold: to incentivize high effort, the competent firm must face low continuations with positive probability. Since consumers are short-lived and pay the firm their expected payoffs, low continuations necessarily require a competent firm to shirk and induce low payments. These low continuations do not vanish in the patient limit.

Low certification standard and efficiency together imply that any FSO scheme must transfer some surpluses from competent to inept firms via some ratings that I call "transfer" ratings. As the preceding discussion makes clear, any positive surplus collected by the certifier is produced by the competent firm. Lemma 2.1 also shows that the certifier can never achieve the payoff $\bar{W}$ by extracting surplus only from the competent type. As the constructions in later sections explicitly illustrate, to fully extract the net market surplus, the certifier must rely on these transfer ratings to motivate the competent firm to exert high effort and then to transfer some surpluses produced by a competent firm to an inept firm and then extract positive net surpluses from *both* types, which sum to the full net surplus.

2.4 Honors Certification: Low Effort Costs

I now turn to construct FSO schemes. By Lemma 2.1, for fixed (μ, ρ, δ), I restrict attention to effort costs $c < \bar{c}$, where $\bar{c}$ is defined in (2.7). Moreover, in this section, I further restrict attention to smaller effort costs

$$c \leq \bar{c}_2 := \delta(1-2\rho)^2 \left(\frac{1-\mu}{1-\mu+\mu\rho} \right) < \bar{c} = \delta(1-2\rho)^2. \qquad (2.8)$$

The case of higher effort costs is analyzed in Section 2.5.

Condition (2.8) allows me to construct a simple FSO scheme with a stationary Markov system that displays only two non-null ratings, denoted by r_0 and r_1, upon a firm's acceptance, highlighting the key insights of my model. I say that a system is stationary if the payment $p_t(q^K, r)$ induced by a consumer's observation of past and current ratings (q^K, r) displayed by the system is independent of time t, and it is Markov if each map S_t of the system depends only on the most recent rating and the most recent outcome. Given a stationary Markov rating system, I say that the associated scheme and ratings are stationary Markov. As will be seen in Section 2.5, a more involved rating structure is required for full-surplus optimality in the case of higher effort costs.

To ease the exposition, I further assume here that each consumer observes only the current rating but not the past ratings, *i.e.* $K = 0$. The main ideas in this section readily extend to the case of $K > 0$, where consumers observe a positive number of past ratings. In this case, the range of costs over which the scheme constructed in this section is FSO is smaller. The case of $K > 0$ is relegated to Section 2.7 as it complicates the exposition. To be sure, even if consumers observe only the current ratings, schemes are defined to be flexible enough to disclose as much information via a current rating as one might wish, including all previous outcomes, ratings and participation decision.

My first main result is:

Proposition 2.1. *Fix* (μ, ρ, δ) *and* $K = 0$*. There exists a stationary Markov scheme* (f_2, S_2) *that is FSO for any* $c \leq \bar{c}_2$*.*

In the remainder of this section, I first present the construction. Then, I interpret the constructed scheme as a "low-standard, honors certification scheme."

Because the ratings r_0 and r_1 are non-null, a firm's rejection payoff is ρ, irrespective of consumers' off-path beliefs on the firm's type (Section 2.2.2). Thus, consumers need *not* believe that a firm who deviates to reject is inept for sure for the construction.

The system S_2 operates as follows. Given a competent report, the initial rating is r_0 with probability $1 - \rho$ and is r_1 with probability ρ. The system thereafter displays rating r_0 upon a good outcome and rating r_1 upon a bad outcome. Given an inept report, the system always displays rating r_1. The fee f_2, given in Section 2.8.1, keeps an accepting firm at her rejection payoff ρ for a given participation rule and collection of equilibria, denoted by π^* and (σ^*, φ^*) and described next, making acceptance individually rational.

The participation rule π^* specifies that a firm accepts with probability one and reports truthfully. The collection of equilibria (σ^*, φ^*) calls for the competent firm to exert high effort given each rating in the game $G(f_2, S_2, C)$, and to always exert zero effort in other induced games. Thus, the ratings r_0 and r_1 induce, respectively, consumers' beliefs 1 and

$$\tilde{\mu} := \frac{\mu\rho}{\mu\rho + 1 - \mu} < 1 \tag{2.9}$$

that the firm is competent in each period. Rating r_0 separates a competent firm from an inept firm, while rating r_1 pools both types. By (2.2), there are two payments, $p(r_0) = 1 - \rho$ and $p(r_1) < 1 - \rho$. (The time subscript is omitted due to stationarity.)

By construction, the collection of equilibria satisfies full extraction and efficiency. To verify that the scheme is FSO, it suffices to verify, first, that the competent firm has no profitable deviation from consistent high effort in the game $G(f_2, S_2, C)$ and optimally exerts zero effort in the games $G(f_2, S_2, I)$ and $G(f_2, S_2, \varnothing)$, and, second, that a firm does not profit by deviating to misreport her type upon acceptance.

Clearly, the competent firm optimally always exerts zero effort in the games $G(f_2, S_2, I)$ and $G(f_2, S_2, \varnothing)$ since the ratings do not reflect her past play in the games. In the game $G(f_2, S_2, C)$, the competent firm consistently exerts high effort to chase the high price $p(r_0)$ and to avoid the low price $p(r_1)$. Since the firm's cost is linear in effort, and because the system S_2 upon a competent report depends only on the recent outcome, the incentive constraint for high effort is that the discounted expected benefit of high effort versus zero effort in the following period outweighs the cost incurred in the current period:

$$\delta\Big[\ \underbrace{\big((1-\rho)p(r_0) + \rho p(r_1)\big)}_{\text{expected payment after high effort}}\ -\ \underbrace{\big(\rho p(r_0) + (1-\rho)p(r_1)\big)}_{\text{expected payment after zero effort}}\ \Big] \geq c.$$

The inequality is equivalent to $c \leq \bar{c}_2$ in (2.8).

A firm does not profit by deviating to misreport upon acceptance. Both types obtain a payoff ρ given the participation rule and the collection of equilibria. A competent firm who deviates to misreport as inept optimally exerts zero effort in the game $G(f_2, S_2, I)$ and earns identical profits as an inept type does. Her payoff upon deviation remains as ρ. An inept type who misreports as competent obtains rating r_0 less frequently and obtains lower profits than the competent firm does. The deviation makes her strictly worse off, as a competent firm obtains exactly ρ after paying the competent fee.[23]

The FSO scheme can be interpreted as an honors certification scheme with a low certification standard. The rating r_0 represents "certified with honors," and the rating r_1 represents "certified." Given the participation rule π^* and the collection of equilibria (σ^*, φ^*), the scheme has a low certification standard, as the inept firm is "certified," *i.e.*, achieves the rating r_1, with probability one. The "certified" rating r_1 is a transfer rating that the certifier relies upon to extract the full net surplus. An inept firm receives a profit $p(r_1) > \rho$, despite failing to produce any surplus above ρ; a competent firm receives a profit $p(r_1) - c < 1 - \rho - c$, despite producing a surplus $1 - \rho - c$. Finally, the honors structure makes the scheme coarse, as the ratings only reveal the last outcome in each period.

2.5 Honors Certification: High Effort Costs

I next turn to present my second main result that concerns full-surplus optimality given higher effort costs. Specifically, I characterize a family of schemes given which the certifier achieves full-surplus optimality when the firm is arbitrarily patient. More-

23. In fact, if $c < \bar{c}_2$, then the participation rule associated with the scheme (f_2, S_2) need not specify *full* inept acceptance, $\pi(I|I) = 1$, for full-surplus optimality. Specifically, in the above construction, replace the pooling rating r_1 by an alternative pooling rating r_1' that induces belief

$$\tilde{\mu}' = \frac{\mu\rho}{\mu\rho + (1-\mu)\pi(I|I)} \geq \tilde{\mu}, \tag{2.10}$$

where the probability $\pi(I|I)$ is chosen such that the competent firm's incentive constraint for high effort

$$\delta\big[\big((1-\rho)p(r_0) + \rho p(r_1')\big) - \big(\rho p(r_0) + (1-\rho)p(r_1')\big)\big] \geq c \tag{2.11}$$

binds. Also adjust the fee such that both types receive a payoff ρ upon acceptance and truthfully reporting. (2.10) follows as it is now less likely that an accepting firm is inept when an inept firm mixes between accepting and rejecting. By construction, when $c = \bar{c}_2$, $\pi(I|I) = 1$. As $\pi(I|I)$ falls, $p(r_1')$ rises, weakening punishment upon a bad outcome and (2.11) holds only for smaller costs. Thus, for any $c \in (0, \bar{c}_2)$, there is a unique $\pi(I|I) \in (0, 1)$ such that (2.11) binds, which is the lowest inept acceptance probability that can be chosen without disrupting efficiency. By (2.11), the competent firm optimally exerts high effort upon acceptance. The inept firm is willing to mix between accepting and rejecting, since she receives a payoff ρ either way.

over, I drop Condition (2.8) that $c \leq \bar{c}_2$ and consider costs $c < (1-2\rho)^2$. The bound $(1-2\rho)^2$ is $\bar{c}$, defined in (2.7) when evaluated in the limit $\delta \to 1$.[24]

Given any $\varepsilon > 0$, I say that a scheme is ε-FSO if the certifier obtains a payoff of at least $\bar{W} - \varepsilon$ by using the scheme given some participation rule and some collection of equilibria in the certifier's problem.

Definition 2.3. *A family* $\{(f_\varepsilon, S_\varepsilon) | \varepsilon > 0\}$ *of schemes is virtually FSO if for each* $\varepsilon > 0$, *there exists* $\underline{\delta} \equiv \underline{\delta}(\mu, \rho, c, K, \varepsilon) \in (0,1)$ *such that for every* $\delta \in [\underline{\delta}, 1)$, *the scheme* $(f_\varepsilon, S_\varepsilon)$ *is* ε*-FSO.*

Given a virtually FSO family of schemes, the certifier achieves full-surplus optimality when the firm is arbitrarily patient. My second main result is:

Proposition 2.2. *Fix* (μ, ρ, c, K) *such that* $c < (1-2\rho)^2$ *and* $K \geq 0$. *Then a virtually FSO family of schemes exists only if*

$$c < (1-2\rho)^2 \frac{1-\mu}{\rho}. \tag{2.12}$$

Conversely, given Condition (2.12), *there exists a virtually FSO family of schemes* $\{(f_\varepsilon^K, S_\varepsilon^K) | \varepsilon > 0\}$ *in which each scheme* $(f_\varepsilon^K, S_\varepsilon^K)$ *is stationary Markov.*

The construction of the scheme $(f_\varepsilon^K, S_\varepsilon^K)$ for each $\varepsilon > 0$ in the proposition calls for a large number of stationary Markov ratings. The number of such ratings decreases in ε, and grows to infinity as ε vanishes.[25] As I show in the next section, each

24. I show that patience is necessary for efficiency when the effort cost is large and close to $(1-2\rho)^2$ in Section 2.9. Intuitively, to achieve efficiency, punishment upon a bad outcome produced by the competent type after any history must be sufficiently harsh. Given discounting, such harsh punishment necessarily involves ratings that induce low enough consumers' beliefs about the firm's competence and are realized soon enough after a bad outcome. But, by Bayes' rule, these ratings that are realized soon enough cannot induce low enough consumers' beliefs.

25. In Section 2.9, I show that when effort cost is sufficiently large, any stationary Markov scheme that is FSO must feature a large number of ratings. Given a stationary Markov structure, the number of ratings restricts the extent to which consumers' beliefs, payments and a firm's continuation profit upon a bad outcome can fall, and in turn restricting the range of costs over which efficiency can be achieved in equilibrium.

such scheme $(f_\varepsilon^K, S_\varepsilon^K)$ also has an "honors certification" implementation that features only three stationary ratings: a null rating that represents "not certified," and two non-null ratings that represent "certified" and "certified with honors," upon a firm's acceptance; moreover, each such scheme has a low certification standard, as an inept firm who participates is "certified" with probability close to one.

If $\rho > 1 - \mu$, then Condition (2.12) strengthens the condition $c < (1 - 2\rho)^2$. Otherwise, Condition (2.12) is vacuous given $c < (1 - 2\rho)^2$. Condition (2.12) is interpreted as follows. For efficiency, there must be top ratings that induce high payments to motivate a competent firm to exert high effort at lower ratings. Top ratings induce high payments in equilibrium only if, at these ratings, the competent firm exerts high effort. To motivate high effort at top ratings, the competent firm who produces bad outcomes must then receive "low" ratings that induce sufficiently low payments. Given efficiency, payments are low only if consumers are unsure of the firm's type. When $\rho > 1 - \mu$ so that monitoring is sufficiently noisy and the effort cost is sufficiently large so that Condition (2.12) fails, low ratings fail to induce sufficiently low payments.

2.5.1 A Stationary Markov Construction

I now present the construction. In the main text, I again assume that each consumer observes only the current rating but no past ratings, *i.e.*, $K = 0$, to ease the exposition. The main ideas readily extend to the case of $K > 0$, which is relegated to Section 2.7. At the end of Section 2.5.2, I briefly comment on why this is so.

Fix $(\mu, \rho, c, \varepsilon)$ such that $c < (1 - 2\rho)^2 \min[\frac{1-\mu}{\rho}, 1]$. Denote the scheme $(f_\varepsilon^0, S_\varepsilon^0)$ to be constructed by $(f_\varepsilon, S_\varepsilon)$ to ease notations. Fix a virtual cost $\hat{c} \in (c, (1 - 2\rho)^2 \min[\frac{1-\mu}{\rho}, 1])$. The set of ratings is $R = \{r_0, \dots, r_N\}$, where each r_n is non-null, and the number N is chosen below. Define a number $\gamma \in (0, 1]$, which is a function of $(\mu, \rho, \hat{c})$, formally given in (2.15) in Section 2.7, and is used below as a transition probability over ratings.

The participation rule of interest, denoted by π_ε, calls for both types to accept and truthfully report with probability one. The collection of equilibria of interest, denoted by $(\sigma_\varepsilon, \varphi_\varepsilon)$, is such that the competent type exerts high effort given a rating $r \neq r_N$ and exerts zero effort given rating r_N in the induced game $G(f_\varepsilon, S_\varepsilon, C)$, and always exerts zero effort in other induced games. The firm's rejection payoff is ρ, as the ratings in R are non-null.

The rating system S_ε operates as follows and is depicted in Figure 2.2. Upon a competent report, the initial rating is drawn from the set R according to a distribution $\lambda^C \equiv (\lambda_n^C)_{n=0}^N \in \Delta(R)$, which is the stationary distribution over R given the participation rule π_ε and the collection of equilibria $(\sigma_\varepsilon, \varphi_\varepsilon)$. Given a current rating r_0, a good outcome leads to rating r_0 with probability one, and a bad outcome leads to rating r_1 with probability one. Given a current rating r_n, $0 < n < N$, a good outcome leads to rating r_{n-1} with probability γ and to rating r_n with probability $1-\gamma$; a bad outcome leads to rating r_{n+1} with probability one. Finally, given a current rating r_N, regardless of the outcome, the rating in the next period is r_{N-1} with probability $(1-\rho)\gamma$ and remains as r_N with complementary probability.

Upon an inept report, the initial rating is drawn from a set $R^I := R \setminus \{r_0\}$ according to a distribution $\lambda^I \equiv (\lambda_n^I)_{n=1}^N \in \Delta(R^I)$, which is chosen such that, given the participation rule and the distribution λ^C, consumers' beliefs upon each rating r_n, $n > 0$, are identical:

$$\underline{\mu} := \varphi(r_n) = \frac{\mu\lambda_n^C}{\mu\lambda_n^C + (1-\mu)\lambda_n^I}, \qquad \text{for each } n > 0. \tag{2.13}$$

The fees $f_\varepsilon(C), f_\varepsilon(I)$ keep each type at the rejection payoff ρ in the collection of equilibria $(\sigma_\varepsilon, \varphi_\varepsilon)$. The fees and distributions are given in Section 2.8.2.

There are two induced beliefs: $\varphi_\varepsilon(r_0) = 1$ and $\varphi_\varepsilon(r_n) = \underline{\mu}$ for each $n = 1, \ldots, N$.

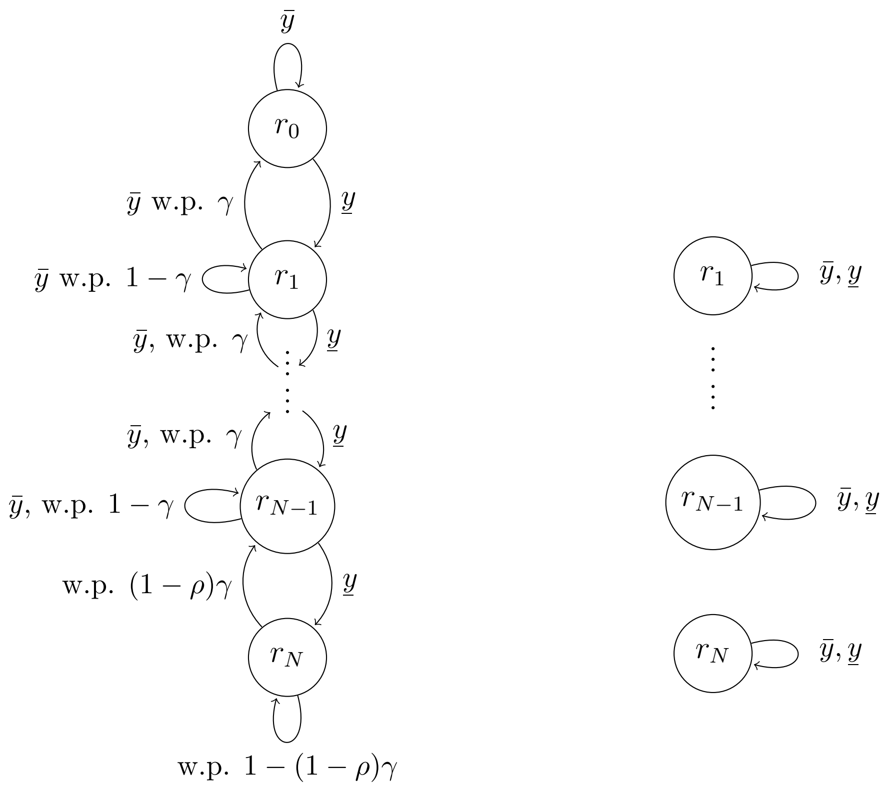

(a) Upon a competent report **(b)** Upon an inept report

Figure 2.2: The rating system S_ε

Rating r_0 is a separating rating. Ratings $r_n, n = 1, \ldots, N$, are pooling ratings. Because the competent firm exerts high effort upon receiving ratings r_n, $n = 0, \ldots, N-1$, and exerts zero effort upon receiving rating r_N, by (2.2), there are three payments associated with the ratings:

$$\bar{p} := p(r_0) = 1 - \rho,$$

$$p_{\underline{\mu}} := p(r_n) = \rho + (1-2\rho)\underline{\mu}, \quad n = 1, \ldots, N-1,$$

$$\underline{p} := p(r_N) = \rho.$$

Observe that $\bar{p} > p_{\underline{\mu}} > \underline{p}$. I now argue that $(\sigma_\varepsilon, \varphi_\varepsilon)$ is a collection of equilibria when the firm is sufficiently patient. The arguments that the firm optimally accepts and reports truthfully, and that the competent firm optimally exerts zero effort in the induced games $G(f_\varepsilon, S_\varepsilon, I)$ and $G(f_\varepsilon, S_\varepsilon, \varnothing)$ are analogous to those in Section 2.4. I focus on the competent firm's effort incentives to "climb up the ladder" for the high payment $\bar{p}$ and to avoid the other payments in the induced game $G(f_\varepsilon, S_\varepsilon, C)$.

The number $N \equiv N(\mu, \rho, \hat{c}, \varepsilon)$ of ratings is chosen to be large enough for two purposes. First, it creates sufficient incentives for high effort when a competent firm receives the top rating r_0. Upon a bad outcome, a competent firm falls down the ladder; with positive probability, she may then produce more bad outcomes and fall far down the ladder, receiving a string of low payments. Second, a large enough N ensures that the competent firm generates large enough surplus for the certifier. The payoff loss from $\bar{W}$ incurred by the certifier corresponds to the event that a competent firm receives rating r_N, upon which she exerts zero effort. The probability that a competent firm receives rating r_N falls as N increases and vanishes as $N \to \infty$. Thus, given any $\varepsilon > 0$, one can fix a sufficiently large N such that the certifier receives a payoff of at least $\bar{W} - \varepsilon$. When N is large, however, the current rating could be very far down the competent ladder and it takes a long sequence of good outcomes to obtain the high payment $\bar{p}$. The firm thus needs to be patient so that this future benefit is not discounted too much, which would disrupt incentives for high effort.

Given a larger μ or ρ, it is more difficult to motivate high effort upon rating r_0, as it is more difficult to lower consumers' beliefs that the firm is competent via the intermediate ratings. The transition probability γ thus strictly decreases in both μ and ρ to lengthen the punishment phase upon a bad outcome at rating r_0, ensuring sufficient effort incentives.

2.5.2 An Honors Certification Implementation

To maintain a Markov structure, the scheme constructed in Proposition 2.2 uses many intermediate ratings that induce identical consumers' payments. Alternatively, as I show next, one can think of the scheme as an honors certification implementation

$(f_\varepsilon, \bar{S}_\varepsilon)$ of the scheme $(f_\varepsilon, S_\varepsilon)$ that features only three stationary ratings for any $\varepsilon > 0$.

The certifier privately uses the system S_ε and displays a rating (to the firm and consumers) via the system $\bar{S}_\varepsilon$ as follows. When the system S_ε delivers rating r_0, the system $\bar{S}_\varepsilon$ displays rating p_1 that represents "certified with honors." When the system S_ε delivers rating r_n, $n = 1, \ldots, N-1$, the system $\bar{S}_\varepsilon$ displays rating $p_{\underline{\mu}}$ that represents "certified." Finally, when the system S_ε delivers rating r_N, the system $\bar{S}_\varepsilon$ displays rating $\underline{p}$ that represents "not certified."

Ratings $\bar{p}, p_{\underline{\mu}}$ and $\underline{p}$ induce beliefs $1, \underline{\mu}$ and $\underline{\mu}$, and payments $\bar{p}$, $p_{\underline{\mu}}$ and $\underline{p}$, respectively. Rating $p_{\underline{\mu}}$ induces belief $\underline{\mu}$ as consumers form expectations over ratings $r_1, \ldots, r_{N-1}$ that the system S_ε may deliver, and each of these ratings induces the belief $\underline{\mu}$.

Upon receiving rating $\bar{p}$, the competent firm infers that her hidden rating is r_0 and therefore optimally exerts high effort. Upon receiving rating $p_{\underline{\mu}}$, the competent firm forms an expectation over her hidden ratings and faces an "average" incentive constraint for high effort. Because the incentive constraints for high effort are satisfied in the stationary Markov construction, the "average" constraint holds. Finally, upon receiving rating $\underline{p}$, the competent firm knows that her hidden rating is r_N, and optimally exerts zero effort.

By construction, both types of the firm earn identical profits from consumers under S_ε and under $\bar{S}_\varepsilon$. Thus, the fees that fully extract the net surplus and keep each type at the rejection payoff ρ under S_ε are identical to those under $\bar{S}_\varepsilon$.

The associated equilibrium outcome given this honors certification implementation aligns with causal observations related to honors certification schemes in practice. Because the number N is chosen to be very large, in each period, the competent firm is very likely to be "certified." Occasionally, upon producing a streak of consecutive good outcomes, the competent firm is "certified with honors" and separates from an inept type. In rare cases where the competent firm produces a long streak of consecutive bad outcomes, her rating becomes "not certified." On the other hand, as a consequence of the low certification standard, in each period, the inept type is very likely to be "certified." The "certified" rating transfers surpluses produced by the competent firm to the inept firm, allowing the certifier to extract (almost) the full net market surplus.

As briefly described in the introduction, a key difficulty in constructing the virtually FSO family of schemes is the conflict between motivating high effort at the separating rating and introducing transfer ratings. To transfer some net surpluses from a competent firm and an inept firm, the transfer ratings must motivate high effort by the competent type and therefore induce a price above ρ. To motivate high effort at the transfer ratings, there must exist a "top" rating that induce a higher price. Thus, the competent firm must exert high effort at this top rating. However, the only channel through which a scheme motivates high effort by the competent firm at the top rating and allows the certifier to fully extract the net surplus is to assign transfer ratings to the firm upon a bad outcome. But if the competent firm obtains the transfer ratings too often, by Bayes' rule, these ratings cannot induce low enough consumers' beliefs, disrupting the incentive for high effort at the top rating. In the construction, the top rating corresponds to "certified with honors," inducing the highest possible price $1-\rho$. The transfer rating "certified" pools all of a firm's track records except a long enough streak of consecutive good outcomes and a long enough streak of consecutive bad outcomes, ensuring that consumers' belief induced by the rating is not too high.[26]

26. Observe that all previous constructions carry over if there are some "intermediate" type of firm whose effort cost exceeds the competent type's effort cost, as long as the certifier optimally incentivizes consistent high effort only from the competent type, and treats the intermediate and inept types as one "inept type" who in equilibrium exerts zero effort. This is the case, for instance, when high effort is inefficient for the intermediate type (*i.e.*, her effort cost is such that the static efficiency condition (2.1) is violated), or when high effort is efficient for the intermediate type but the prior belief that a firm is an intermediate type is sufficiently small. Otherwise, the certifier optimally motivates consistent high effort by both the competent and the intermediate types, and leaves some surplus for the competent type to prevent her from misreporting to be an intermediate type (*i.e.*, to satisfy incentive compatibility in the certifier's problem). This disrupts full extraction, but the

Finally, I remark on why much of the structure in the above constructions extends to the case of $K > 0$. When consumers observe more past ratings, the competent firm faces a stronger temptation to exert zero effort at the top rating r_0. Even if the firm produces a bad outcome at rating r_0, the next K consumers see that the firm received rating r_0 before. Given (virtually) consistent high effort, these consumers pay the firm the highest payment $\bar{p}$, even if the current ratings are not r_0. To motivate high effort at rating r_0, the certifier can lower the transition probability γ to some carefully chosen $\gamma^K < \gamma$, as defined in (2.15) in Section 2.7, to effectively lengthen the punishment phase upon a bad outcome at rating r_0.

2.6 Concluding Remarks

To conclude, this paper begins with the observation that a key role of firm certification schemes is to address adverse selection and dynamic moral hazard problem. Yet, certification schemes often appear counterproductive to do so. Specifically, in practice, honors certification schemes are pervasive, and many such schemes feature low certification standards. My results show how a certifier uses low-standard, honors certification schemes that trade off screening and incentives for quality provisions to maximize profits. Importantly, the results show that such schemes fully alleviate dynamic moral hazard and can maximize social welfare.

I close the paper with a brief discussion of some of the assumptions that I made in my analysis. First, I have followed the mechanism-design approach to certification in the literature (see, *e.g.*, Albano and Lizzeri, 2001), and allowed the certifier to offer a menu of schemes that discriminates against different types of firms. In contrast, following the traditional "industrial organization" approach, one might be interested in schemes in which the fee and the system are independent of a firm's type, as is

main insights on the honors certification structure and low certification standards carry over.

often the case in practice when certifying the quality of a finished product. I construct such a scheme in Section 2.9. The range of costs over which the scheme can induce efficiency is smaller, as the scheme does not feature a rating that separates both types and creates stronger incentives for high effort.

Moreover, I have considered a setting in which the certifier fully extracts the net surplus from the participating firms. Full extraction by a monopoly certifier is a familiar result from the certification literature. To be clear, there are different frictions in practice that may rule out full extraction, but they are beyond the scope of this paper. For example, full extraction requires a firm to accept a scheme when she is indifferent between accepting and rejecting. One might reasonably argue that a firm rejects for sure whenever she is indifferent. This might be due to reasons outside of the model, say, because signing up for certification is costly. To attract participation, the certifier might then offer a firm some additional surplus that compensates for the cost, and the remainder of the present construction carries over. Another setting that possibly rules out full extraction and is also beyond the scope of this paper, is competition between multiple certifiers.

Finally, as a first step towards examining the design of dynamic firm certification, my analysis has focused on full-surplus optimality. The solution to the certifier's profit maximization problem when full-surplus optimality is not even virtually attainable, *i.e.*, when the effort cost is such that $c > (1 - 2\rho)^2$, is an open question. In this case, the exact solution to the certifiers problem requires keeping track of histories after which the certifier would like a competent firm to exert high effort and those after which the certifier would like the firm to shirk and then optimizing over all schemes. (Of course, when the cost is sufficient large, the certifier cannot induce any high effort and thus obtain a revenue of zero.) I leave this for future research.

2.7 Appendix A: Past Rating Observations

In this appendix, I assume that each consumer observes the current rating as well as a finite number K of most recent past ratings. To illustrate the additional difficulty that the certifier faces when consumers observe some past ratings, I begin by re-examining the simple, binary rating structure in Section 2.4 with the assumption that $K = 1$. I then generalize the construction of the virtually FSO family of schemes in the case of $K = 0$ below Proposition 2.2 in the main text to the case of any finite $K \geq 0$.

2.7.1 Finitely Many Past Observations

Example 1 (Binary ratings with $K = 1$)**.** Consider a scheme (f_2', S_2), where the system S_2 is as in Proposition 2.1 and the fee f_2' is such that the certifier extracts a firm's revenue from consumers up to her rejection payoff ρ given the participation rule and collection of equilibria described below. The fee f_2' is different from the fee f_2 in Proposition 2.1, as consumers' beliefs and hence payments are different when $K = 1$. The participation rule again calls for each type to accept and report truthfully. Moreover, the collection of equilibria is again such that the competent type consistently exerts high effort upon accepting and truthfully reporting.

Contrary to the case of $K = 0$, each consumer forms a belief that the firm is competent given a current rating and the most recent past rating. Similar to the case of $K = 0$, there are two stationary beliefs on the firm being competent and two (stationary) payments. The belief is $\tilde{\mu}_1 := \frac{\mu\rho^2}{\mu\rho^2+1-\mu}$ if both ratings are the pooling rating r_1, and the belief is 1 otherwise. Payment, as given by (2.2), is $\rho + \tilde{\mu}_1(1 - 2\rho)$ if both ratings the consumer observes are r_1 and is $1 - \rho$ otherwise. The following holds:

Proposition 2.3. *Fix* (μ, ρ, δ). *The scheme* (f_2', S_2) *is FSO if and only if*

$$c \leq \bar{c}_2' := \frac{\delta^2(1-\mu)\rho(1-2\rho)^2}{1-\mu(1-\rho^2)}. \tag{2.14}$$

The proof of this proposition, containing the detailed calculations, is relegated to Section 2.9. Note that $\bar{c}_2' < \bar{c}_2$, where $\bar{c}_2$ is the cost cutoff for the scheme (f_2, S_2) to be FSO when $K = 0$, because it is more difficult to incentivize consistent high effort by the competent firm when consumers also observe the most recent past rating. Specifically, here, the competent firm is more tempted to exert zero effort when she receives a current (separating) rating r_0. Upon producing a bad outcome such that the next rating becomes r_1, the next consumer still sees that the firm has received a current rating r_0 in the past. This consumer also believes that the firm is competent who will exert high effort and thus pays the firm a high payment $1 - \rho$. This weakens the punishment upon a bad outcome and in turn weakens the competent firm's incentive to exert high effort when she receives a current rating r_0. ♦

The weakened incentives for high effort when receiving a separating rating can be restored by strengthening punishment upon a bad outcome and requiring a firm to be sufficiently patient, as in the construction for the case of $K = 0$ in Proposition 2.2 in the main text. I now present the construction of a virtually FSO family $\{(f_\varepsilon^K, S_\varepsilon^K) | \varepsilon > 0\}$ of schemes when consumers observe the current ratings as well as any finite K most recent past ratings, generalizing the construction in Section 2.5 in the main text.

Construction. Fix $(\mu, \rho, c, K, \varepsilon)$ such that $c < (1-2\rho)^2 \min[\frac{1-\mu}{\rho}, 1]$. Fix a virtual cost $\hat{c} \in (c, (1-2\rho)^2 \min[\frac{1-\mu}{\rho}, 1])$. In the scheme $(f_\varepsilon^K, S_\varepsilon^K)$, the set of ratings is $R = \{r_0, \ldots, r_N\}$ and $r_n \neq \varnothing$ for each n, where $N > K+1$ is chosen to be sufficiently large for reasons as outlined in the main text. The participation rule and the collection of equilibria are as in the case of $K = 0$. Specifically, the participation rule, denoted by π_ε^K, calls for both types to accept with probability one and then truthfully report. The collection of equilibria, denoted by $(\sigma_\varepsilon^K, \varphi_\varepsilon^K)$, is such that in the induced game upon a competent report, the competent firm exerts high effort given a current rating $r \neq r_N$ and exerts zero effort given a current rating r_N; in other induced games, the competent firm always exerts zero effort.

Rating transitions upon each report are depicted in Figure 2.3 below. Define

$$\gamma^K(\mu, \rho) = \begin{cases} 1, & \text{if } \hat{c} < \frac{(1-\mu)(1-2\rho)^2(\rho+\rho^K(1-2\rho))}{(1-\rho)(1-\mu-\rho+2\mu\rho)}, \\ \frac{(1-2\rho)^2(1-\rho^{K+1})-\mu+\mu\rho((4-\hat{c})(1-\rho)+(1-2\rho)^2\rho^K)}{(1-\mu)(1-\rho)(\hat{c}(1-\rho)+(1-2\rho)^2(1-\rho^K))}, & \text{otherwise.} \end{cases} \tag{2.15}$$

As explained below, the transition probability γ in the construction below Proposition 2.2 in the main text is γ^K evaluated at $K = 0$. By definition, $\gamma^K \in (0, 1]$, for each K. The initial rating upon a competent report is drawn according to the stationary distribution over the set of ratings R, given the participation rule π_ε^K and the collection of equilibria $(\sigma_\varepsilon^K, \varphi_\varepsilon^K)$. Given a current rating r_0, the next rating is r_0 upon a good outcome and is r_1 upon a bad outcome. Given a current rating r_n, $n = 1, \dots, N-1$, upon a good outcome, the next rating is r_{n-1} with probability γ^K and is r_n with probability $1 - \gamma^K$; upon a bad outcome, the next rating is r_{n+1}. Finally, given a current rating r_N, the next rating is r_{N-1} with probability $(1-\rho)\gamma^K$ and is r_N with probability $1 - (1-\rho)\gamma^K$, independent of the realized current outcome.

The initial rating upon an inept report is drawn according to the stationary distribution over the set $R \backslash \{r_0\}$, given the participation rule π_ε^K and the collection of equilibria $(\sigma_\varepsilon^K, \varphi_\varepsilon^K)$. All rating transitions are exogenous, *i.e.*, independent of outcomes produced by the firm. Given a current rating r_1, the next rating is r_1 with probability $1-\rho$ and is r_2 with probability ρ. Given a current rating r_n, $n = 2, \dots, N-1$, the next rating is r_{n-1} with probability $(1-\rho)\gamma^K$, r_n with probability $(1-\rho)(1-\gamma^K)$ and

r_{n+1} with probability ρ. Finally, given a current rating r_N, the next rating is r_{N-1} with probability $(1-\rho)\gamma^K$ and is r_N with probability $1-(1-\rho)\gamma^K$.

This general construction when evaluated at $K=0$ is identical to the construction below Proposition 2.2 in the main text, except that the rating transitions upon an inept report are different. Specifically, in the main text, ratings are drawn once and for all upon an inept report, and the initial distribution over the inept ratings equals the stationary distribution given the transitions upon an inept report in the current general construction. In fact, these two specifications are equivalent in the case of $K=0$ for the analysis, because, upon screening a firms type via her report in an FSO scheme, the rating transitions after an inept report matter only to the extent that the corresponding stationary distribution affects consumers beliefs on the firms type. In particular, the transition probability γ in the construction below Proposition 2.2 in the main text is plainly γ^K, defined in (2.15) when evaluated at $K=0$.

As discussed in the main text, the transition probability γ^K are chosen to ensure that the competent type optimally exerts high effort at rating r_0. When the virtual cost $\hat{c}$ is sufficiently small, pure transitions ($\gamma^K = 1$) between ratings suffice given any $c < (1-2\rho)^2$ if the firm is sufficiently patient. In contrast, when the virtual cost is large (but smaller than $(1-2\rho)^2$), the system must feature mixed transitions to effectively lengthen the punishment phase upon a bad outcome at the top rating r_0.

Again, the payoff loss incurred by the certifier from the upper bound $\bar{W}$ corresponds to the event that a competent firm receives rating r_N, upon which she exerts zero effort. The probability that a competent firm receives rating r_N strictly decreases in N, and vanishes as $N \to \infty$. The number of ratings N is chosen to be large enough to ensure that the payoff loss incurred by the certifier falls short of ε, as in the main text.

Observe that much of the rating structure from the case of $K=0$ carry over. First, rating r_0 remains a separating rating. Second, rating transitions upon an inept

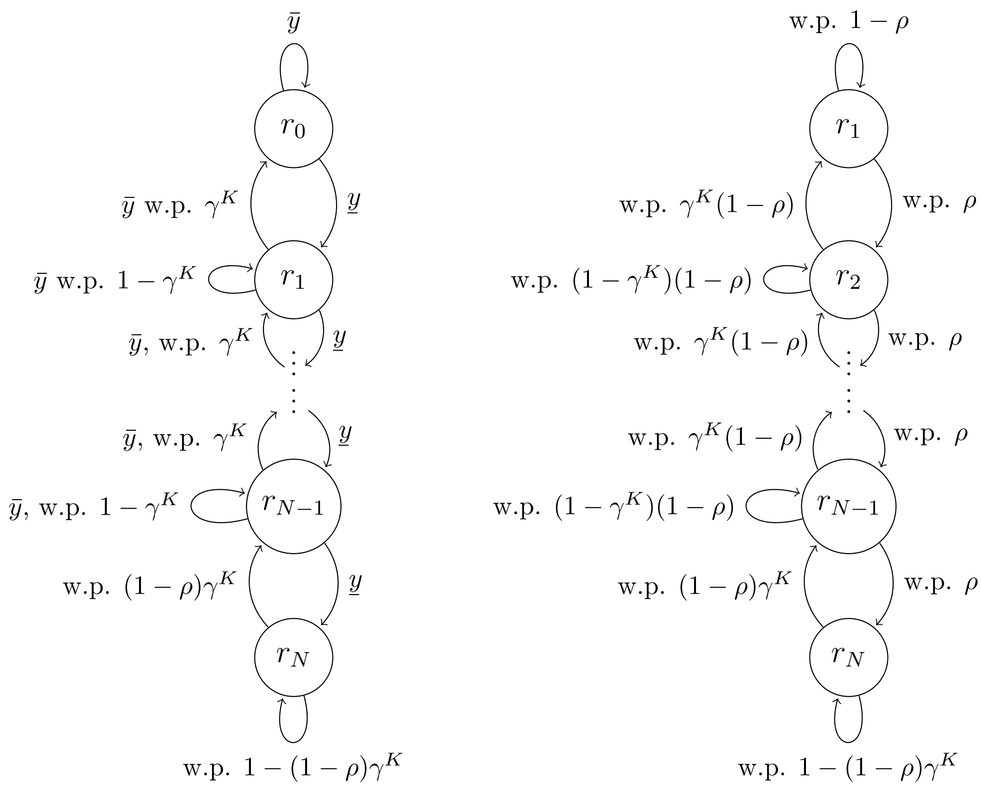

(a) Upon a competent report **(b)** Upon an inept report

Figure 2.3: Rating transitions

report are exogenous. Finally, these transitions are such that, whenever the ratings that a consumer observes do not include rating r_0, its belief that the firm is competent is identical and equal to some $\underline{\mu}^K$, where

$$\underline{\mu}^K = \mu\left[1 + \frac{(1-\mu)(\gamma^K(1-\rho)-\rho)}{\rho\big(1-(\frac{\rho}{\gamma^K(1-\rho)})^N\big)}\right]^{-1}. \tag{2.16}$$

Again, there are two (stationary) beliefs on the firm's type being competent, 1 and $\underline{\mu}^K$. There are three payments: $\bar{p} = 1-\rho$ when a consumer observes a rating r_0 (among the ratings it observes), $p_{\underline{\mu}^K} = \rho + \underline{\mu}^K(1-2\rho)$ when a consumer does not observe a rating r_0 and the current rating is not r_N, and $\underline{p}$ if the current rating is r_N. The scheme thus again has a simple stationary non-Markovian implementation with three non-null ratings.

The reason why much of the rating structure from the case of $K = 0$ easily extends is as follows. As mentioned, the key difficulty when consumers observe more past ratings is that the competent firm faces a stronger temptation to shirk when receiving a current rating r_0. Given the prescribed equilibrium in the induced game upon a competent report, even if the firm produces a bad outcome when receiving rating r_0, the next K entering consumers still see that the firm has received the rating r_0 before. These consumers also pay the firm the highest payment $\bar{p}$, even when the current ratings are not r_0. To ensure sufficient effort incentives when the firm receives a current rating r_0, the transition probability γ^K to climb up strictly decreases in the number of past observations K. Upon producing a bad outcome at rating r_0, the firm finds it harder to climb back up to the top rating r_0. The periods in which the firm receives the payment $\bar{p}$ also more likely terminate, at which point the firm starts collecting lower payments. The competent firm who is receiving a lower payment (again) finds it harder to climb up the ladder and receive the highest payment $\bar{p}$. A sufficiently patient competent firm thus optimally exerts high effort at rating r_0.

2.8 Appendix B: Omitted Details

2.8.1 The Scheme (f_2, S_2)

It remains to specify the fees:

$$f_2(C) = (1-2\rho)\left(\rho\left(\frac{\mu\rho}{\mu\rho+1-\mu}\right)+1-\rho\right)-c, \tag{2.17}$$

$$f_2(I) = (1-2\rho)\left(\frac{\mu\rho}{\mu\rho+1-\mu}\right). \tag{2.18}$$

2.8.2 The Scheme $(f_\varepsilon^K, S_\varepsilon^K)$

It remains to specify the fees and the distributions:

$$f_\varepsilon^K(C) = \frac{1-\rho^{K+1}}{1-\rho}\lambda_0^C(1-\rho-c) + \left(1 - \frac{1-\rho^{K+1}}{1-\rho}\lambda_0^C - \lambda_N^C\right)(p_{\underline{\mu}^K} - c) + \lambda_N^C\rho - \rho,$$

$$f_\varepsilon^K(I) = (1-\lambda_N^I)p_{\underline{\mu}^K} + \lambda_N^C\rho - \rho,$$

$$\lambda_n^C = \left(\frac{\rho}{\gamma^K(1-\rho)}\right)^n \left(\frac{\gamma^K(1-\rho)-\rho}{\gamma^K(1-\rho) - \rho(\frac{\rho}{\gamma^K(1-\rho)})^N}\right), \quad n = 0, \ldots, N$$

$$\lambda_n^I = \left(\frac{\rho}{\gamma^K(1-\rho)}\right)^{n-1} \left(\frac{\gamma^K(1-\rho)-\rho}{\gamma^K(1-\rho)(1-(\frac{\rho}{\gamma^K(1-\rho)})^N)}\right), \quad n = 1, \ldots, N.$$

2.9 Appendix C: Supplementary Details

This supplement is organized as follows. First, I provide a further discussion on the construction of the virtually FSO family of schemes in the main text. I complement the discussion with four numerical examples in the case of $K = 0$. I then consider the case of $K = \infty$, in which consumers observe an infinite history of past ratings. In this case, I show that full-surplus optimality fails (even virtually). Finally, I construct a FSO scheme that is independent of the firm's type report.

2.9.1 Virtual Full-Surplus Optimality

In this section, I discuss the economic forces behind the construction of the virtually FSO family of schemes in Proposition 2.2.

Moral hazard and rating structure. To understand the use of many stationary Markov ratings in Proposition 2.2, Proposition 2.4 below derives a cost upper bound given (μ, ρ, δ) above which efficiency fails when schemes are restricted to be stationary Markov.

Proposition 2.4. *Fix a scheme (f, S) that is stationary Markov given a participation rule π and a collection of equilibria (σ, φ). Let n be the number of ratings that are displayed with positive probability. If (σ, φ) is efficient, then*

$$c \leq \delta(1-2\rho)^2\left(1 - \frac{\mu\rho^{n-1}}{1-\mu+\mu\rho^{n-1}}\right) < \delta(1-2\rho)^2 = \bar{c}.$$

Finiteness and a stationary Markov structure restrict the extent to which consumers' beliefs, and hence payments and a firm's continuation profit upon a bad outcome can fall. They in turn restrict the range of cost over which efficiency can be incentivized in equilibrium.

When $n = 2$, the cost bound in Proposition 2.4 coincides with $\bar{c}_2$, which is defined in (2.8). Thus, the cutoff $\bar{c}_2$ is tight in the sense that there is a stationary Markov FSO scheme that display two ratings if and only if $c \leq \bar{c}_2$. For $n > 2$, the problem becomes less tractable and the cost bound in Proposition 2.4 is not tight in general. In Section 2.9.2, I present four numerical examples in the case of $K = 0$ to illustrate several insights. I show that the tight cost upper bound $\bar{c}_n$ (given μ, ρ and δ) for a stationary Markov FSO scheme with n ratings to exist equals the value of a minimization problem over two $n \times n$ probability transition matrices upon a competent report and a n-dimensional probability vector upon an inept report. The solutions to the problem, namely, the two matrices and the vector, are nontrivial functions of (μ, ρ, δ). Moreover, fixing (μ, ρ, δ), $\bar{c}_n$ need not strictly increase in n. By using more stationary Markov ratings, the certifier can induce lower payments for punishment and create stronger effort incentives, but he also introduces new incentive constraints for high effort upon the additional ratings. Nonetheless, in Section 2.9.2, I also show that for several parameter values (μ, ρ), $\bar{c}_n$ strictly increases when moving from $n = 2$ to 3, and when the incentive constraints are evaluated in the limit $\delta \to 1$. Such monotonicity due to patience is a virtue and motivates my focus on virtual full-surplus optimality for high effort costs.

The role of patience. To further illustrate the role that patience plays for efficiency, I now turn to a negative result. I show that, for any fixed (μ, ρ, δ), there exist a costs large enough (but strictly smaller than $\bar{c}$) such that efficiency fails in equilibrium:

Proposition 2.5. *Fix (μ, ρ, δ). There is $\eta \equiv \eta(\mu, \rho, \delta) > 0$ such that for some $\hat{c}(\eta) < \bar{c}$, where $\bar{c}$ is defined in* (2.7)*, the certifier's payoff is less than $\mu(1 - 2\rho - c) - \eta > 0$ whenever $c \geq \hat{c}(\eta)$.*

To motivate consistent high effort by a competent firm when c is arbitrarily close to $\bar{c}$, punishment upon a bad outcome after any history must be made as harsh as possible. Because of discounting, such harsh punishments include the use of ratings that induce sufficiently low consumers' beliefs that the firm is competent soon enough after a bad outcome. But if these ratings are realized soon enough upon a bad outcome after any history by the competent type, consumers' beliefs that the firm is competent upon such ratings cannot be too small, yielding a contradiction.

The incentive constraints for high effort can be relaxed by introducing a "bad" rating that induces zero effort by the competent firm, and hence a payment ρ, as in the construction in Proposition 2.2. But Proposition 2.5 says that the certifier cannot obtain a payoff arbitrarily close to $\bar{W}$ when c is sufficiently close to $\bar{c}$. This would require that the probability of receiving the bad rating be arbitrarily close to zero, in which case this bad rating would have a negligible effect on the competent firm's effort incentives.

2.9.2 Numerical Examples

In the examples, I focus on the competent firm's incentive constraints for high effort upon accepting and reporting truthfully. Given the certifier's ability to screen, it is then straightforward to choose a participation rule and fees to incentivize participation by both types and fully extract a participating firm, obtaining a FSO scheme as in Propositions 2.1 and 2.2.

Finally, note that in the case of $K = 0$, it is without loss to assume that the system upon an inept report specifies only absorbing states when constructing stationary Markov schemes. This is because, upon screening a firms type via her report in a FSO scheme, the rating transitions after an inept report matter only to the extent that the corresponding stationary distribution affects consumers beliefs on the firms type.

Three Stationary Markov Ratings

Example 2 illustrates that given (μ, ρ, δ), the certifier may achieve full-surplus optimality for a range of costs beyond $\bar{c}_2$.

Example 2 (Three ladder ratings)**.** Fix $(\mu, \rho, \delta) = (\frac{9}{10}, \frac{1}{10}, \frac{9}{10})$, so that $\bar{c}_2 \approx 0.303$. There are three ratings, labeled r_0, r_1 and r_2. Upon a competent report, rating transitions are depicted below in Figure 2.4. The initial distribution on $\{r_0, r_1, r_2\}$ is chosen to be the stationary distribution $(\frac{81}{91}, \frac{9}{91}, \frac{1}{91})$.

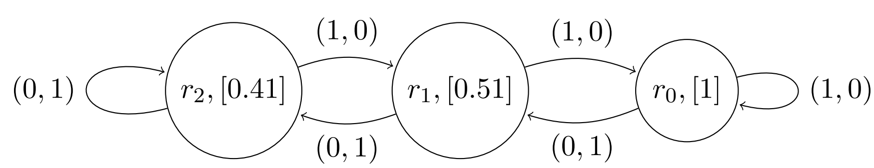

Figure 2.4: Transitions upon a competent report

Each node is labeled by the rating, followed by the corresponding belief in squared brackets. Transition probabilities are written in the form (x, y) and rounded up to two decimal places, where x denotes the probability upon a good outcome and y denotes the probability upon a bad outcome.

Transitions upon an inept report is depicted by Figure 2.5 below. The initial distribution over the ratings $\{r_0, r_1, r_2\}$ is $(0, 0.86, 0.14)$. The ratings are absorbing.

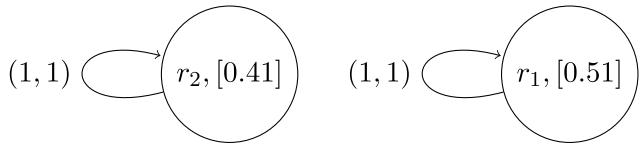

Figure 2.5: Transitions upon an Inept Report

The continuation profit V_n upon each rating r_n by the competent firm in an efficient equilibrium satisfies the following Bellman equation:

$$V_n = (1-\delta)(p(r_n)-c) + \delta((1-\rho)V_{\max(n-1,0)} + \rho V_{\min(n+1,2)}), \quad n = 0, 1, 2.$$

The incentive constraint for high effort upon each rating r_n is

$$\frac{\delta}{1-\delta}(1-2\rho)(V_{\max(n-1,0)} - V_{\min(n+1,2)}) \geq c.$$

The maximum cost below which consistent high effort holds in equilibrium is therefore

$$\frac{\delta(1-2\rho)}{1-\delta}\min(V_0 - V_1, V_0 - V_2, V_1 - V_2) \approx 0.31 > \bar{c}_2.$$

♦

To see why more ratings may relax the cost threshold, it is instructive to compare the transitions in Figure 2.4 to those upon a competent report in the scheme (f_2, S_2), depicted in Figure 2.6 below.

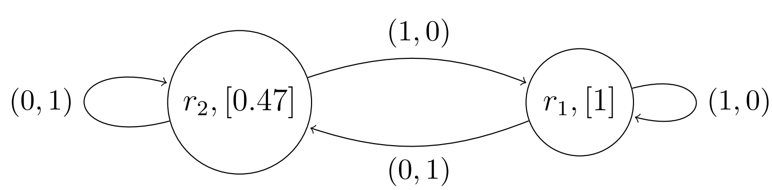

Figure 2.6: Transitions in the system S_2 upon a competent report

The use of more ratings allows the certifier to "push" consumers' beliefs towards

0 by statistically revealing more bad outcomes in a firm's track record, lowering consumer payments associated with "low" ratings to create stronger effort incentives. The downside is the inevitable introduction of additional incentive constraints for high effort upon observing the additional ratings. Ratings that induce intermediate beliefs may adversely create a "cushion" that weakens effort incentives. Of course, the certifier can always use a smaller amount of ratings than the available ones by not allowing some ratings to be realized with positive probability.

Interestingly, the certifier can achieve full-surplus optimality for an even larger range of costs if he dispenses with the ladder structure. Example 3 below illustrates.

Example 3 (Three ratings with mixed transitions)**.** Departing from Example 2, consider instead the following transitions upon a competent report depicted in Figure 2.7 below, with the initial distribution being the stationary distribution $(0.79, 0.2, 0.01)$. Transitions upon an inept report are depicted in Figure 2.8 below, with the initial distribution being $(0, 0.76, 0.24)$.

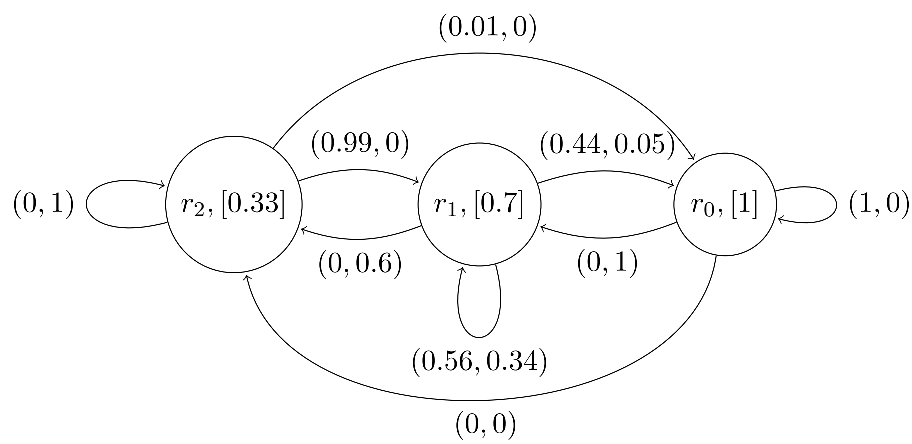

Figure 2.7: Transitions upon competent report

Let α_{mn} denote the transition probability from rating r_m to rating r_n upon a good outcome in the competent system, and let β_{mn} denote the counterpart upon a bad

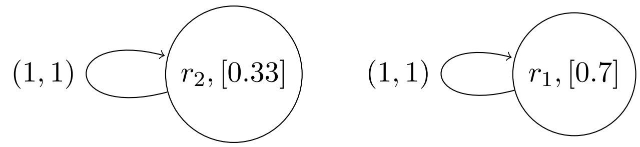

Figure 2.8: Transitions upon inept report

outcome. The competent firm's continuation profit V_n upon each rating r_n satisfies

$$V_m = (1-\delta)(p(m)-c) + \delta \sum_{n=0}^{2}((1-\rho)\alpha_{mn} + \rho\beta_{mn})V_n, \quad m, n = 0, 1, 2.$$

Using the fact that $\alpha_{m2} = 1 - \alpha_{m0} - \alpha_{m1}$ and $\beta_{m2} = 1 - \beta_{m0} - \beta_{m1}$, the incentive constraints for high effort by the competent firm upon rating r_m is

$$\frac{\delta(1-2\rho)}{1-\delta}\left[\sum_{n=0,1}(\alpha_{mn} - \beta_{mn})(V_0 - V_n)\right] \geq c.$$

The maximum cost below which consistent high effort holds in equilibrium is therefore

$$\min_{m\in\{0,1,2\}} \frac{\delta(1-2\rho)}{1-\delta}\left[\sum_{n=0,1}(\alpha_{mn} - \beta_{mn})(V_0 - V_n)\right] \approx 0.35 > 0.31.$$

♦

Example 4 shows, however, three stationary and Markov ratings need not improve upon two.

Example 4 (Three ratings with mixed transitions and low prior)**.** Departing from Example 3, keep $\delta = \frac{9}{10}$ and $\rho = \frac{1}{10}$, but suppose that $\mu = \frac{1}{100}$. This gives $\bar{c}_2 \approx 0.575$. Choosing new transition probabilities $(\alpha_{mn}, \beta_{mn})_{m,n=0,1,2}$ and new initial distribution λ^I over inept ratings (and fixing the initial distribution λ^C over competent ratings to be the stationary distribution given the new transitions), the maximum cost below which consistent high effort holds in equilibrium is

$$\max_{(\alpha_{mn},\beta_{mn})_{m,n=0,1,2},(\lambda^I_m)_{m=0,1}} \min_{m\in\{0,1,2\}} \frac{\delta(1-2\rho)}{1-\delta}\left[\sum_{n=0,1}(\alpha_{mn}-\beta_{mn})(V_0-V_n)\right] \approx 0.572,$$

which does not improve over $\bar{c}_2$. ♦

The key force behind why no improvement obtains is the low prior μ. Beliefs and hence prices induced by the two pooling ratings are too close, and a firm finds it suboptimal to exert costly effort for a negligible change in consumer payments.

Nonetheless, I show numerically in Example 5 below that for many parameter values (μ,ρ), an improvement obtains when moving from two to three stationary Markov ratings, if the incentive constraints are evaluated at the patient limit $\delta \to 1$. An arbitrarily patient competent firm still wishes to exert effort to achieve a rating that induces a higher consumer's belief, despite this increase in belief may be small.

Example 5 (Cost cutoff and monotonicity at the patient limit)**.** The above examples reveal that in general, the maximum cost below which efficiency obtains in equilibrium with n ratings is given by

$$\bar{c}_n(\mu,\rho,\delta) := \max_{(\alpha_{ms},\beta_{ms})^{n-1}_{m,s=0},(\lambda^I_m)^{n-1}_{m=0}} \min_{m=0,\ldots,n-1} \frac{\delta(1-2\rho)}{1-\delta}\left[\sum_{s=0}^{n-1}(\alpha_{ms}-\beta_{ms})(V_0-V_s)\right]. \tag{2.19}$$

When the incentive constraints for high effort are evaluated at the patient limit $\delta \to 1$, the limit cost cutoff is given by

$$\bar{c}^*_n(\mu,\rho) := \max_{(\alpha_{ms},\beta_{ms})^{n-1}_{m,s=0},(\lambda^I_m)^{n-1}_{m=0}} \min_{m=0,\ldots,n-1}(1-2\rho)\left[\sum_{s=0}^{n-1}(\alpha_{ms}-\beta_{ms})\left(\lim_{\delta\to 1}\frac{V_0-V_s}{1-\delta}\right)\right]. \tag{2.20}$$

Table 2.1 below computes numerically the cutoff (2.20) at $(\mu,\rho) \in \{0.1,0.3,0.5,0.7,0.9\} \times \{0.1,0.2,0.3,0.4\}$. The cutoff (2.20) increases as the certifier moves from two to three stationary Markov ratings.[27] It suggests that as $n \to \infty$, the cutoff (2.20) approaches $(1-2\rho)^2 \min[\frac{1-\mu}{\rho}, 1]$, which is the bound on effort cost identified in Proposition 2.2. ♦

		ρ			
		0.1	0.2	0.3	0.4
		0.63	0.35	0.15	0.038
μ	0.1	0.63	0.36	0.16	0.038
		0.64	0.36	0.16	0.04
		0.61	0.33	0.14	0.034
	0.3	0.63	0.35	0.14	0.034
		0.64	0.36	0.16	0.04
		0.58	0.3	0.12	0.029
	0.5	0.6	0.33	0.14	0.029
		0.64	0.36	0.16	0.040
		0.52	0.25	0.094	0.021
	0.7	0.6	0.29	0.11	0.023
		0.64	0.36	0.16	0.03
		0.34	0.13	0.043	0.0087
	0.9	0.43	0.14	0.047	0.0087
		0.64	0.18	0.053	0.01

Table 2.1: Limiting cost cutoffs (2.20)

Each cell of the table consists of, from top to bottom, the limit cutoff $\bar{c}_n^*$ for $n = 2, 3$ and the limit upper bound $(1-2\rho)^2 \min[\frac{1-\mu}{\rho}, 1]$, rounded to two decimal places.

Monotonicity due to patience is a virtue. As the first three examples illustrate, the solution to (2.19) is in general a non-trivial function of (μ, ρ, δ): a ladder structure does not suffice, and requires solving for non-trivial transition probabilities $(\alpha_{ms}, \beta_{ms})_{m,s=0}^{n-1}$ as well as the inept stationary distribution. In fact, these observations motivate the constructions in the main result of the paper, Proposition 2.2.

27. The details involved in evaluating the unnormalized difference in continuation profits in (2.20) are contained in the proof of Proposition 2.2.

2.9.3 Infinite Past Ratings Observation

In this section, I show that full-surplus optimality fails (even virtually) if consumers observe the current ratings as well as all past ratings, *i.e.*, $K = \infty$.

Proposition 2.6. *Fix (μ, ρ, c, K) such that $c < (1-2\rho)^2$ and $K \leq \infty$. Suppose that for $\varepsilon > 0$, there is a scheme given which, there is a sequence $\{\delta_m\}_{m=0}^{\infty}$ converging to 1 such that for every $m \geq M_\varepsilon$ for some $M_\varepsilon > 0$, the certifier's optimal payoff is at least $\bar{W} - \varepsilon$ given a discount factor δ_m. Then $K < \infty$.*

For the certifier to virtually obtain $\bar{W}$, the competent firm must virtually consistently exert high effort on path. Given $K = \infty$, consumers' beliefs that the firm is competent tend to drift up and converge to 1. The competent firm's payoff after any history on path in the patient limit must then equal $1 - \rho - c$. However, once consumers' beliefs become close to one, to motivate high effort at any given history, the competent firm must face a positive-probability punishment phase in which she receives low payments, and such punishment phase must not vanish in the patient limit. The competent firm's continuation payoff after any history on path in the patient limit must therefore be strictly smaller than $1 - \rho - c$. This yields a contradiction.

The solution to the certifier's problem when $K = \infty$ thus requires the certifier to incentivize high effort by the competent firm after specific histories and shirking after other histories. Solving for the certifier's optimal payoff might become intractable, but the fundamental insight that the certifier wishes to maximize a firm's expected value to consumers via coarse and opaque ratings remains.

2.9.4 A Type-Independent, Full-Surplus-Optimal Scheme

In this section, I construct a type-independent scheme (f_2', S_2') that is FSO whenever

$$c \leq \vec{c}_2'(\mu, \rho, \delta) = \frac{\delta(1-\mu)\mu(1-2\rho)^3}{(1-\rho-\mu(1-2\rho))(\rho+\mu(1-2\rho))}. \tag{2.21}$$

The system S_2' displays a rating r_0 with probability 1 upon a good outcome with some probability α, and a rating r_1 otherwise in each period. The probability α is

$$\alpha := \frac{c(\rho+\mu(1-2\rho))}{c(\rho+\mu(1-2\rho))^2 + \delta(1-\mu)\mu(1-2\rho)^3}, \tag{2.22}$$

which belongs to $(0,1]$ when Condition (2.21) below holds. The fee f_2' is

$$f_2' = (1-2\rho)\left[\mu - \frac{\delta\alpha(1-\mu)\varphi^2(1-2\rho)^2}{(\rho+\mu(1-2\rho))(1-\alpha(\rho+\mu(1-2\rho)))}\right]. \tag{2.23}$$

The participation rule specifies that both types accept and report truthfully. The competent firm always exerts high effort in the induced game $G(f_2', S_2', C)$, and always shirks in other induced games. The participation fee $f_2' := f_2'(C) = f_2'(I)$ is set to keep a participating firm at her rejection payoff ρ.

Proposition 2.7. *The scheme* (f_2', S_2') *is FSO whenever* (2.21) *holds.*

Rating r_1 in period 0 does not reflect any past outcome, thus $\varphi_0(r_1) = \mu$. In each period $t \geq 1$, the two ratings statistically reveal the outcome in the last period, inducing two possible levels of reputations:

$$\varphi_t(0) = \frac{\mu(1-\alpha(1-\rho))}{\mu(1-\alpha(1-\rho)) + (1-\mu)(1-\alpha\rho)} < \varphi_t(1) = \frac{\mu(1-\rho)}{\mu(1-\rho)+(1-\mu)\rho}.$$

Consistent high effort again implies that $p_t(1) > p_t(0)$ for each $t \geq 1$. When (2.21) holds, the competent firm consistently exerts high effort for the higher payment $p_t(1)$.

Note that $\bar{c}_2' < \bar{c}_2$.[28] This is because the two ratings in this construction pool both types of firms, making it more difficult to motivate consistent high effort by the competent firm.

Notice also that the scheme must induce a competent firm to be indifferent between high or low effort after every history upon acceptance (and must induce both types to accept). The reason is that, for full extraction, a fee independent of the report must bind both types individual rationality constraints, keeping them at the rejection payoff. This requires both types to extract identical expected profits from consumers. However, efficiency requires that the competent firm be motivated to secure higher profits by separating herself from an inept type.

2.10 Appendix D: Proofs

2.10.1 Proof of Lemma 2.1

It remains to show the converse and that (2.7) is necessary for efficiency. I first show the converse. Fix a scheme and a participation rule that solve the certifier's problem. Since a firm truthfully reports upon acceptance given incentive compatibility, efficiency implies that the participation rule satisfies full competent acceptance. By (2.6), the certifier achieves $\bar{W}$.

I now show that (2.7) is necessary for efficiency. Recall from the main text that efficiency implies full competent acceptance. Thus, given efficiency, the incentive constraint for high effort after each history h must hold in the induced game upon a competent report. By exerting high effort, the competent firm's continuation profit

28. To see this, note that the inequality can be simplified to $(1-\rho)/(1-2\rho) > (2-\mu)\mu$. The left side is strictly larger than 1 for $\rho \in (0, \frac{1}{2})$. The right side is bounded above by 1.

following a history h in period t in which a rating r realizes is

$$V = (1-\delta)(p_t(r) - c) + \delta((1-\rho)V(h\bar{y}) + \rho V(h\underline{y})),$$

where $V(hy)$ denotes the expected continuation profit upon producing an outcome y after history h. It suffices to verify that it is not profitable for the firm to deviate to zero effort, since cost is linear in effort. A one-stage deviation to zero effort gives

$$V' = (1-\delta)p_t(r) + \delta(\rho V(h\bar{y}) + (1-\rho)V(h\underline{y})).$$

The incentive constraint, $V - V' \geq 0$, can therefore be simplified as

$$V(h\bar{y}) - V'(h\underline{y}) \geq \frac{(1-\delta)c}{\delta(1-2\rho)}. \tag{2.24}$$

Recursively using (2.24),

$$\begin{aligned} &V(h\bar{y}) \leq (1-\delta)(1-\rho-c) + \delta((1-\rho)V(h\bar{y}\bar{y}) + \rho V(h\bar{y}\underline{y})) \\ &\leq (1-\delta)(1-\rho-c) + \delta\left(V(h\bar{y}\bar{y}) - \frac{(1-\delta)\rho c}{\delta(1-2\rho)}\right) \leq 1-\rho-c-\frac{\rho c}{1-2\rho}. \end{aligned} \tag{2.25}$$

Similarly,

$$\begin{aligned} &V(h\underline{y}) > (1-\delta)(\rho-c) + \delta((1-\rho)V(h\underline{y}\bar{y}) + \rho V(h\underline{y}\underline{y})) \\ &\geq (1-\delta)(\rho-c) + \delta\left(V(h\underline{y}\underline{y}) + \frac{(1-\delta)(1-\rho)c}{\delta(1-2\rho)}\right) > \cdots \geq \rho - c + \frac{(1-\rho)c}{1-2\rho}. \end{aligned} \tag{2.26}$$

This inequality is strict because payments, given in (2.2), are always strictly bigger than ρ given consistent high effort. (2.24), (2.25) and (2.26) together imply that $c < \delta(1-2\rho)^2$.

2.10.2 Proof of Proposition 2.1

Given the fees and the participation rule, it is straightforward to compute that the certifier's expected payoff equals $\bar{W}$. I show that (σ^*, φ^*) is a collection of equilibria when (2.8) holds. Consider the game $G(f_2, S_2, C)$. The competent firm's play in $G(f_2, S_2, C)$ can be represented by a two-state automaton, with the set of states given by the set of ratings. Here, a state r collects all firm's histories with the most recent rating being r. Let $V(r)$ be the competent firm's continuation profit in state r:

$$V(r) = (1-\delta)(p(r) - c) + \delta[(1-\rho)V(r_0) + \rho V(r_1)],$$

where, by (2.2),

$$p(r_1) = \rho + \frac{(1-2\rho)\mu\rho}{\mu\rho + 1 - \mu} \quad \text{and} \quad p(r_0) = 1 - \rho.$$

A deviation to shirk yields

$$V(r; \underline{e}) := (1-\delta)p(r) + \delta[\rho V(r_0) + (1-\rho)V(r_1)].$$

The incentive constraint for high effort $V(r) - V(r; \underline{e}) \geq 0$ in each state r can be simplified to $\delta(1-2\rho)(p(r_0) - p(r_1)) \geq c$, or equivalently, (2.8).

The competent type finds no profitable deviation from the participation rule π^*. By accepting and reporting truthfully, the competent firm's payoff equals ρ by construction. By deviating to misreport upon acceptance, consumers pay the firm $p(r_1)$ in each period because only rating r_1 is displayed in the induced game $G(f_2, S_2, I)$, and the competent firm's payoff equals ρ. By deviating to reject, the firm's payoff equals ρ. Finally, the inept type also finds no profitable deviation from the participation rule π^*. By accepting and reporting truthfully, the inept firm's payoff equals ρ.

By accepting and misreporting, the inept firm's payoff is at most ρ by construction, whenever $c \leq \bar{c}_2$. Finally, by deviating to reject, the inept firm's payoff equals ρ.

2.10.3 Proof of Proposition 2.2

Necessity of Condition (2.12). Fix $\varepsilon > 0$ and a scheme given which there is a sequence $\{\delta_m\}_{m=0}^{\infty}$ converging to 1 such that for every $m \geq M_\varepsilon$ for some $M_\varepsilon > 0$, the certifier's payoff is at least $\bar{W} - \varepsilon$ given a discount factor δ_m for some participation rule and some collection of equilibria. Fix such a sequence $\{\delta_m\}_{m=0}^{\infty}$, M_ε and $m \geq M_\varepsilon$. Next, fix a sequence $\{x_m\}_{m \geq M_\varepsilon}$ where each $x_m \in (0,1)$ is chosen below to be sufficiently small, with the properties that $\liminf_m x_m =: x > 0$ and $\delta_m(1 - x_m) < \delta_{m+1}(1 - x_{m+1})$ for each m. I first argue that given any histories h, h' in the induced game upon a competent report, the competent firm's continuation payoffs $V(h), V(h')$ at the histories upon accepting and truthfully reporting must satisfy

$$\left| \frac{V(h) - V(h')}{1 - \delta_m} \right| < D := \frac{1 - \rho(4(1-\rho) + c)}{\rho x} \tag{2.27}$$

where $D > 0$ is an upper bound on the value of the following program:

$$D_m := \max_{\alpha \in [0,1]} \frac{V_0 - V_1}{1 - \delta_m}$$

subject to

$$V_0 = (1 - \delta_m)(1 - \rho - c) + \delta((1 - \rho\alpha)V_0 + \rho\alpha V_1), \tag{2.28}$$

$$V_1 = (1 - \delta_m)(\rho - c) + \delta_m(x_m V_0 + (1 - x_m)V_1), \tag{2.29}$$

$$\delta_m \alpha (1 - 2\rho)(V_0 - V_1) \geq (1 - \delta_m)c. \tag{2.30}$$

The program can be interpreted as follows. The values V_0, V_1 are upper and lower bounds on the competent firm's continuation payoff after any history in the induced game upon acceptance and truthfully reporting. There are two states, 0 and 1, associated with the continuation values V_0 and V_1. In state 0, the firm receives the stage payoff $1 - \rho - c$, as if she exerts high effort, subject to incentive constraint for high effort (2.30). In state 1, the firm receives a payoff $\rho - c$, which is strictly lower than her lowest stage payoff. Given state 0, with probability $1-\rho+\rho(1-\alpha) = 1-\rho\alpha$, the next state is 0 and with complementary probability, the next state is 1. This is as if, given high effort in state 0, the next state is 0 with probability one upon a good outcome, or with probability $1 - \alpha$ upon a bad outcome. Given state 1, with probability x_m, the next state is 0 and with complementary probability, the next state is 1. Since the certifier's optimal payoff is at least $\bar{W} - \varepsilon$ given a discount factor δ_m, to derive the upper and lower bounds on $V(h)$ after any history h, it suffices to consider to a transition probability $x_m > 0$, which is chosen to be sufficiently small and satisfy the properties mentioned above to ensure that V_1 is a lower bound on the competent firm's continuation payoff after any history in the induced game upon a competent report. Such a sequence $\{x_m\}_{m \geq M_\varepsilon}$ exists because $\lim_{x_m \to 0} V_1 = \rho - c$, which is strictly smaller than a competent firm's lowest equilibrium continuation profit ρ in the induced game upon a competent report.

At the optimum of the program, the constraint (2.30) binds, and the optimal value of α is such that given a current state 0, the probability of staying at state 0 is maximized without violating the constraint (2.30). Hence, at the optimum, V_0 (resp., V_1) gives an upper (resp., lower) bound on the competent firm's continuation payoff after any history: $V_1 < V(h) \leq V_0$ after any history h in the induced game upon a competent report. The lower bound is strict because the competent firm's payoff in each period is strictly above $\rho - c$. Thus, $\left|\frac{V(h)-V(h')}{1-\delta_m}\right| < D_m$. Solving (2.28), (2.29) and (2.30) simultaneously for V_0, V_1, α gives

$$D_m = \frac{1 - \rho(4(1-\rho) + c)}{(1-2\rho)(1-\delta_m(1-x_m))}.$$

This value strictly increases in m, and hence

$$D_m < \limsup_{m\to\infty} D_m = \frac{1-\rho(4(1-\rho)+c)}{\rho x}.$$

Now, fix an arbitrary $\psi > 0$. Given the bound (2.27), by Ross (2014, Theorem 2.2, p.95), there exists a bounded function $\ell(h)$ for each history h and a constant g satisfying

$$\begin{aligned} g + \ell(h_t) &= p(h_t) - c + (1-\rho)\mathbf{E}^{P_C^C}_{h_{t+1}}[\ell(h_{t+1})|h_t, \bar{y}] + \rho\mathbf{E}^{P_C^C}_{h_{t+1}}[\ell(h_{t+1})|h_t, \underline{y}], \\ g &= \mathbf{E}^{P_C^C}_{h}[p(h)] - c, \\ \ell(h) &= \lim_{m\to\infty} \frac{V(h) - V(h^*)}{1-\delta_m}, \end{aligned}$$

where $p(h)$ denotes consumer's payment at history h and h^* is chosen at which the competent firm exerts high effort and such that

$$0 \leq \sup_h \ell(h) - \ell(h^*) \leq \psi. \tag{2.31}$$

The incentive constraint for high effort at history h^* is

$$\delta_m(1-2\rho)\frac{(V(h^*\bar{y}) - V(h^*\underline{y}))}{1-\delta_m} \geq c.$$

Thus,

$$\begin{aligned} & g = p(h^*) - c + (1-\rho)\ell(h^*\bar{y}) + \rho\ell(h^*\underline{y}) \\ \iff \quad & \ell(h^*\bar{y}) - \ell(h^*\underline{y}) = \frac{p(h^*) - c - g + \ell(h^*\bar{y})}{\rho}. \end{aligned} \tag{2.32}$$

It follows from the incentive constraint that

$$\delta_m(1-2\rho)(\ell(h^*\bar{y})-\ell(h^*\underline{y})) = \delta_m(1-2\rho)\lim_{m\to\infty}\frac{V(h^*\bar{y})-V(h^*\underline{y})}{1-\delta_m} \geq c.$$

Simplifying,

$$\delta_m(1-2\rho)(\ell(h^*\bar{y})-\ell(h^*\underline{y})) \geq c \quad \iff \quad \delta_m\left(\frac{p(h^*)-c-g+\ell(h^*\bar{y})}{\rho}\right) \geq \frac{c}{1-2\rho}.$$

Simplifying the right inequality,

$$\delta_m\big((1-2\rho)(1-\mathbf{E}_h^{P_C^C}[\varphi(h)])+\ell(h^*\bar{y})\big) \geq \frac{\rho c}{1-2\rho},$$

where $\varphi(h)$ denotes consumers' belief that the firm is competent at history h in the induced game upon a competent report. By construction, $\ell(h^*\bar{y}) \leq \psi$ and $\mu \leq \mathbf{E}_h^{P_C^C}[\varphi(h)]$. The above inequality implies that

$$\delta_m\big((1-2\rho)(1-\mu)+\psi\big) \geq \frac{\rho c}{1-2\rho}.$$

Since $\psi > 0$ is chosen arbitrarily, it must hold that

$$\delta_m(1-2\rho)(1-\mu) \geq \frac{\rho c}{1-2\rho} \quad \iff \quad c \leq \frac{1-\mu}{\rho}\delta_m(1-2\rho)^2 < \frac{1-\mu}{\rho}(1-2\rho)^2.$$

This proves the necessity of Condition (2.12).

Verify that the construction constitutes a virtually FSO family of schemes. Fix some $N > K+1$, chosen below to be sufficiently large. State variables that govern a competent firm's continuation payoff are given by a vector (k, r), consisting of the current rating r and the number $k \geq 1$ of past periods since rating 0 is most recently realized, as these two numbers pin down consumers' beliefs and the firm's effort choices. The state transitions are as follows. Given a current state (k, r_0) for any k, the next state is $(1, r_0)$ upon a good outcome and is $(1, r_1)$ upon a bad outcome. Given a current state (k, r_n) for any k, and $n = 1, \dots, N-1$, upon a good outcome, the next state is $(k+1, r_{n-1})$ with probability γ^K and is $(k+1, r_n)$ with probability $1-\gamma^K$; upon a bad outcome, the next state is $(k+1, r_{n+1})$. Finally, given a current state (k, r_N) for any k, the next state is $(k+1, r_{N-1})$ with probability $\gamma^K(1-\rho)$ and is $(k+1, r_N)$ with complementary probability.

Let $V_{(k,n)}$ denote the continuation payoff of the competent firm given a state (k, r_n), the participation rule and the collection of equilibria. The continuation payoffs satisfy

$$\begin{aligned} V_{(k,0)} &=: V_0, \quad k = 1, \dots, K, \\ V_{(k,n)} &=: V_n, \quad k \geq K+1, n = 1, \dots, N, \end{aligned}$$

and the system of Bellman equations

$$\begin{aligned} V_0 &= (1-\delta)(1-\rho-c) + \delta((1-\rho)V_0 + \rho V_{(1,1)}), \\ V_{(k,n)} &= (1-\delta)(1-\rho-c) + \delta((1-\rho)\gamma^K V_{(k+1,n-1)} \\ &\quad + (1-\rho)(1-\gamma^K)V_{(k+1,n)} + \rho V_{(k+1,n+1)}), \quad k = 1, \dots, K, n = 1, \dots, k, \\ V_n &= (1-\delta)(p_{\underline{\mu}^K} - c) + \delta((1-\rho)\gamma^K V_{n-1} \\ &\quad + (1-\rho)(1-\gamma^K)V_n + \rho V_{n+1}), \quad n = 1, \dots, N-1, \\ V_N &= (1-\delta)\rho + ((1-\rho)\gamma^K V_{N-1} + (1-(1-\rho)\gamma^K)V_N). \end{aligned}$$

This linear system of $K(K+1)/2+N+1$ equations has $K(K+1)/2+N+1$ unknowns, namely, the continuation payoffs, and admits a unique solution.

It suffices to verify that the incentive constraints for high effort in the patient limit hold strictly, when the number N is (fixed and) chosen to be sufficiently large:

$$\lim_{\delta\to 1}\delta(1-2\rho)\left(\frac{V_0-V_{(1,1)}}{1-\delta}\right)>c, \tag{2.33}$$

$$\lim_{\delta\to 1}\delta(1-2\rho)\left(\frac{\gamma^K V_{(k+1,n-1)}+(1-\gamma^K)V_{(k+1,n)}-V_{(k+1,n+1)}}{1-\delta}\right)>c, \quad k=1,\dots,K, n=1,\dots,k, \tag{2.34}$$

$$\lim_{\delta\to 1}\delta(1-2\rho)\left(\frac{\gamma^K V_{n-1}+(1-\gamma^K)V_n-V_{n+1}}{1-\delta}\right)>c, \quad n=1,\dots,N-1. \tag{2.35}$$

Since the strategy profile σ_ε^K induces an irreducible Markov chain over a finite set of $K(K+1)/2+N+1$ states upon a competent report, Ross (2014, Theorem 2.4) shows that the absolute unnormalized difference of continuation profits $\left|\frac{V_{(k,n)}-V_0}{1-\delta}\right|$ is uniformly bounded over k,n and $\delta\in(0,1)$. Moreover, by Ross (2014, Theorem 2.2), the following holds. First, there exists bounded $\ell^N_{(k,n)}$ and g^N such that

$$\ell^N_{(k,n)}=\lim_{\delta\to 1}\frac{V_{(k,n)}-V_0}{1-\delta},$$

$$\ell^N_{(k,n)}:=\ell^N_n, \quad k\geq K+1, n=1,\dots,N$$

$$g^N=\lim_{\delta\to 1}V_0=\frac{1-\rho^{K+1}}{1-\rho}\lambda_0^C(1-\rho-c)+\left(1-\frac{1-\rho^{K+1}}{1-\rho}\lambda_0^C-\lambda_N^C\right)(p_{\underline{\mu}^K}-c)+\lambda_N^C\rho.$$

Second, the following system

$$g^N=1-\rho-c+\rho\ell^N_{(1,1)},$$

$$g^N+\ell^N_{(k,n)}=1-\rho-c+(1-\rho)\gamma^K\ell^N_{(k+1,n-1)}$$
$$+(1-\rho)(1-\gamma^K)\ell^N_{(k+1,n)}+\rho\ell^N_{(k+1,n+1)}, \quad k=1,\dots,K, n=1,\dots,k.$$

$$g^N+\ell^N_n=p_{\underline{\mu}^K}-c+(1-\rho)\gamma^K\ell^N_{n-1}+(1-\rho)(1-\gamma^K)\ell^N_n+\rho\ell^N_{n+1}, \quad n=1,\dots,N-1,$$

$$g^N+\ell^N_N=\rho+(1-\rho)\gamma^K\ell^N_{N-1}+(1-(1-\rho)\gamma^K)\ell^N_N,$$

holds and admits a unique solution $(((\ell^N_{(n,k)})^k_{n=1})^K_{k=1}, (\ell_n)^N_{n=1})$. Verifying (2.33)–(2.35) is then equivalent to verifying

$$-\ell^N_{(1,1)} > \frac{c}{1-2\rho}, \tag{2.36}$$

$$\gamma^K \ell^N_{(k,n-1)} + (1-\gamma^K)\ell^N_{(k,n)} - \ell^N_{(k,n+1)} > \frac{c}{1-2\rho} \quad k = 1, \ldots, K, n = 1, \ldots, k, \tag{2.37}$$

$$\gamma^K \ell^N_{n-1} + (1-\gamma^K)\ell^N_{n} - \ell^N_{n+1} > \frac{c}{1-2\rho}, \quad n = 1, \ldots, N-1. \tag{2.38}$$

Solving the above system for $(((\ell^N_{(n,k)})^k_{n=1})^K_{k=1}, (\ell_n)^N_{n=1})$, it holds that

$$\begin{aligned}\lim_{N\to\infty} -\ell^N_{(1,1)} &= \begin{cases} \frac{(1-\mu)(1-2\rho)(\rho+\rho^K-2\rho^{K+1})}{(1-\rho)(1-\rho-\mu(1-2\rho))}, & \text{if } \hat{c} < \frac{(1-\mu)(1-2\rho)^2(\rho+\rho^K(1-2\rho))}{(1-\rho)(1-\mu-\rho+2\mu\rho)}, \\ \frac{\hat{c}}{1-2\rho}, & \text{otherwise.} \end{cases} \\ &> \frac{c}{1-2\rho}.\end{aligned}$$

It also holds that

$$\begin{aligned} &\lim_{N\to\infty} \left[\gamma^K \ell^N_{n-1} + (1-\gamma^K)\ell^N_n - \ell^N_{n+1}\right] \\ &= \lim_{N\to\infty} \left[\gamma^K \ell^N_{(k,n-1)} + (1-\gamma^K)\ell^N_{(k,n)} - \ell^N_{(k,n+1)}\right], \quad k = 1, \ldots, K, n = 1, \ldots, k, \\ &= \begin{cases} \frac{(1-\mu)(1-2\rho)\zeta(K,n)}{(1-\rho)(1-\mu-\rho+2\mu\rho)}, & \text{if } \hat{c} < \frac{(1-\mu)(1-2\rho)^2(\rho+\rho^K(1-2\rho))}{(1-\rho)(1-\mu-\rho+2\mu\rho)}, \\ \frac{\zeta_1'(K)+\zeta_2'(K)\zeta_3'(K,n)}{(1-\rho)^2(1-2\rho)(1-\mu+\mu\rho-\rho^{K+1})}, & \text{otherwise.} \end{cases} \\ &> \frac{c}{1-2\rho}, \end{aligned}$$

where

$$\zeta(K,n) := 2 + (K(1-\rho)-1)\left(-1+\frac{1}{r}\right)^n + \left(\left(-1+\frac{1}{r}\right)^n - 2\rho\right)\rho^K,$$

$$\zeta_1'(K) := (1-\rho^{K+1})(\hat{c}(1-\rho)(1-\mu-\rho) + (1-\mu)(1-2\rho)^2(2-\rho-\rho^K)),$$

$$\zeta_2'(K) := (1-\mu)\rho(1-K+K\rho-\rho^K)(\hat{c}(1-\rho) + (1-2\rho)^2(1-\rho^K)),$$

$$\zeta_3'(K,n) := \left(\frac{(1-2\rho)^2(1-\rho^{K+1}) - \mu(1-\rho((4-\hat{c})(1-\rho)+(1-2\rho)^2\rho^K))}{(1-\mu)\rho(\hat{c}(1-\rho)+(1-2\rho)^2(1-\rho^K))}\right)^{n+1}.$$

Because the incentive constraints hold strictly in the limit $N \to \infty$, there exists N sufficiently larges such that (2.36)–(2.38) hold.

Because the competent firm only shirks upon rating r_N, the loss from the payoff upper bound $\mu(1-2\rho-c)$ occurs with probability

$$\mu\lambda_N^C + (1-\mu)\lambda_N^I = \mu\left(\frac{1-\mu-\rho}{(1-\mu-\mu\rho)(\frac{1-\mu-\mu\rho}{(1-\mu)\rho})^N - \rho + \mu\rho}\right) + (1-\mu)\left(\frac{1-\mu-\rho}{(1-\mu)\rho((\frac{1-\mu-\mu\rho}{(1-\mu)\rho})^N - 1)}\right),$$

which strictly decreases in N and vanishes as $N \to \infty$. This completes the proof.

2.10.4 Proof of Proposition 2.3

It is analogous to construct the fees such that each type of firm optimally accepts and reports truthfully given the collection of equilibria.

Here, I focus on verifying the competent firm's incentives to exert high effort upon accepting and truthfully reporting. Let $(r_{n'}, r_n)$ denote a state, where $r_{n'}$ is the most recent past rating and r_n is the current rating. Let $V_{(n',n)}$ be the competent type's continuation payoff upon a state $(r_{n'}, r_n)$ in the induced game $G(f_2', S_2, C)$ upon accepting and truthfully reporting given the collection of equilibria. The continuation payoffs solve the system of Bellman equations:

$$
\begin{aligned}
V_{(1,1)} &= (1-\delta)(1-\rho-c) + \delta[(1-\rho)V_{(1,1)} + \rho V_{(1,0)}] \\
V_{(0,1)} &= (1-\delta)(1-\rho-c) + \delta[(1-\rho)V_{(1,1)} + \rho V_{(1,0)}] \\
V_{(1,0)} &= (1-\delta)(1-\rho-c) + \delta[(1-\rho)V_{(0,1)} + \rho V_{(0,0)}] \\
V_{(0,0)} &= (1-\delta)(\rho + \tilde{\mu}_1(1-2\rho) - c) + \delta[(1-\rho)V_{(0,1)} + \rho V_{(0,0)}].
\end{aligned}
$$

The incentive constraints to exert high effort at each state are

$$
\begin{aligned}
\delta(1-2\rho)\frac{V_{(1,1)} - V_{(1,0)}}{1-\delta} \geq c, \quad &\iff \quad \frac{\delta^2(1-\mu)\rho(1-2\rho)^2}{1-\mu(1-\rho^2)} \geq c \\
\delta(1-2\rho)\frac{V_{(0,1)} - V_{(0,0)}}{1-\delta} \geq c. \quad &\iff \quad \frac{\delta(1-\mu)(1+\delta\rho)(1-2\rho)^2}{1-\mu(1-\rho)^2} \geq c.
\end{aligned}
$$

These two constraints hold simultaneously if and only if $c \leq \bar{c}_2'$, where $\bar{c}_2'$ is defined in (2.14) in the main text.

2.10.5 Proof of Proposition 2.4

The proof is identical to that in Lemma 2.1 in the main text, except that here the competent firm's continuation profit after a history h upon a bad outcome satisfies

$$
V(h\underline{y}) \geq \rho + (1-2\rho)\left(\frac{\mu\rho^n}{\mu\rho^n + 1 - \mu}\right) - c + \frac{(1-\rho)c}{1-2\rho},
$$

since

$$
\rho + (1-2\rho)\left(\frac{\mu\rho^n}{\mu\rho^n + 1 - \mu}\right)
$$

is the lowest equilibrium payment given efficiency.

2.10.6 Proof of Proposition 2.5

Suppose, towards a contradiction, that for every $\eta > 0$ and $c < \bar{c}$, the certifier's optimal payoff is at least $\mu(1-2\rho-c)-\eta$. The incentive constraint (2.24) necessarily holds after some history h_t in the induced game upon a competent report given some collection of equilibria. Following the proof of Lemma 2.1,

$$V(h_t\bar{y}) \leq 1 - \rho - c - \frac{\rho c}{1-2\rho}.$$

Similarly,

$$V(h_t\underline{y}) \geq \rho + (1-2\rho)(1-\delta)\sum_{s=0}^{\infty}\delta^s \underline{\varphi}_{s+t+1}\mathbf{E}^{P_C^C}[\sigma^C_{s+t+1}|h_t\underline{y}] - c + \frac{(1-\rho)c}{1-2\rho},$$

where $\underline{\varphi}_{s+t+1}$ denotes the lowest equilibrium consumer's belief on path in period $s+t+1$, and $\mathbf{E}^{P_C^C}[\sigma_{s+t+1}|h_t\underline{y}]$ denotes the competent firm's probability of high effort in period $s+t+1$ (where the expectation is taken over continuation histories conditional on $h_t\underline{y}$.). These two inequalities imply that a necessary condition for full-surplus optimality is

$$c \leq \delta(1-2\rho)^2\left(1-(1-\delta)\sum_{s=0}^{\infty}\delta^s \underline{\varphi}_{s+t+1}\mathbf{E}^{P_C^C}[\sigma^C_{s+t+1}|h_t\underline{y}]\right). \tag{2.39}$$

As $\eta \to 0$, it must be true that $\mathbf{E}^{P_C^C}[\sigma^C_{s+t+1}|h_t\underline{y}] \to 1$. Let

$$x(\eta) := \delta(1-2\rho)^2(1-\delta)\sum_{s=0}^{\infty}\delta^s \underline{\varphi}_{s+t+1}\mathbf{E}^{P_C^C}[\sigma^C_{s+t+1}|h_t\underline{y}].$$

By construction, (2.39) is violated for any $c \in (\bar{c}-x(\eta), \bar{c})$, yielding a contradiction.

2.10.7 Proof of Proposition 2.6

Suppose not. Then it must be true that the competent type (virtually) consistently exerts high effort. Because $K = \infty$, consumers' beliefs that the firm is competent converge to 1 as time grows conditional on the firm being competent, and the competent firm's stage payoffs converge to $1 - \rho - c$ given histories at which she is expected to exert high effort. Thus, given any history h upon acceptance and truthful reporting by the competent type, and letting $V(h)$ be the competent type's continuation payoff,

$$\lim_{\delta \to 1} V(h) = 1 - \rho - c. \tag{2.40}$$

Now, fix a history h' upon acceptance and truthful reporting by the competent firm, at which the competent firm exerts high effort. The incentive constraint for high effort must hold (recall that $\bar{y}$ denotes a good outcome and $\underline{y}$ denotes a bad outcome):

$$V(h\bar{y}) - V(h'\underline{y}) \geq \frac{(1-\delta)c}{\delta(1-2\rho)}.$$

Using this inequality, the competent firm's continuation payoff at history h' must satisfy

$$\begin{aligned} V(h') &\leq (1-\delta)(1-\rho-c) + \delta((1-\rho)V(h'\bar{y}) + \rho V(h'\underline{y})) \\ &\leq (1-\delta)(1-\rho-c) - \frac{(1-\delta)\rho c}{1-2\rho} + \delta V(h'\bar{y}). \end{aligned}$$

Similarly, fix another history h'' upon acceptance and truthful reporting by the competent firm, at which the competent firm shirks. The incentive constraint for high effort must be (weakly) violated:

$$V(h''\bar{y}) - V(h''\underline{y}) \leq \frac{(1-\delta)c}{\delta(1-2\rho)}.$$

Using this inequality, the competent firm's continuation payoff at history h'' must satisfy

$$\begin{aligned} V(h'') &\leq (1-\delta)\rho + \delta(\rho V(h''\bar{y}) + (1-\rho)V(h''\underline{y})) \\ &\leq (1-\delta)\rho - \frac{(1-\delta)\rho c}{1-2\rho} + \delta V(h''\underline{y}). \end{aligned}$$

Thus, at any history h, the competent type's continuation payoff upon acceptnace and truthful reporting satisfies (and recall that $1-\rho-c > c$ since high effort is efficient)

$$V(h) \leq (1-\delta)(1-\rho-c) - \frac{(1-\delta)\rho c}{1-2\rho} + \delta \max(V(h\bar{y}), V(h\underline{y})).$$

Recursively applying this inequality to the continuation payoff, it holds that

$$V(h) \leq 1-\rho-c - \frac{\rho c}{1-2\rho}.$$

Taking the limit $\delta \to 1$ on both sides yield

$$\lim_{\delta\to 1} V(h) \leq 1-\rho-c - \frac{\rho c}{1-2\rho}.$$

Because $c > 0$ and $\rho < 1/2$, (2.40) is violated, yielding a contradiction as desired.

2.10.8 Proof of Proposition 2.7

The proof is analogous to that of Proposition 2.1 in the main text. Here, I show that the competent type consistently exerts high effort upon participating. Consider the induced game $G(f_2', S_2', C)$. The firm's continuation play from $t \geq 1$ can be represented by a two-state automaton, with the set of states given by $\{0, 1\}$. State r collects the firm's histories with the most recent rating being r. The competent firm's continuation profit in state r is

$$V_C(r) = (1-\delta)(p_t(r) - c) + \delta((1-\rho)\alpha V_C(1) + (1-(1-\rho)\alpha)V_C(0)),$$

where

$$p_t(0) = \rho + (1-2\rho)\varphi_t(0) = \rho + \frac{(1-2\rho)\mu(1-(1-\rho)\alpha)}{\mu(1-(1-\rho)\alpha) + (1-\mu)(1-\rho\alpha)}, \tag{2.41}$$

$$p_t(1) = \rho + (1-2\rho)\varphi_t(1) = \rho + \frac{(1-2\rho)\mu(1-\rho)\alpha}{\mu(1-\rho)\alpha + (1-\mu)\rho\alpha}. \tag{2.42}$$

A deviation to shirk yields

$$V_C(r; \underline{e}) := (1-\delta)p_t(r) + \delta(\rho\alpha V_C^{\sigma}(1) + (1-\rho\alpha)V_C^{\sigma}(0)).$$

The incentive constraint $V_C(r) - V_C(r; \underline{e}) \geq 0$ in each state r can be simplified as

$$\delta\alpha(1-2\rho)^2\left(\frac{\mu(1-\rho)}{\mu(1-\rho) + (1-\mu)\rho} - \frac{\mu(1-\alpha(1-\rho))}{\mu(1-\alpha(1-\rho)) + (1-\mu)(1-\alpha\rho)}\right) \geq c. \tag{2.43}$$

By definition of α, given in (2.22), the left hand side equals c. Therefore, the incentive constraint holds. In period 0, each type of firm receives a payment

$$p_0(0) = \rho + \mu(1-2\rho) \tag{2.44}$$

upon accepting, and the competent firm faces the same incentive constraint (2.43).

Chapter 3

Strategic Manipulation in Tournament Games

3.1 Introduction

Tournaments are commonly used as an incentive scheme to elicit costly effort from economic agents.[1] However, many tournament designs incentivize shirking behavior from players, *irrespective of any effort cost.* Consider a simple example. Two players, a and b, play against each other in stage 1 of a tournament. The winner plays against c while the loser plays against d in stage 2. The winner of each match in stage 2 qualifies to be a final winner of the tournament. Suppose that when each player plays to win, a always beats b and d, but loses to c. Further, d always beats b and c when he plays to win, while b beats c. Each player prefers qualifying to a higher stage. If they play to win throughout, then a is eliminated in stage 2 and d is the final winner. However, a one-stage deviation by a to shirk and lose against b in stage 1 manipulates the qualification process and makes him the final winner at the expense

1. For example, they are used in sports events, labor promotions and innovation contests. Seminal papers include Lazear and Rosen (1981), Green and Stokey (1983) and Rosen (1986). Konrad (2009) provides a comprehensive survey.

of d, who is now eliminated in stage 2. Importantly, effort cost is irrelevant in reaching the decision to shirk. A famous real-world example is the 2012 London Olympics, where eight badminton players were disqualified for deliberately losing matches to meet preferred opponents in the next round.

The goal of the present paper is to characterize the set of tournaments immune to such manipulation. This is of interest for three reasons. First, the manipulation problem is in general detrimental to the profitability of organizing a tournament and the reputation of its organizers, and may arise in any economic setting with multistage competition. Second, it is often against a designer's incentive to simply disqualify any player who shirks in any tournament. In particular, it is often the stronger players who profit by shirking, and they shirk to secure their participations in later stages. The matches in later stages that feature stronger players tend to generate larger revenues. Finally, the outcome of a tournament is supposed to reflect the relative strengths of players. If the outcome is manipulated so that a player fails to qualify in an earlier stage than some weaker player, the outcome becomes a noisy indicator of the players' strengths.

To set the stage, I develop a simple model that allows for an *arbitrary* design of a tournament. This contrasts with existing work in the contest design literature, which largely studies the prize spread or assignment of opponents to players given a *specific* design. For example, see Rosen (1986), Moldovanu and Sela (2001), Groh, Moldovanu, Sela, and Sunde (2012) and Fu and Lu (2013). In particular, I build upon the traditional approach in the mathematics literature by defining a tournament as a pair consisting of a set of players and a binary relation on the set which captures the relative strength between any two players.[2] A designer then chooses a design for any given tournament, inducing a *tournament game* – a multistage game with observable effort. Succinctly, in a tournament game, each stage is defined as a partition of the

2. For example, see Rubinstein (1980). Laslier (1997) provides an extensive survey.

set of active players, while each element in the partition is a *group* of players, within which the players play pairwise matches against all others. The result of each match depends on the efforts exerted by the players and according to the match results, a ranking of players is established for each group. Qualification of each player to a higher stage depends on his rank in a group, subject to a quota. To illustrate, a *round-robin* tournament game has one stage of active play, in which all players belong to the same group. A *single-elimination* tournament game has a sequence of stage partitions, each of which contains a number of groups of two, with the winner in each group qualifying to the next stage. The 2014 FIFA World Cup had eight groups of four teams from which two qualified in stage 1, and played as a single-elimination tournament beginning from stage 2.

I characterize the set of *incentive-compatible* tournament games which are tournament games with the following two properties, irrespective of players' strengths. First, to capture sequential rationality, full effort exertion is a *subgame perfect equilibrium*, so that there exists at least one equilibrium with no manipulation. Second, the outcome under *each* subgame perfect equilibrium must coincide with the outcome prescribed by the *full-effort social choice function*, which maps the relative strengths between players and a given design to an outcome under which the strongest players from a group qualify and among these qualifying players a stronger player ranks higher. Put simply, this protects the outcome from manipulation even when a deviation away from full effort exertion to another equilibrium takes place. Plainly, the main result we are going to prove is the following: *Every incentive-compatible tournament game allows only the top-ranked player from each group to qualify.* In the proof, we shall see that in tournament games with multiple qualifiers in some group, some player can often profit by shirking to qualify with a lower rank. Indeed, in *every* tournament game with multiple qualifiers in some group, there always exists an equilibrium with an outcome different from the outcome prescribed by the full-effort social choice function.

This paper differs from related work by providing an analytical *game-theoretic foundation* and allowing for an *arbitrary design*. For instance, Dagaev and Sonin (2013) borrows tools from social choice theory to show that multiple qualifier systems, where players compete in several local tournaments to qualify to international tournaments, are manipulable. On the other hand, Pauly (2014) uses a computer-assisted proof to show that the designs of the 2012 Olympics Badminton and the 2014 FIFA World Cup give rise to manipulation.

3.2 Model

Let N be a finite set of players. A tournament is a pair $(N, \rightarrow)$, where $\rightarrow$ is a complete, asymmetric and irreflexive binary relation on N, so that $i \rightarrow j$ if i is *(relatively) stronger than* j. Given N, a designer fixes a *design* D_N which induces a tournament as a multistage game, referred to as a *tournament game*. A design consists of four components:

1. a finite number of stages and groups in each stage, as well as the size and a qualification quota of each group;

2. a partition P^1 of the set N, where each element $G \in P^1$ is a group in stage 1;

3. a sorting rule f^t for each stage $t > 1$ that sorts players qualifying from stage $t-1$ into groups in stage t;

4. a tie-breaker $\succ_{\text{tie}}$, which is a strict linear order on N.

We label each group i in stage t by G_i^t and simply write G for an arbitrary group unless there is a possibility of ambiguity. The order $\rightarrow$ on N is known to the players and the designer only *after* the design is fixed. A player is *active* before he is eliminated. In every stage t, each active player is sorted into some group G.

Within G, every player i plays a *match* against each opponent j, where each *costlessly* chooses an effort level from $[0, 1]$ against his opponent. The player who exerts a higher effort wins the match. When both exert the same effort, i wins if $i \to j$. Let $w_i(G) \equiv |\{\iota \in G : i \text{ wins against } \iota\}|$ denote i's number of wins in G. Define a ranking of players $\succ_G$ on G such that $i \succ_G j$, or i *ranks higher than* j *in* G, if $w_i(G) > w_j(G)$, or if $w_i(G) = w_j(G)$ and $i \succ_{\text{tie}} j$. Also let $w_i^{FE}(G) \equiv |\{\iota \in G : i \to \iota\}|$ denote i's number of wins in G if all players exert the same effort (e.g., full effort, or effort 1).

The q_G highest ranked players in G by $\succ_G$ qualify from G to the next stage. Let $|G| \geq 2$ and $1 \leq q_G \leq |G| - 1$, so that at least one player is eliminated and at least one qualifies from each group. In the next stage, active players are sorted into groups by a sorting rule. A sorting rule f^t partitions the set of active players in stage t according to the label of the groups they qualify from in stage $t-1$ and their ranks in the groups, with the partition denoted by P^t such that each $G \in P^t$ is a group. The game ends when players from stage $T \geq 1$ qualify to stage $T+1$, where no sorting of players and no match take place.

Efforts are observed at the end of each stage and there is perfect recall. Let H^t be the set of histories in the beginning of stage t, which contains a typical element $h^1 = \emptyset$ or $h^t = (a^1, \ldots, a^{t-1})$ for $t \geq 2$, where $\emptyset$ denotes the null history and a^t is the vector of efforts exerted by players in stage t. Let $A_i(h^t)$ be player i's set of feasible effort choices in stage t when the history is h^t. A pure strategy by i is a sequence of maps $(s_i^t)_{t=1}^T$, where each $s_i^t : H^t \to \cup_{h^t \in H^t} A_i(h^t)$, so that $s_i^t(h^t) \in A_i(h^t)$ for every h^t. Each i chooses a strategy $s_i = (s_i^t)_{t=1}^T$ to maximize his payoff defined by $u_i(s_i, s_{-i}) = t$, given the strategy s_{-i} by all players except i, where t is the highest stage player i qualifies to. Of particular interest is the profile s^*, where players exert full effort throughout (i.e., after all histories).

Let G' denote the set of players qualifying from a group G. An *outcome* of a tournament game is a collection of ordered sets $(G', \succ_G)$ for each group G in each stage $t = 1, \ldots, T$. Define $g(s)$ as the outcome yielded by some strategy profile s of the game. A *social choice function* (henceforth, SCF) F maps $(\rightarrow, D_N)$ into an outcome $F(\rightarrow, D_N)$. Let F^{FE} be the *full-effort* SCF where $F^{FE}(\rightarrow, D_N)$ is the outcome such that for each group $G \in P^t$ in each stage $t = 1, \ldots, T$, for any two players i, j who qualify (that is, $i, j \in G'$) and any player k who is eliminated from G,

1. $\min\{w_i^{FE}(G), w_j^{FE}(G)\} \geq w_k^{FE}(G)$;
2. $i \succ_G j$ if and only if $w_i^{FE}(G) \geq w_j^{FE}(G)$, where the tie breaker $i \succ_{\text{tie}} j$ applies when the inequality holds with equality.

A pair (F, s) of a SCF F and a strategy profile s is *subgame perfect implementable* if s is a subgame perfect equilibrium (henceforth, SPE) and $g(s') = F(\rightarrow, D_N)$ for *any* SPE s'. There is at least one SPE because a tournament game is a finite game of perfect information.

A tournament game is *incentive-compatible* (henceforth, IC) if (F^{FE}, s^*) is subgame perfect implementable for every order $\rightarrow$ on N. Intuitively, the universal quantification is desirable because the relative strengths are unknown before a design is fixed. Examples 10 and 11 below show that the two requirements in the implementation of (F^{FE}, s^*) are independent.

Remark 3.1. The fact that effort is costless and that players derive a unit marginal payoff upon qualifying over each stage, albeit stylized, serves to emphasize that effort cost and prize spread can be of little relevance to players' effort choices when a tournament is not IC. Moreover, pairwise matching stands in contrast to a more general assumption where multiple players can possibly be matched simultaneously. For our purpose, pairwise matching delivers an advantage by its clear application: a player often chooses a specific opponent against whom he shirks, and pairwise matching allows us to capture such action.

Remark 3.2. The outcome prescribed by the full-effort SCF "imitates" the full effort outcome in any equilibrium. Naturally, the outcome prescribed by a SCF contains the elements that the designer cares about. According to the full-effort SCF, he is concerned with the identity of and the observed rankings of the players who qualify from each group in a stage. A less restrictive requirement would be that he cares about their identities but not the observed rankings, or cares about the observed rankings only if they affect subsequent partitions. The current definition of a full-effort SCF has the advantage over the two less restrictive ones of revealing information about the relative quality among the qualifying players in each group. We illustrate this by the following simple example.

Example 6. Consider a tournament $(N, \to)$ where $N = \{a, b, c\}$, $a \to b$, $a \to c$ and $b \to c$. The design D_N is as follows. The first stage partition is $P^1 = \{A\}$, where $A = N$, with $q_A = 2$. The two players qualifying from A are the final winners. Clearly, players maximize their utilities by exerting full effort in each match such that $a \succ_A b \succ_A c$, and s^* is a SPE. Each player obtains the same utility in another SPE equivalent to s^* except that a deviates to lose against b such that $b \succ_A a \succ_A c$. If the second equilibrium ranking is observed, one could potentially conclude that ranking reveals that the relative quality of player b is the highest among the participating players, when this is not the case. ♦

Furthermore, observe that the full-effort SCF imposes a restriction on the rankings of *qualifying* players instead of *all* players in each group. The example below shows that the latter requirement is so restrictive that even a simple round-robin tournament is not subgame perfect implementable. Denote by $G[r]$ the r^{th} ranked player in a group G.

Example 7. Consider a tournament $(N, \to)$, where $N = \{a, b, c, d, e, f\}$ and

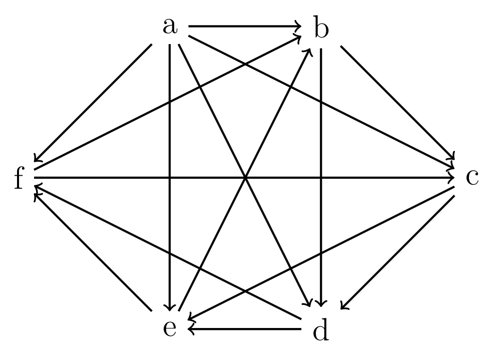

.

The design is one of *round-robin*: $P^1 = \{G\}$ where $G = \{a, b, c, d, e, f\}$, and $G[1]$ is the final winner. Consider a strategy profile s equivalent to s^* except that a shirks against f in G. The tie-breaker satisfies $a \succ_{\text{tie}} b \succ_{\text{tie}} c \succ_{\text{tie}} d \succ_{\text{tie}} e \succ_{\text{tie}} f$. It is straightforward to check that s is a SPE. In particular, under s we have $(w_a, w_b, w_c, w_d, w_e, w_f) = (4, 2, 2, 2, 2, 3)$, where $w_i \equiv w_i(G)$ for each player i. As a result, f ranks strictly above all other players except a in G. Consider a social choice function F_*^{FE} that requires that the ranking of all players, as opposed to only the winner as does F^{FE}, are equal to the ranking under full effort. Then F_*^{FE} designates that $i \succ_G f$ for every $i \in N \backslash \{f\}$, so $g(s) \neq F^{FE}(\to, D_N)$. Intuitively, because a is far stronger than his opponents in the group, he is indifferent between winning all matches or winning all except losing to f. Losing to f, however, manipulates the ranking. ♦

3.3 Incentive-Compatible Tournament Games

In this section, we formally state and prove the main theorem. Perhaps surprisingly, as far as IC is concerned, all components in a design except the qualification quota in each group are irrelevant.

Theorem 3.1. *Fix a tournament $(N, \to)$. A tournament game induced by a design*

D_N *is IC if and only if the design designates that* $q_G = 1$ *for every group* G *in any stage* $t \leq T$.

Before proceeding to the proof, we study a few examples to illustrate the model and to build intuition underlying the result. Example 8 illustrates a design which, with multiple qualifiers allowed in some group, induces a tournament game where s^* fails to be a SPE.

Example 8. Consider a tournament $(N, \to)$, where $N = \{a, b, c, d, e, f\}$ and

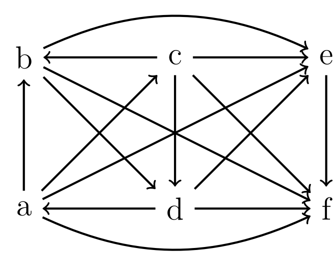

.

The design D_N is as follows. The first stage partition is $P^1 = \{A, B\}$, where $A = \{a, b, e\}$ and $B = \{c, d, f\}$, with $q_A = q_B = 2$. Then $P^2 = \{C, D\}$, where $C = \{A[1], B[2]\}$ and $D = \{A[2], B[1]\}$, with $q_C = q_D = 1$. The players $C[1]$ and $D[1]$ are the final winners. Consider the outcome under s^*. In stage 1, $a \succ_A b \succ_A e$ and $c \succ_B d \succ_B f$. In stage 2, $C = \{a, d\}$ and $D = \{b, c\}$, with $d \succ_C a$ and $c \succ_D b$. The final winners are c and d. If a unilaterally deviates to shirk against b in stage 1, then the ranking in A becomes $b \succ_A a \succ_A e$, so that $C = \{b, d\}$ and $D = \{a, c\}$ in stage 2. Because $a \succ_D c$, a qualifies as a final winner and profits from the deviation. ♦

The key insight from the example is that we can always find a relation $\to$ on N such that s^* fails to be a SPE whenever a player can "choose" which group to qualify to. Nonetheless, this should not be confused with the statement that one can always find a relation $\to$ on N such that s^* fails to be a SPE whenever multiple qualifiers are allowed in some group. In particular, Example 9 below shows that the latter claim is false. Furthermore, building on the same setup, Example 10 shows that while s^* is a

SPE, there exists another SPE s inducing an outcome $g(s)$ that does not agree with any outcome prescribed by the full-effort SCF.

Example 9. Consider a tournament $(N, \to)$ where $N = \{a, b, c\}$ and $\to$ is arbitrary. The design D_N is as follows. The first stage partition is $P^1 = \{A\}$, where $A = \{a, b, c\}$ and $q_A = 2$. Then $P^2 = \{B\}$, where $B = \{A[1], A[2]\}$ and $q_B = 1$. $B[1]$ is the final winner. Given any history h^2, exerting full effort in stage 2 by each active player is clearly a Nash equilibrium. Consider the first stage. Depending on the precise specification of $\to$ on N, there are two possibilities in A when players exert full effort. First, the first-ranked player has 2 wins, the second-ranked player has 1, while the third-ranked player has 0. Second, each of them has 1 wins, and they are ranked according to some given tie-breaker. Clearly, in either case, no player would be strictly better off by deviating to shirk in any match. Therefore, s^* is a SPE for any $\to$. ♦

Example 10. Consider the tournament game given above in Example 9, and suppose that $a \to b$, $b \to c$ and $a \to c$. The outcome implemented by the SCF F^{FE} is such that $a \succ_A b \succ_A c$; $B = \{a, b\}$ and $a \succ_B b$, implying that a is the final winner. It is easy to verify that s^* is a SPE. Further, consider a strategy profile s equivalent to s^* except that a shirks against b in stage 1. Under s, $b \succ_A a \succ_A c$. Then $B = \{a, b\}$, $a \succ_B b$ and a is the final winner. Being the final winner, a has no profitable deviation. Neither does b, because by shirking against a in either stage he cannot change his payoff, and by shirking against c in stage 1, he can possibly be eliminated a stage earlier, depending on the specified tie-breaker. Finally, it is clear that c has no profitable deviation. Thus there exists a SPE s such that $g(s) \neq F^{FE}(\to, D_N)$. ♦

Conversely, Example 11 below illustrates a setting where $g(s) = F^{FE}(\to, D_N)$ for any SPE s, but s^* is not a SPE in the tournament game. Together with Example 10, this shows that the requirements of s^* being a SPE and $g(s) = F^{FE}(\to, D_N)$ for every SPE s are independent.

Example 11. Consider a tournament $(N, \rightarrow)$, where $N = \{a, b, c, d, e, f, g\}$ and

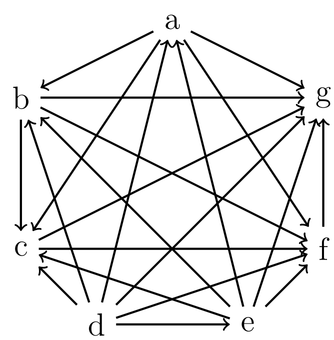

.

Let the design D_N designate that $P^1 = \{A, B, C\}$ where $A = \{a, b, c\}$ with $q_A = 2$, $B = \{d, e\}$ with $q_B = 1$ and $C = \{f, g\}$ with $q_C = 1$. Then $P^2 = \{D, E\}$, where $D = \{A[1], B[1]\}$ and $E = \{A[2], C[1]\}$. The players $D[1]$ and $E[1]$ are the final winners. Finally, suppose that the tie-breaker follows $c \succ_{\text{tie}} b \succ_{\text{tie}} a \succ_{\text{tie}} d \succ_{\text{tie}} e \succ_{\text{tie}} f \succ_{\text{tie}} g$. By construction of $\rightarrow$, by exerting full effort in group E, $A[2]$ ultimately becomes a final winner, while $A[1]$ is eliminated from group D in stage 2 as long as $B[1]$ exerts full effort. Moreover, in *any* SPE, both a and b never rank the lowest in A and are eliminated, for a always qualifies by exerting full effort in every match in A, while b always qualifies by exerting full effort against c. So in equilibrium, either $a \succ_A b \succ_A c$ or $b \succ_A a \succ_A c$. If the latter is true, then b must have 2 wins, a has 1, and c has 0, by construction of $\succ_{\text{tie}}$. Thus a must have shirked against b. But then b can be strictly better off by shirking against a and qualify as $A[2]$, contradicting subgame perfection. So it must be true that $a \succ_A b \succ_A c$, and a has 2 wins, b has 1 and c has 0 by construction of $\succ_{\text{tie}}$. In particular, b must have shirked against a, for otherwise a could shirk against him and profit by qualifying as $A[2]$. On the other hand, it should be clear that the stronger player in each group B, C, D and E would exert full effort in equilibrium. Thus $g(s) = F^{FE}(\rightarrow, D_N)$ in any SPE s. But s^* is clearly not a SPE in the induced tournament game, for a would profit by unilaterally deviating to shirk against b in A to qualify as $A[2]$. ♦

One key assumption underlying Theorem 3.1 is that the sorting rule in each stage relies on the ranks of each qualified player instead of their numbers of wins in the previous stage. Indeed, the latter measure provides an immediate medium for a player to manipulate his standing and manipulating behavior may arise even when the qualification quota is one in each group as shown Example 12 below.

Example 12. Fix some stage in an arbitrary tournament game. Suppose that if all players exert full effort then player i ranks first with w wins in some group G the current stage, qualifies to group A in the next stage and is eliminated. Suppose also that no other player in G has more than $w-3$ wins in the current stage, so i would still qualify if he lost one match in G. Suppose that when i ranks first in G with $w-1$ wins he qualifies to group B in the next stage and then ranks first in B as long as he exerts full effort. Then exerting full effort after all histories is not a SPE. ⧫

We are now ready to prove the main result.

3.3.1 Proof of Theorem 3.1

Fix a tournament $(N, \rightarrow)$ and a design D_N with only one qualifier from each group in each stage. Fix s^* as the strategy profile. To show that it is a SPE, it suffices to show that each player is best responding to others' strategies from the beginning of each stage given any history, since every stage begins with a proper subgame. Clearly, this is true for any final winner. This is also true for a player eliminated in any stage, because any deviation would see him remain eliminated since only one player can qualify from each group and whenever qualified the player would be sorted in the same group in the next stage upon deviation. Next, we show that $g(s) = F^{FE}(\rightarrow, D_N)$ for any SPE s. Observe that, the outcome designated by F^{FE} satisfies that i qualifies from G if and only if $i \in \text{argmax}_{j \in G} w_j^{FE}(G)$ and $i \succ_{\text{tie}} k$ for any other maximizer k. If $g(s) \neq F^{FE}(\rightarrow, D_N)$ for some SPE s, then there must exist a group G in some stage t after some history h^t under s, and two players $i, j \in G$ so that $j \succ_G i$ but $i \in \text{argmax}_{l \in G} w_l^{FE}(G)$ and $i \succ_{\text{tie}} k$ for any other maximizer k. Because $q_G = 1$, i is eliminated in stage t. But the fact that i is a maximizer implies that i must have shirked in at least one match in G under s. If he deviates to exert full effort in those matches, then he must rank first and qualify, contradicting that s is a SPE.

Conversely, it suffices to show for each design D_N with $q_G > 1$ in a group G in some stage t that, for some order $\to$ on N, either s^* is not a SPE or that $g(s) \neq F^{FE}(\to, D_N)$ for some SPE s. Fix a history h^t, after which $G = \{1, 2, \ldots, g\}$. Let the order $\to$ on N designate that $i \to j$ if $i < j$ for any two players $i, j \in G$. Thus $i \succ_G j$ if $i < j$ under s^*, and both 1 and 2 qualify from G. Upon qualification, two possibilities arise, depending on the sorting rule f^{t+1}. The first possibility entails that $1, 2 \in A$ for some group A in stage $t+1$. Let s^* be a SPE, for otherwise the claim is proved in this case. Consider a one-stage deviation s from s^* in stage t where 1 shirks against 2 in G. The rankings of all groups in P^t are unchanged, except that the ranks of 1 and 2 switch in G so that $g(s) \neq F^{FE}(\to, D_N)$. We argue that s is a SPE: first, the stage partitions $(P^t)_{t=1}^T$ are the same under s and s^*, which implies that players face the same opponents in each stage under both profiles; second, s^* is a SPE. Therefore IC fails. It remains to consider the second possibility, where 1 and 2 are sorted into two separate groups, say B and C respectively, in stage $t+1$. Further let the order $\to$ on N designate that, $i \to 1$ for each $i \in B \backslash \{1\}$ and $1 \to j$ for each $j \in C$. Clearly, 1 ranks the lowest in B and is eliminated in stage $t+1$. However, s^* is not a SPE because a one-stage deviation by 1 to shirk against 2 in G would allow 1 to qualify to C in stage $t+1$, in which he ranks the first and qualifies further. This completes the proof.

3.4 Concluding Remarks

The result has the flavor of an impossibility result by implying that many tournament games in practice are not IC. Nonetheless, in practice, regulations that allow only one qualifier from each group may not be desirable. The appeal of IC may therefore be put into question. Imagine an unknown tennis player being drawn against Roger Federer in the same group in the first stage of a single-elimination tournament. The tournament game is IC, but is perhaps too harsh on the athletes and their fans. Of course, there are many factors other than IC, for instance time constraint and the number of matches, that a designer needs to take into account when designing a tournament. To be clear, the aim of the present paper is *not* to provide a general selection criterion of tournament designs. Instead it seeks to provide a baseline framework for future work. For instance, the current solution concept is *binary*: a tournament game is either IC or not. Devising a method to quantify IC would allow for the study of trade-offs between IC and other potentially desirable factors.

Bibliography

G. L. Albano and A. Lizzeri. Strategic Certification and Provision of Quality. *International Economic Review*, 42(1):267–283, 2001.

C. D. Aliprantis and K. C. Border. *Infinite Dimensional Analysis: A Hitchhiker's Guide*. 3rd edition. Springer-Verlag, Berlin, 2006.

A. Atkeson, C. Hellwig, and G. Ordoñez. Optimal Regulation in the Presence of Reputation Concerns. *The Quarterly Journal of Economics*, 130(1):415–464, 2015.

C. N. Avery and J. A. Chevalier. Herding Over The Career. *Economics Letters*, 63 (3):327–333, 1999.

M. Backus and A. T. Little. I Dont Know. *American Political Science Review*, 114 (3):724–743, 2020.

M. Baddeley. *Copycats & Contrarians: Why We Follow Others... and When We Don't*. Yale University Press, 2018.

H. Bar-Isaac and S. Tadelis. Seller Reputation. *Foundations and Trends® in Microeconomics*, 4(4):273–351, 2008.

R. Benabou and G. Laroque. Using Privileged Information to Manipulate Markets: Insiders, Gurus, and Credibility. *The Quarterly Journal of Economics*, 107(3): 921–958, 1992.

V. Bhaskar and C. Thomas. Community Enforcement of Trust with Bounded Memory. *The Review of Economic Studies*, 86(3):1010–1032, 2019.

S. Board and M. Meyer-ter Vehn. Reputation for Quality. *Econometrica*, 81(6): 2381–2462, 2013.

A. Bonatti and J. Hörner. Career Concerns with Exponential Learning. *Theoretical Economics*, 12(1):425–475, 2017.

A. Bourgoin and J.-F. Harvey. Professional Image Under Threat: Dealing with Learning-Credibility Tension. *Human Relations*, 71(12):1611–1639, 2018.

Z. Bozanic, J. Chen, and M. J. Jung. Analyst Contrarianism. *Journal of Financial Reporting*, 4(2):61–88, 2019.

A. Brad, A. Delemare, N. Hurley, V. Lenikus, R. Mulrenan, N. Nemes, U. Trunk, and N. Urbancic. The False Promise of Certification. *Changing Markets Foundation, Utrecht*, 2018.

A. Brandenburger and B. Polak. When Managers Cover Their Posteriors: Making The Decisions The Market Wants To See. *The RAND Journal of Economics*, pages 523–541, 1996.

Y.-K. Che and J. Hörner. Recommender Systems as Mechanisms for Social Learning. *The Quarterly Journal of Economics*, 133(2):871–925, 2018.

G. Cisternas. Two-sided Learning and the Ratchet Principle. *The Review of Economic Studies*, 85(1):307–351, 2018.

J. C. Coffee. Brave New World?: The Impact(s) of the Internet on Modern Securities Regulation. *The Business Lawyer*, pages 1195–1233, 1997.

D. Dagaev and C. Sonin. Winning by Losing: Incentive Incompatibility in Multiple Qualifiers. *Discussion paper series*, Centre for Economic Policy Research, 2013.

R. Deb and C. Stewart. Optimal Adaptive Testing: Informativeness and Incentives. *Theoretical Economics*, 13(3):1233–1274, 2018.

R. Deb, M. Pai, and M. Said. Evaluating Strategic Forecasters. *American Economic Review*, 108(10):3057–3103, 2018.

R. Deb, M. F. Mitchell, and M. Pai. (Bad) Reputation in Relational Contracting. 2020.

M. Dewatripont, I. Jewitt, and J. Tirole. The Economics of Career Concerns, Part I: Comparing Information Structures. *The Review of Economic Studies*, 66(1): 183–198, 1999a.

M. Dewatripont, I. Jewitt, and J. Tirole. The Economics of Career Concerns, Part II: Application to Missions and Accountability of Government Agencies. *The Review of Economic Studies*, 66(1):199–217, 1999b.

D. Dranove and G. Z. Jin. Quality Disclosure and Certification: Theory and Practice. *Journal of Economic Literature*, 48(4):935–63, 2010.

M. R. Effinger and M. K. Polborn. Herding and Anti-herding: A Model of Reputational Differentiation. *European Economic Review*, 45(3):385–403, 2001.

M. Ekmekci. Sustainable Reputations with Rating Systems. *Journal of Economic Theory*, 146(2):479–503, 2011.

J. Ely and J. Välimäki. Bad Reputation. *The Quarterly Journal of Economics*, 118 (3):785–814, 2003.

G. Eyal. *The Crisis of Expertise*. John Wiley & Sons, 2019.

D. P. Foster and R. V. Vohra. Asymptotic Calibration. *Biometrika*, 85(2):379–390, 1998.

Q. Fu and J. Lu. The Optimal Multi-stage Contest. *Economic Theory*, 51(2):351–82, 2013.

C. Funk, M. Hefferon, B. Kennedy, and C. Johnson. Trust and Mistrust in Americans Views of Scientific Experts. *Pew Research Center*, 2, 2019.

M. Gentzkow and J. M. Shapiro. Media Bias And Reputation. *Journal of Political Economy*, 114(2):280–316, 2006.

R. J. Goodland. *Ecolabeling: Opportunties for Progress Toward Sustainability*. Consumer's Choice Council, 2002.

J. Green and N. Stokey. A Comparison of Tournaments and Contracts. *Journal of Political Economy*, 91(3):349–64, 1983.

C. Groh, B. Moldovanu, A. Sela, and U. Sunde. Optimal Seedings in Elimination Tournaments. *Economic Theory*, 49(1):59–80, 2012.

M. Gurri. *The Revolt of the Public and the Crisis of Authority in the New Millennium*. Stripe Press, 2018.

M. Halac and I. Kremer. Experimenting with Career Concerns. *American Economic Journal: Microeconomics*, 12(1):260–88, February 2020.

R. Harbaugh and E. Rasmusen. Coarse Grades: Informing the Public by Withholding Information. *American Economic Journal: Microeconomics*, 10(1):210–35, 2018.

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

H. Hong and J. D. Kubik. Analyzing the Analysts: Career Concerns and Biased Earnings Forecasts. *The Journal of Finance*, 58(1):313–351, 2003.

H. Hong, J. D. Kubik, and A. Solomon. Security Analysts' Career Concerns and Herding of Earnings Forecasts. *The Rand Journal of Economics*, pages 121–144, 2000.

H. Hopenhayn and M. Saeedi. Optimal Quality Ratings and Market Outcomes. *NBER Working Paper*, 26221, 2019.

J. Hörner. Reputation and Competition. *American Economic Review*, 92(3):644–663, 2002.

J. Hörner and N. Lambert. Motivational Ratings. *The Review of Economic Studies*, 11 2020. ISSN 0034-6527. doi: 10.1093/restud/rdaa070. URL `https://doi.org/10.1093/restud/rdaa070`.

G. Howells. The Potential and Limits of Consumer Empowerment by Information. *Journal of Law and Society*, 32(3):349–370, 2005.

N. Klein and T. Mylovanov. Will Truth Out? An Advisors Quest to Appear Competent. *Journal of Mathematical Economics*, 72:112–121, 2017.

K. Konrad. *Strategy and Dynamics in Contests*. Oxford University Press, New York, 2009.

S. Kovbasyuk and G. Spagnolo. Memory and Markets. 2018.

J.-F. Laslier. *Tournament Solutions and Majority Voting*, volume 7 of *Studies in Economic Theory*. Springer, 1st edition, 1997.

E. Lazear and S. Rosen. Rank-order Tournaments as Optimum Labor Contracts. *Journal of Political Economy*, 89(5):841–64, 1981.

G. Levy. Anti-herding and Strategic Consultation. *European Economic Review*, 48 (3):503–525, 2004.

W. Li. Changing One's Mind When the Facts Change: Incentives of Experts and the Design of Reporting Protocols. *The Review of Economic Studies*, 74(4):1175–1194, 2007.

Q. Liu. Information Acquisition and Reputation Dynamics. *The Review of Economic Studies*, 78(4):1400–1425, 2011.

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

A. Lizzeri. Information Revelation and Certification Intermediaries. *The RAND Journal of Economics*, 30:214–231, 1999.

A. Lizzeri, M. A. Meyer, and N. Persico. The Incentive Effects of Interim Performance Evaluations. 2002.

C. Lorecchio and D. Monte. Dynamic Information Design under Constrained Communication Rules. 2021.

G. Mailath and L. Samuelson. Who Wants a Good Reputation? *The Review of Economic Studies*, 68(2):415–441, 2001.

S. Majumdar and S. W. Mukand. Policy Gambles. *American Economic Review*, 94 (4):1207–1222, 2004.

I. Marinovic, M. Ottaviani, and P. N. Sørensen. Forecasters Objectives and Strategies. In *Handbook of Economic Forecasting*, volume 2, pages 690–720. Elsevier, 2013.

I. Marinovic, A. Skrzypacz, and F. Varas. Dynamic certification and reputation for quality. *American Economic Journal: Microeconomics*, 10(2):58–82, 2018.

A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic Theory*, volume 1. Oxford University Press New York, 1995.

R. McLachlin. Service Quality in Consulting: What is Engagement Success? *Managing Service Quality: An International Journal*, 2000.

B. Moldovanu and A. Sela. The Optimal Allocation of Prizes in Contests. *American Economic Review*, 91:542–58, 2001.

A. Momparler, P. Carmona, and C. Lassala. Quality of Consulting Services and Consulting Fees. *Journal of Business Research*, 68(7):1458–1462, 2015.

S. Morris. Political correctness. *Journal of Political Economy*, 109(2):231–265, 2001.

R. B. Myerson. Multistage Games with Communication. *Econometrica*, pages 323–358, 1986.

T. Nichols. *The Death of Expertise: The Campaign Against Established Knowledge and Why It Matters*. Oxford University Press, 2017.

OECD. Actual Effect of Selected Programs. *Organization for Economic Co-Operation and Development, GD*, 97, 1997.

W. Olszewski. Calibration and Expert Testing. In *Handbook of Game Theory with Economic Applications*, volume 4, pages 949–984. Elsevier, 2015.

M. Ottaviani and P. N. Sørensen. Professional Advice. *Journal of Economic Theory*, 126(1):120–142, 2006a.

M. Ottaviani and P. N. Sørensen. Reputational Cheap Talk. *The Rand Journal of Economics*, 37(1):155–175, 2006b.

M. Ottaviani and P. N. Sørensen. The Strategy of Professional Forecasting. *Journal of Financial Economics*, 81(2):441–466, 2006c.

M. Pagano and P. Volpin. Credit ratings failures and policy options. *Economic Policy*, 25(62):401–431, 2010.

F. Partnoy. How And Why Credit Rating Agencies Are Not Like Other Gatekeepers. 2006.

F. Partnoy. What's (Still) Wrong with Credit Ratings. *Wash. L. Rev.*, 92:1407, 2017.

M. Pauly. Can Strategizing in Round-robin Subtournaments be Avoided? *Social Choice and Welfare*, 43(1):29–46, 2014.

F. Pavesi and M. Scotti. Good Lies. 2019.

A. Prat. The Wrong Kind of Transparency. *American Economic Review*, 95(3): 862–877, 2005.

C. Prendergast. A Theory of "Yes Men". *The American Economic Review*, pages 757–770, 1993.

C. Prendergast and L. Stole. Impetuous Youngsters and Jaded Old-timers: Acquiring a Reputation for Learning. *Journal of Political Economy*, 104(6):1105–1134, 1996.

R. Radner. Repeated Principal-Agent Games with Discounting. *Econometrica*, 53 (5):1173–1198, 1985.

S. Rosen. Prizes and Incentives in Elimination Tournaments. *American Economic Review*, 76(4):701–15, 1986.

S. M. Ross. *Introduction to Stochastic Dynamic Programming.* Academic Press, 2014.

A. Rubinstein. Ranking the Participants in a Tournament. *Journal of the Society of Industrial and Applied Mathematics*, 38:108–11, 1980.

J. Rüdiger and A. Vigier. Learning About Analysts. *Journal of Economic Theory*, 180:304–335, 2019.

D. S. Scharfstein and J. C. Stein. Herd Behavior and Investment. *American Economic Review*, 80(3):465–479, 1990.

A. Smirnov and E. Starkov. Timing of Predictions in Dynamic Cheap Talk: Experts vs. Quacks. *University of Zurich, Department of Economics, Working Paper*, (334), 2019.

J. Sobel. A Theory of Credibility. *The Review of Economic Studies*, 52(4):557–573, 1985.

K. Stahl and R. Strausz. Certification and Market Transparency. *The Review of Economic Studies*, 84(4):1842–1868, 2017.

N. L. Stokey, R. Lucas, and E. Prescott. Recursive Methods in Economic Dynamics. *Cambridge, MA: Harvard University*, 1989.

B. Trueman. Analyst Forecasts and Herding Behavior. *The Review of Financial Studies*, 7(1):97–124, 1994.

F. Varas, I. Marinovic, and A. Skrzypacz. Random Inspections and Periodic Reviews: Optimal Dynamic Monitoring. *The Review of Economic Studies*, 03 2020. ISSN 0034-6527.

N. Vellodi. Ratings Design and Barriers to Entry. 2019.

A. Zapechelnyuk. Optimal Quality Certification. *American Economic Review: Insights*, 2(2):161–76, 2020.

J. Zwiebel. Corporate Conservatism and Relative Compensation. *Journal of Political Economy*, 103(1):1–25, 1995.