Mind Changes and the Design of Reporting Protocols*

Wei Li†

June, 2004

Abstract

In organizational, political, and financial settings, information is collected and reported by experts as it is received over time. This paper studies, in such dynamic situations, the incentives of an expert with reputational concerns to reveal his most recent information and the reporting protocol that induces the most truthful revelation. A principal receives sequential reports from an agent of privately known ability, who privately observes signals about the state of the world. The agent's signals are of different initial quality and, in contrast to previous work, also of different quality *improvement*. First, when the talented agent also improves faster, "mind changes" (inconsistent reports) may be a sign of high ability, yet a mediocre agent still tends to repeat his early report. Second, requiring sequential reports creates an incentive to misreport the final, more accurate signal, but requiring a single report can only extract the agent's final, not interim, opinion. As a result, sequential reports dominate when the principal's optimal decision is very sensitive to the reports' accuracy. A single report dominates when either the mediocre agent's signals improve faster, or when the agent is very unlikely to be talented.

Keywords: sequential information transmission, reputational concerns, cheap talk games.
JEL classification: D82, C70, M50

*This paper is adapted from chapter three of my MIT dissertation. I thank George Akerlof, Abhijit Banerjee, Mathias Dewatripont, Glenn Ellison, Botond Kőszegi, Jean Tirole and Muhamet Yildız for insightful comments and encouragement. In addition, this paper has benefited from the input of seminar participants at MIT, UCR, UCSD, and UCLA.

†University of California, Riverside. Email: wei.li@ucr.edu.

1 Introduction

In most economic models of communication, information is collected once and transmitted in a single piece from a sender to a receiver (Crawford and Sobel 1982, Aghion and Tirole 1997, Morris 2001). In many realistic settings, however, a sender receives *multiple* pieces of information over time and is asked to convey his opinion *multiple* times as more information comes in.[1] Formal sequential reports are observed frequently in congressional committees, accounting (Dye and Verrecchia 1995), capital budgeting (Arya and Sivaramakrishnan 1997), and the financial market (Penno 1985). Informally, consultants, doctors, and other professionals are often asked to convey their early opinions before giving a final report.

In many of these environments, the sender's ability to observe the underlying true state of the world and thus the quality of his information improves over time, as he becomes more familiar with the task at hand. Moreover, the sender typically cares about how his reports reflect on his ability. This paper investigates how an agent of privately known ability reacts strategically to improvement in the quality of his information under a sequential reports system. It also applies these insights to the optimal choice of reporting protocols. Namely, it identifies conditions under which the principal should require a report after the agent has received all the information, and conditions under which she should ask for sequential reports instead.

In the basic model, an agent delivers an interim report and a final report about the state of the world, based on his sequence of private signals. To reflect improvement in the agent's ability to observe the true state of the world, he receives signals of increasing accuracy. After each signal, the agent sends a report to the principal, who makes a decision after the final report. Next, the true state becomes observable to all. The same game is repeated in the second stage. The agent can be of two privately known types: smart (type H) or average (type L). A smart agent and an average one differ not only in the *level* of signal quality, but also in the *slope* of signal quality improvement. A smart agent learns about the true state of the world with higher initial accuracy than an average one, but his signal quality improvement may be higher or lower than that of

[1] Throughout the paper, the receiver of information who then makes decisions ("the principal") is female and the sender who receives and transmits information ("the agent") is male.

an average agent. The agent is paid the expected value of his information, since the reports are assumed to be unverifiable. As a result, the agent has an incentive to improve his future reputation, i.e., how smart he is perceived to be in the second stage.

The first main insight emerging from the basic model is that mind changes, or inconsistent reports, may signal high ability in equilibrium. This can happen when a smart agent's signals improve faster than an average one's. Since the smart agent is more likely to receive and report an accurate first signal, an average agent might want to "defend" his early report even when he receives conflicting signals. Thus, similar to Prendergast and Stole (1996), an agent may stick to a position that he gradually realizes is likely to be wrong because changing his mind may make him appear incapable of finding the true state of the world earlier. However, unlike in some existing models (Scharfstein and Stein 1990, Prendergast and Stole 1996), an average agent is more likely to give consistent reports even though such consistency per se may indicate *low* ability in equilibrium.

The reason for this paradoxical result is that both consistency and accuracy matter in the determination of the agent's future wage, which is shown to be a convex function of his perceived ability. Being consistently right leads to the highest possible wage in the second stage, whereas being consistently wrong leads to the lowest. Thus, choosing to repeat one's early signal is a "gamble" to receive the highest wage. Since the type L agent's improvement in signal quality is smaller, he is more willing to take on this gamble. Intuitively, it takes confidence in one's information *improvement* to change one's mind and to admit an early mistake.[2] Even for the average agent, however, the final signal is more accurate. Thus if he gives consistent reports, he is more likely to be consistently wrong. Therefore, he will choose to give consistent reports only when the wage is sufficiently convex in the principal's perception of ability—which happens when her decision is very sensitive to the available information.

The second main insight of the model is therefore, when the principal's second stage decision depends strongly on the agent's reports, mind changes are valued more as a sign of ability. When the principal's optimal decision is independent of type, i.e., her decision only depends on the reports

[2] Some experimental and sociological evidence for increasing commitment to a wrong project is consistent with this prediction of the model. See for example Staw (1976, 1981, 1992), and the references within. Wicklund and Braun (1987) show that people who are more confident in their ability seem to be less committed to their early positions than the less confident ones.

she receives regardless of the agent's type, then in equilibrium consistent reports are valued more. This result suggests that mind changes are valued more in professions where optimal decisions are very sensitive to information accuracy.

The model so far demonstrates that truthful revelation in a sequential report setting may not be in the interest of an average agent due to reputational concerns. The principal, however, is only concerned with the accuracy of the reports and may want to choose a reporting protocol to encourage truthful revelation. One natural question is whether the principal should require sequential reports at all, given that the average agent may repeat his initial signal and thus fail to convey his latter, higher quality information. It would seem that requiring a single report (or a vector of reports) after the agent has received all signals is optimal because it eliminates the average agent's incentive to appear consistent.

An answer to the above question, and the third main insight of this paper is that which reporting protocol is optimal in term of truthful revelation depends on how sensitive the principal's decision problem is to the accuracy of reports. The advantage of a final report system is that the agent will only be judged on its accuracy, thus he will report his best estimate of the state (his final signal) truthfully. The disadvantage is that the agent ignores his still informative initial signal. Therefore if the principal's optimal decision depends only on which state of the world is more likely, a final report is optimal: the principal does not need to worry about the distortion in the average agent's final report. Moreover, a final report system is also preferable when the average agent's signals improve faster.

On the other hand, the sequential reports system is optimal when the exact likelihood of each state is crucial to the principal's optimal decision. The two reports (even though the final report may be distorted) offer the principal finer information and may lead to better decision than one truthful report under the final report system. Moreover, the sequencing of reports in the first stage (whether the reports are consistent or not) may provide a better estimate of the agent's type in the second stage.

It is important to emphasize, that this result hinges on the *timing*, not the *number*, of the agent's reports. Despite the seeming similarity, it is shown that the sequential reports system cannot be replicated by requiring two reports at the end. Under the sequential reports system, an agent always reports his initial signal truthfully, even though he may distort his final report to appear consistent. If the principal requires both reports at the end, then the agent simply repeats his final signal, which is the agent's best estimate of the state, in order to appear both consistent and accurate. As a result, his first signal is lost in equilibrium, just like when one final report is required.

Finally, an extension is considered in which the agent gives one additional report based on a new (third) signal. For the principal, the direct effect of the third signal is always positive due to its additional informativeness. The indirect effect of the additional signal on truthtelling incentives, however, is the main focus of the three signal model. The fourth insight of the paper concerns a tradeoff emerging from the longer sequence of reports: early commitment to a position versus late commitment.

On one hand, the three signal model shows that, the principal may want to request the third report when the improvement in the smart agent's signal quality levels off. The reason is that in this case, an average agent may lie against his true third signal in equilibrium if he has lied against his second signal to appear consistent. This "escalation effect", however, improves the average agent's incentive to tell the true second signal. Intuitively, an average agent wants to appear less consistent than he would have in a two signal model because he may have to lie more in the next report, and thus suffer from a big loss in accuracy. On the other hand, the principal may not want to require the last report when the smart agent's signal quality improves a lot in the final signal but the average one's does not. The reason is that in equilibrium, an average agent reports his true final signal if he has lied in his second report to appear consistent. While improving accuracy of the final report, the possibility of reporting truthfully later worsens the truthtelling incentive in his second report. Intuitively, an average agent wants to appear more consistent than in the two signal model because he can change his mind later and appear accurate. This negative indirect effect may outweigh the information provided by the third signal.

Previous research has shown that in a multi-agent setting, economic agents may want to be consistent with some early movers or existing consensus because they want to increase the market's perception of their ability (Scharfstein and Stein 1990).[3] In reputational herding models such as Scharfstein and Stein (1990), an agent wants to conform to the early mover because of a "smart people think alike" effect: smart agents receive signals that are correlated conditional on the state. Controlling for any information learned from the earlier mover, then if smart agents' observations are independent conditional on the state, each agent will report according to his own signal and there will be no reputational herding or the incentive to appear consistent. Here, both reports are associated with the agent and thus reputational concerns distort his reports even when signals are conditionally independent.

Ottaviani and Sorensen (2001) analyze a static reputational cheap talk game with very general distributions of the state and the agent's (expert's) type. They find that full revelation, or truthtelling is generically impossible in this type of game. Either no informative equilibrium exists, or the expert can only communicate part of their information, for example, "high" or "low" despite a rich signal and message space. The current model adopts simple distributions of the agent's type and the state to zoom in on the dynamic aspect of the agent's incentive problem. That is, the focus is on the interaction between the agent's interim report and future report given his reputational concerns.

More closely related to this paper, Prendergast and Stole (1996) consider a reputational concerns model in which a manager with privately known ability receives noisy signals about the true profitability of his investments over time, and the more capable manager receives signals with higher precision. In their model, the market infers each manager's precision from the period to period change in his investment choices. Initially, large changes indicate high quality information and therefore high precision relative to the prior, and each manager exaggerates out of reputational concerns. But eventually changes in investment indicate (many) past errors and everyone becomes too conservative. Therefore, exaggeration is beneficial only because the agent has no reputational stake in the prior, and "admitting" that a previous investment choice was bad always hurts reputation. One of the main insights of the present paper is that due to improvement in signal quality, admitting a previous mistake can indicate high ability in equilibrium. More generally, a goal of this

[3] Another reason is that they incorporate the information contained in the earlier actions before making their own decisions (statistical herding, see Banerjee (1992)).

paper is how such improvement affects incentives to report truthfully. Finally, this paper is also interested in what type of reporting protocol can elicit the most truthful reports given different types of agent's signal structure.

The paper proceeds as follows: Section 2 sets up the model and discusses important assumptions. Section 3 characterizes the equilibria and shows that, with improving signal quality, mind changes may signal high ability in equilibrium. Section 4 explores theoretically when the principal may make better decisions using sequential reports. Section 5 introduces an extension while section 6 concludes. All proofs are collected in Appendix A.

2 The Two Signal Model

In the basic model of sequential reports, a principal needs to make a decision based on an agent's signals. Although the model is clearly more general, this paper will couch it in a concrete story: the owner of a company needs to make an investment decision after reviewing a consultant's initial report (m_0) and final report (m_1) on a project's profitability. The profitability depends on the true state of world s, which is ex ante good or bad ($s \in \{g, b\}$) with equal probability. It is easiest to equate state with profitability: no investment yields zero, while investment brings profit g and b (net of investment cost) when the state is g and b respectively.[4] Moreover, it is assumed that the principal does not invest without further information, i.e., $g + b \leq 0$.

Before setting up the sequential reports game more formally, it may be useful to provide some real world examples of the types of situations this model describes. First, in an application to stock markets, an analyst (the agent) receives multiple pieces of information about a company over time and releases multiple stock recommendations. Eventually the company's true profitability becomes known and the investors (the principal) make a judgment about the analyst's ability. Second, in an application to the political arena, a politician (the agent) announces a policy initiative according to his private information. He receives new information and needs to decide whether to maintain the initiative or to change course. The voters (the principal) need to decide whether to support such an initiative. Later the truth becomes observable and the politician can be disciplined by elections.

[4] The nontrivial case is when $g > 0$, $b < 0$.

2.1 Environment and Information

The agent works in two stages $N = 0, 1$.[5] In each stage, the true state of the world s is, independently, either good (g) or bad (b). Events *within stage* 0 proceed as follows:

- At $t = 0$: the agent gets a fixed wage w_0 and then receives his first signal $i_0 \in \{g, b\}$;
- At $t = 0.5$: the agent sends an initial report $m_0 \in \{g, b\}$ as to which state his initial signal indicates;
- At $t = 1$: the agent receives his second signal $i_1 \in \{g, b\}$;
- At $t = 1.5$: the agent sends a final report $m_1 \in \{g, b\}$ as to which state his second signal indicates;
- At $t = 2$: the principal makes the investment decision $a \in \{0, 1\}$ based on the reports;
- At $t = 2.5$: the true state of world becomes observable to all but not verifiable.

Stage 1 repeats the above process: the agent receives a fixed wage at the beginning. Later, he delivers two sequential reports, the principal makes her investment decision and the game ends. Timing of this game is illustrated in Figure 1. Notice that even though in the timeline, the principal only decides after the agent's reports, it is not crucial to the results of the model as long as she can observe the second report. In the real world examples above, the investors or the voters may take some action based on the early report. But as long as they can observe the final report, incentives similar to the current model's would arise because the principal can still use both reports and, later, the realized state to evaluate the agent.

The agent is one of two types: $\theta \in \{H, L\}$. An agent is smart (type H) with probability η and merely average (type L) with probability $1 - \eta$. While the distribution of the state and that of the agent's type are common knowledge, only the agent knows his type. The agent receives two private

[5] Some career concern models such as Scharfstein and Stein (1990) employ a reduced form second stage in which the agent's wage is his posterior probability of being talented. Modeling two full stages, however, makes it possible to study explicitly the *shape* of the agent's wage in the second stage, which influences the agent's truthtelling incentives in the first period. See Lemma 1 for characterization of the wage function in the second stage.

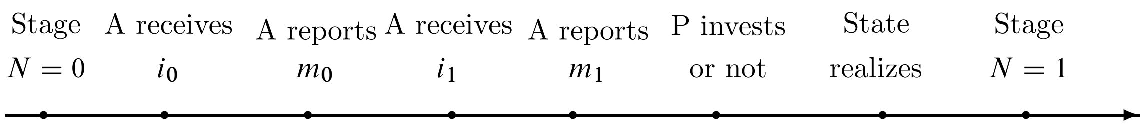

Figure 1: Timeline

signals about the underlying true state. The signals an agent receives are independent conditional on the state. The qualities of these signals depend on the agent's type:

$$Pr(i_1 = s|H, s) = p_1 \geq Pr(i_0 = s|H, s) = p_0 > \frac{1}{2}; \quad Pr(i_1 = s|L, s) = r > Pr(i_0 = s|L, s) = \frac{1}{2}.$$

This specification means that, first, both types of agent receive a more accurate second signal.[6] Second, combined with the symmetric distribution of the state,[7] the assumption guarantees that the initial signal i_0 itself is not informative about ability, i.e., $Pr(i_0 = g|H) = Pr(i_0 = g|L) = \frac{1}{2}$. Thus it enables the analysis to focus on the dynamic incentive problems due to the improvement in signal quality, because in equilibrium the agent is not tempted to lie in his first report.

2.2 Payoffs

The principal and the agent are risk neutral, but the principal cannot transfer the ownership of the project to the agent (e.g. due to credit constraints). Let m^N be the history of reports in stage N. Also let $\hat{\eta} \equiv Pr(H|m^0, s)$ be the principal's posterior estimate of the agent being type H, given his first stage reports as well as the observed state. Let Π^N be the stage N profit, then:

$$\Pi^0 = \sum_{m^0}[Pr(g, m^0)g + Pr(b, m^0)b]a(m^0; \eta); \quad \Pi^1 = \sum_{m^1}[Pr(g, m^1)g + Pr(b, m^1)b]a(m^1; \hat{\eta}).$$

The principal is assumed to choose action $a : m^N \to \{0, 1\}$ to maximize her net expected profit: $E\Pi = \Pi^0 - w_0 + E[\Pi^1 - w_1|m^0, s]$.

The agent cannot be paid conditional on the accuracy of reports because the true state of the world is assumed to be unverifiable (though observable), and as is standard in the cheap talk liter-

[6] The assumption that type L's first signal is completely uninformative simplifies the analysis. As long as type H's initial signal is more accurate, the main results hold with slight modifications.

[7] Allowing asymmetric state distribution introduces potential lying in the agent's first report in addition to the dynamic incentive problems. For example, when state g is much more likely than state b, the smart agent is more likely to observe state g because his initial signal is more accurate. This gives type L an incentive to report $s = g$ with some probability even when his first signal is b. This problem is similar to Prendergast (1993).

ature, no contract can be written on the reports (messages), therefore the agent is only motivated by his wage in the second stage. Assume the principal operates in a perfectly competitive environment, the agent's second stage wage is simply the expected value of his information conditional on the principal's updated belief of him being type H. Thus all rents accrue to the agent if he is perceived as talented and the principal's maximization problem simplifies to maximizing her first stage profit.[8]

Assume that the agent reports truthfully in the second stage (later shown to be part of an equilibrium). Let $a^*(m^1, \hat{\eta})$ denote the principal's optimal action given $m^1, \hat{\eta}$, then the wage of the agent is $w(\hat{\eta}) = \Pi^1(m^1, \hat{\eta})|_{a=a^*} \equiv V(a^*(\hat{\eta}))$.[9] Moreover,

Lemma 1 *(1). $w(\hat{\eta})$ is a convex and piecewise-linear function of $\hat{\eta}$, the posterior probability that the agent is smart; (2). $w(\hat{\eta})$ is linear in $\hat{\eta}$ if the principal's optimal action $a^*(m^1, \hat{\eta})$ is independent of the agent's posterior ability.*

As shown by Blackwell (1953), the value function of the agent's information is convex in the principal's beliefs about the agent's type. The reason is that the principal can make better (and potentially different) decisions given two different posterior distributions of the agent's type, than she can if constrained to make the best decision given a convex combination of these two type distributions. Intuitively, imagine that the agent's type is known in the second stage, then for any given report sequence, the principal can choose the most profitable action given the agent's type. Thus she can do no worse than if she has to choose an action knowing only the agent's type distribution. In this model, due to the binary state distribution and the binary signal structure, the wage function is also piecewise linear. The exact shape of the payoff function depends on the difference of signal quality between types as well as the project-specific values g, b.

Example 1: A simple convex payoff function. Suppose that state $s = b$ is sufficiently bad that the principal is only willing to invest if she believes that $s = g$ is very likely.[10] Then the payoff

[8] As a result, the principal is only concerned about maximizing her first stage profit in choosing optimal reporting protocol. This is discussed more in section 4.

[9] The agent's wage is the value of his information over what the principal would obtain by default, which is zero because his optimal decision without further information is assumed to be no investment.

[10] Formally, this requires $\frac{r}{2}g + \frac{1-r}{2}b \leq 0$, and $(1 - p_0)p_1 g + p_0(1 - p_1)b > 0$. The first inequality means that report sequence (b, g) from type L is not good enough news about the state to warrant investment, while the second inequality means that the same sequence from type H is.

function takes the form of a kinked function:

$$w(\hat{\eta}) = \begin{cases} 0, & \text{if } \hat{\eta} < \eta_1 \\ \tau_1(\hat{\eta} - \eta_1), & \text{if } \hat{\eta} \in [\eta_1, \eta_2) \\ \tau_2(\hat{\eta} - \eta_2 + \eta_1), & \text{if } \hat{\eta} \in [\eta_2, 1] \end{cases}$$

where $\tau_1 < \tau_2$. The implicit incentive system here is straightforward: the agent is "fired" at the end of the first stage if $\hat{\eta} < \eta_1$. If $\hat{\eta} \geq \eta_1$, then he is retained and his wage depends on which segment $\hat{\eta}$ falls in: he gets either a good or a star wage. Intuitively, even the best possible news from an agent who is very likely average is insufficient to change the principal's decision from no investment (the default) to investment, while good news from an agent quite likely to be smart may induce her to invest and get higher expected profit. □

The second part of the lemma shows that $w(\hat{\eta})$ is linear when the principal's optimal action depends only on the reports she receives regardless of type. This occurs when $gr + b(1 - r) = 0$ such that report sequences (g, g) or (b, g) from an average agent yield expected profit of exactly zero, thus the principal will invest if there is *any* probability that the agent is smart and the reports are positive. In this case, $w(\hat{\eta}) = (g - b)(p_1 - r)\hat{\eta}$.

Albeit simple, this lemma shows that the reduced form approach used in many reputational concerns models, where the agent maximizes the posterior probability that he is smart because his future wage is linear in such probability, is a special case (Scharfstein and Stein 1990, Prendergast and Stole 1996). Such a reduced form approach implicitly assumes that the principal's future decision problem is not very sensitive to the agent's forecasting accuracy. One economic implication, to be explored partially below, is that in professions where key information is provided by experts driven primarily by reputational concerns, the implicit incentive itself may be convex. Therefore even if the agents themselves are risk neutral, the implicit incentive structure encourages risk-taking behaviors. Moreover, the higher are the premiums on the accuracy of expert's advice, the more convex the implicit incentive system becomes and more risk may be taken on the part of experts.

2.3 Equilibrium

The principal needs to infer the true signals the agent has received and to update her belief about the agent's type from his reports. Thus the principal's decision depends on each type of agent's strategy and the agent's strategy depends on her inference. The ensuing analysis adopts the concept of Perfect Bayesian Equilibrium (PBE), in which the agent's strategy is a function that maps his type, his signals, as well as the history of reports, if any, to new report(s). Thus, in the first stage, the strategy of the agent is $m_0 : \Theta \times I_0 \to \Delta(g,b)$ and $m_1 : \Theta \times I_0 \times M_0 \times I_1 \to \Delta(g,b)$. His strategy in the second stage is similarly defined. The equilibrium consists of a triple $(m^*, a^*, \hat{\eta})$ such that: $m^*(\theta, I) = \text{argmax}_m \; Ew(\hat{\eta})$; and $a^* = \text{argmax}_{a \in \{0,1\}} \; \Pi(a, m)$, where $\hat{\eta}$ is the principal's posterior belief that the agent is smart, given the agent's strategy. This belief is updated by Bayes' rule whenever possible.

Two well known equilibrium multiplicity problems exist in cheap talk games. First, there always exist "babbling" equilibria in which all messages are taken to be meaningless and ignored by the receiver.[11] Second, there exists an unimportant type of multiplicity of equilibrium: because the meaning of messages in cheap talk games is endogenously determined in equilibrium, any permutation of messages across meanings yields another equilibrium. Given these two problems, this paper restricts attention to *informative* equilibria in which both the principal and agent use and understand the *literal* meaning of the reports, but whether they think the reports are credible depends on the equilibrium strategies (see also footnote 13).[12]

3 Equilibrium Information Revelation

Section 3.1-3.2 categorize the equilibrium strategies of the principal and the agent, focusing on how information revelation depends on both the initial difference and the improvement in the types'

[11] Papers such as Farrell (1993) argue that the babbling equilibria are frequently implausible, especially in games with some common interest. In an evolutionary setting, Blume, Kim and Sobel (1993) show that the babbling equilibrium is often unstable in the long run. However, in some settings (Morris (2001)), the babbling equilibrium is the only equilibrium because no information can be transmitted due to incentive problems.

[12] Myerson (1989) and Farrell (1993) show that the second type of multiple equilibria disappears if we separate the meaning of a message from its credibility. That is, in a rich language such as English, both the sender and the receiver use the literal meaning of a message but may not believe its content.

signal quality. Section 3.3 studies the case when the agent's type is symmetric information to illustrate that models in which the agent knows how smart he is (as in the current model) and models in which he does not (as in many existing models) yield very different predictions.

Observe that there always exists a truthtelling equilibrium in the second stage such that the agent of either type reports truthfully. The reason is that the agent's wage $w(\hat{\eta})$ does not depend on his second stage performance and he has no further reputational concerns because the second stage is the end of his career. Since there is no conflict of interest between the principal and the agent, it is assumed that this truthtelling equilibrium is always played in the second stage. What is interesting is the agent's equilibrium behavior in the first stage.

In the first stage, assume that both types of agents report the first signal i_0 truthfully (shown later to be part of an equilibrium strategy). Without loss of generality, the agent's continuation pure strategy after receiving i_1 is either: always report true i_1; or always repeat $m_0 = i_0$.[13] Observe that it cannot be an equilibrium for type H to always report i_1 truthfully and for type L to always repeat his first report regardless of i_1, or vice versa. Suppose so, then type H and type L can be distinguished perfectly on the equilibrium path when $i_0 \neq i_1$, in which case L has a strong incentive to deviate and pretend to be H. Therefore there can be at most three possible continuation equilibria: a "full revelation equilibrium" in which both types of agent report their second signal truthfully; a "full pooling equilibrium" in which both types simply repeat their initial report, and finally, a "partial revelation equilibrium" in which the agent plays a convex combination of the above two types of strategies.

3.1 Signal Quality Improvement and the Agent's Equilibrium Incentives

This subsection focuses on how the agent's incentive to report his second, more informative signal truthfully depends on his type and the signal quality improvement. Since both type H and type L agent receive signals of increasing quality, it is necessary to define a measure of signal quality

[13] By restricting attention to the literal meaning of messages, a lot of uninteresting equilibria are eliminated. For example, here the agent can use other strategies such as always reporting the opposite of i_0 or i_1. But it does not change the essence of the equilibrium if each type uses an opposite strategy because one can simply redefine i_0 or i_1. That is, suppose that there exists a full revelation equilibrium in which everyone reports the opposite of their true signals and the principal knows that the reports are the opposite of the signals. Such an equilibrium is equivalent to one in which everyone just reports the true signals.

improvement. A smart agent is considered to improve faster than an average one if the following condition holds:

$$\frac{1-r}{r} > \frac{p_0(1-p_1)}{p_1(1-p_0)}, \tag{1}$$

while an average agent is considered to improve faster if inequality (1) does not hold. The above inequality compares the confidence of an agent in his second signal *relative* to the first when the two signals disagree.[14] The left hand side of the inequality measures the probability ratio that a type L agent's second signal is wrong vs. his second signal is right and the right hand side is the same ratio for type H. When this inequality holds, type H trusts his second signal more than type L when he receives conflicting signals.

To begin with, consider a benchmark case when both types of agent report truthfully. A comparison of posterior probabilities that the agent is smart given his reports and the observed true state suggests that both the accuracy and consistency of reports indicate high ability:

$$Pr(H|i_0 = i_1 = s) \geq Pr(H|i_0 \neq s, i_1 = s) \geq Pr(H|i_0 = s, i_1 \neq s) \geq Pr(H|i_0 \neq s, i_1 \neq s).$$

Denote the above four posterior probabilities respectively as (CR), (R), (W), and (CW) such that CR stands for consistently right; R for a right change of mind; W for a wrong change of mind, and lastly, CW stands for consistently wrong. Then the above inequalities show that, given the correct final report, a change of mind is bad for agent's reputation because it means that he is wrong at the beginning; however, being consistently wrong is worse than a wrong change of mind because L is more likely to get two wrong signals in a row.

In the current model, however, the agent may not report truthfully due to their reputational concerns: both types want to appear smart for the principal. The following proposition characterizes the equilibria with sequential reports:

Proposition 1 *There exists a* $\overline{\eta} \in [\frac{1}{3}, 1)$ *such that when* $\eta < \overline{\eta}$,[15]

(1.1). When type H *improves relatively faster, there exists a cutoff value* p_0^L *such that for* $p_0 < p_0^L$*, a full revelation equilibrium exists. For* $p_0 \geq p_0^L$*, a partial revelation equilibrium exists. In this*

[14] Note that inequality (1) implies that type H's second signal is also better than that of type L: or $p_1 \geq r$.

[15] This condition guarantees the monotonicity of the mixing probability with p_0. In the most restrictive case, this requires that the fraction of type H does not exceed $\frac{1}{3}$. In other cases, for example, when $p_0 \approx 1$, $\overline{\eta} \approx 1$ and the restriction is trivial.

equilibrium, $m_0 = i_0$ for both types of agent. In the second report, type H always reports truthfully. Type L reports truthfully if $i_0 = i_1$, but repeats m_0 with probability $\pi^* \in (0, 1)$ if $i_1 \neq i_0$. Moreover, π^* increases with p_0.

(1.2). When type L improves relatively faster, there exists a cutoff value p_0^H such that for $p_0 < p_0^H$, a full revelation equilibrium exists. For $p_0 \geq p_0^H$, a full pooling equilibrium exists in which both type H and L report $m_0 = m_1 = i_0$.

Proposition 1 shows that, first, when p_0 is relatively low, the agent tries to deliver an accurate final report, which is quite important to the principal's updated belief of the agent's type. When the agent's signals disagree, the more he believes in his later (and better) signal, the less attractive repeating his first report becomes. Intuitively, lying and repeating the first report is likely to lead to a consistently wrong sequence of reports, which yields the lowest reputational payoff. Therefore, when p_0 is sufficiently close to $\frac{1}{2}$, both H and L have (almost) uninformative first signals and the final report is the key indicator of ability, and both types of agent will report their second signal truthfully. *Ceteris paribus*, the faster an agent's signals improve, the more value he attaches to the second signal because it is "better late than never": he can look like H who is unlucky in the first signal but finds out about the true state later.

Second, the higher quality the smart agent's first signal is, the more likely the average agent prefers repeating his first report. The reason is that as p_0 becomes higher, a correct first report is increasingly more likely to reflect high ability. As a result, type L is tempted to repeat his first report with a positive probability after receiving conflicting signals to appear smart when type H's first signal is quite accurate. Consider an extreme example where type H's first signal is perfect ($p_0 = 1$), then regardless of a type L agent's second signal, he repeats his first report with probability one. Any mind change shows that he is type L for sure, while repeating his first report make him appear smart with some probability.

More subtly, however, higher relative improvement in signal quality is crucial for type H agent to report truthfully. It is not sufficient that type H receives better signals than type L in *absolute* terms. Rather, a smart agent is more truthful only when he improves faster as defined by inequality (1). For example, suppose $p_0 \approx p_1$ and the signals differ, type L believes that his second signal is correct with probability r, which is larger than $\frac{1}{2}$, the approximate probability H places on his second signal being correct. In this case, type H, despite the fact that *both* his signals are more accurate than type L, has less relative confidence in his second signal and is thus more tempted to repeat his first report. Once type H repeats his first report, type L will imitate, therefore in equilibrium both types repeat their initial report, as described in part 3 of Proposition 1.

How does this model relate to models without improvement in signal quality? Suppose that the agent's initial signal is exactly as accurate as his second one, then both types of agent have the same estimate of the true state after receiving conflicting signals: the two signals exactly offset each other (that is, $Pr(g|i_0^g, i_1^b; \theta) = \frac{1}{2}$). Therefore reporting the true second signal does not increase his probability of giving a correct final report and, in turn, the principal's posterior that he is smart. Hence type H is better off repeating his first report so that he may obtain $w(CR)$ with probability $\frac{1}{2}$. Therefore there is no need to require the second report.

This partial revelation equilibrium may be highly inefficient: the principal's information may deteriorate significantly even if the probability that the agent is smart is exceedingly small. The reason is that an average agent may repeat his first uninformative report with a high probability to appear smart despite a high quality second signal. Consider the following example:

Example 2: One good apple may ruin the barrel. Suppose that $\eta = 0.001, p_1 = 1, r = 0.9$ and $w(\hat{\eta}) = \hat{\eta}$. That is, the agent is extremely likely to be average, and type L's second signal is very accurate. In equilibrium, however, type L agent repeats his first report with probability $\pi^*(p_0) = 9.982p_0 + 0.078p_0^2 - 9$. Clearly, π^* increases in p_0. When $p_0 = 0.95, \pi^* = \frac{1}{2}$. Hence a type L agent lies against his highly informative signal i_1 and uses his totally uninformative signal i_0 half of the time, even though the prior probability he is smart is only one out of a thousand. □

3.2 Value of Inconsistent Reports

This section focuses on whether the sequencing of the reports alone may carry information about the agent's type, before he can be judged based on the accuracy of his reports, i.e., before the true state of the world in the first stage becomes observable. Formally, the market is considered to value consistency more if $Pr(H|m_0 = m_1) > Pr(H|m_0 \neq m_1)$ and to value mind changes more otherwise. This is interesting because a major insight of the reputational herding models is that consistency (with early movers or existing consensus) is valued by the market as a sign of talent.

Proposition 1 shows that the principal may receive consistent reports because the average agent pretends to be consistent despite conflicting signals. The following proposition describes when consistent reports signal higher ability and when mind changes do.

Proposition 2 *In the first stage, suppose that there exists a partial revelation equilibrium in which type L reports $m_1 = m_0$ with probability π^* when he receives conflicting signals while type H always report truthfully. Then, the following is true:*
(2.1) the market values consistency more than mind changes before observing the state when $\pi^* \leq (2p_0 - 1)(2p_1 - 1)$, *which occurs when* $w(\hat{\eta})$ *is linear.*
(2.2) the market values mind changes more than consistency before observing the state when $\pi^* > (2p_0 - 1)(2p_1 - 1)$. *This occurs when the payoff function* $w(\hat{\eta})$ *is sufficiently convex.*

Figure 2: Partial Information Revelation Equilibrium with a Convex Payoff Function

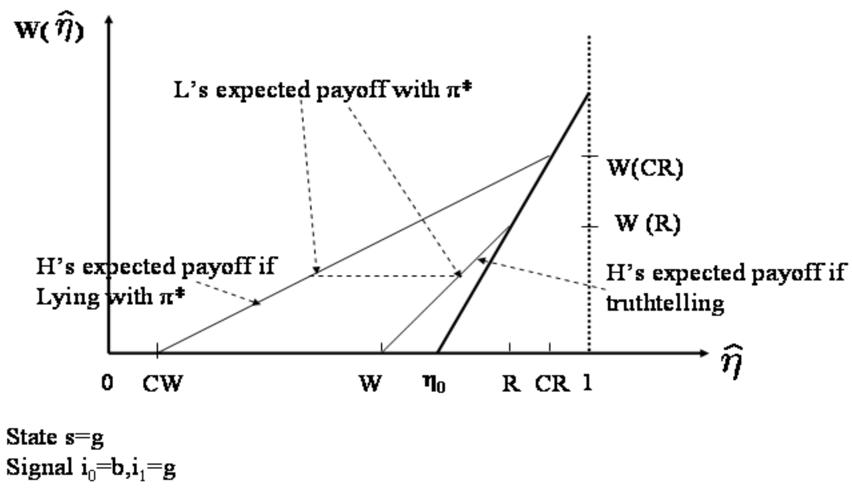

The second part of the above proposition may appear counterintuitive: if the principal does not value consistency in equilibrium, then there seems to be no reason for a type L agent to lie against his second, more informative signal to appear consistent. Instead, a type L agent should simply tell the truth when he receives conflicting signals. However, when the principal values highly accurate reports disproportionally more than somewhat accurate ones, the wage function is very convex. As can be seen from Figure 2, for a type L agent who receives inconsistent signals, repeating his first report leads to the best future wage $w(CR)$ with probability $1-r$ or the worst future wage $w(CW)$ with probability r. If he follows his true second signal and gives inconsistent reports, he receives $w(R)$ or $w(W)$ with probability r and $1-r$ instead. Although consistent reports are more likely to be wrong, the cost is relatively small given type L's lack of confidence in the second signal relative to type H. The benefit of consistent reports, however, is that they are riskier than mind changes and can give type L higher reputational payoff in expectation. The smart agent, on the other hand, reports his second signal truthfully because he has a different probability distribution over outcomes CW, W, R and CR. Namely, his second signal is so much better than his first that repeating his first signal is very likely to lead to the worst payoff of all: being consistently wrong.

Example 3: Market values mind changes. Type L's equilibrium mixing probabilities change when the payoff function becomes increasingly convex. Let $\eta = 0.1$, $p_0 = 0.8$, $p_1 = 0.9$ and $r = 0.6$. For simplicity, use $w(\hat{\eta}) = \hat{\eta}^{\alpha}$ as an approximation for the agent's wage function, then the following table shows that when $\alpha = 5$, the market values mind changes more because the equilibrium mixing probability $\pi^* = 0.49$ is higher than 0.48, the cutoff value given in Proposition 2.

Convexity: α	1	1.6	2	3	5
Mixing probability: π^*	.32	.39	.42	.46	.49

Moreover, notice that the principal may make better investment decisions in the first stage by updating her beliefs of the agent's type, using the sequencing of his reports. Recall Example 1: when no report sequence from an average agent can convince the principal that $s = g$ is sufficiently likely to invest, the second stage wage is a piece-wise linear function with two kinks. Then, if type H's signals improve faster and $p_0 \geq p_0^L$, an average agent may repeat his first report with such a high mixing probability π^* that mind changes signal higher ability and $Pr(s = g|b, g) > Pr(s = g|g, g)$. Using the parameter values given above, simple calculations can show that at $\alpha = 5$ and $\pi^* = 0.49$, $Pr(s = g|b, g) \approx 0.62 > Pr(s = g|g, g) \approx 0.47$. As a result, the principal may invest when she hears reports (b, g), which is more likely to reflect true signals from a smart manager, but not to invest otherwise. □

3.3 When Ability is Symmetric Information

The model above shows that the smart agent reports more truthfully because he is more confident in the improvement in his signal quality. In some professions (or stages of one's career), the agent may not have superior knowledge of his ability. Turning to a two signal model with symmetric information, this subsection shows that the premium on mind changes is not only due to type H's faster improvement in signal quality, but also to his private knowledge of ability.

All the modeling assumptions remain the same except that now both the principal and the agent only know the agent's type distribution. Thus the agent needs to infer his own type from his signals before sending his final report.[16] That is, his private belief of the state, $Pr(s|i_0, i_1)$, depends on his updated belief of how smart he is given the signals. Assume $w(\hat{\eta})$ is linear to simplify analysis. The findings when ability is symmetric information are summarized below:

Proposition 3 *When $\eta \leq \overline{\eta} \in [\frac{1}{3}, 1)$, there exists a $\hat{p}_0$ such that:*

(3.1) If $p_0 < \hat{p}_0$, there exists a full revelation equilibrium in which the agent always reports $m_0 = i_0, m_1 = i_1$. Moreover, $\hat{p}_0 > p_0^L$ if type H improves faster, and $\hat{p}_0 > p_0^H$ if type L improves faster.

(3.2) If $p_0 \geq \hat{p}_0$ and η is not close to 0, then there exists a full pooling equilibrium in which the agent always repeats his first report.

(3.3) The market always values consistent reports more than mind changes in equilibrium, i.e., $Pr(H|m_0 = m_1) > Pr(H|m_0 \neq m_1)$.

First, there exists a cutoff point $\hat{p}_0$ such that for all $p_0 < \hat{p}_0$, the agent always reports truthfully. The intuition is similar to the full information revelation equilibrium in Proposition 1: when $p_0 \approx \frac{1}{2}$, type H's first signal is not very accurate and conflicting signals are relatively frequent for both types, thus giving a correct final report is much more important in judging one's ability. However, for any given parameters, the cutoff point $\hat{p}_0$ is larger than the corresponding cutoff value in the asymmetric information case. The reason is that if the agent (of either type H or L) receives inconsistent signals in the symmetric information case, he believes that with some probability he

[16] Recall that his first signal is uninformative about type by assumption.

is average and his second signal is not too accurate, and with some probability that he is smart and the second signal is very accurate. Therefore he has more incentive to report the second signal truthfully than in the asymmetric information case.

Second, when $p_0 \geq \hat{p}_0$, the expected payoff from consistent reports becomes too high for the agent to report truthfully. With symmetric information, however, this means that both types of agent repeat the first report with some probability. Lying of the type H agent further increases the payoff from consistent reports until the agent repeats his first report with probability one. Thus the most important insight from the symmetric information case is that consistent reports signal high ability in equilibrium. The key to understand this result is that both the principal and the agent himself believe that H is more likely to be consistent. In fact, $Pr(H|m_0 \neq m_1) \leq \eta$ for all p_0. In contrast, it is easy to see that η is the minimum probability that the agent is smart conditional on consistent reports when both types repeat their initial report.

This subsection shows whether consistency or mind changes is more valuable as a sign of talent also depends on how well the agent knows his own ability. In professions where one's talent is unknown to all parties, consistency is more valued. High quality information and fast improvement are not enough to ensure that a smart agent changes his mind and acknowledge a (likely) early mistake: the agent needs to know how good he is.

4 Optimal Reporting Protocol

Section 3 shows that when the agent's signals exhibit different levels of improvement, requiring sequential reports may induce an average agent to lie and report too consistently. Example 2 shows that the inefficiency due to type L's reputational concerns can be quite significant. A frequently observed alternative is for the principal to require report(s) *after* the agent has received all of his signals. This section compares these two reporting systems by first investigating incentives generated when the principal requires a final report. Based on this, it shows when and which of the reporting systems enable the principal to make better decisions.

Recall that in the second stage, the agent receives the full expected value of his information. Therefore the optimal reporting protocol here is one that elicits the most truthful reports and leads to the highest first stage profit for the principal.[17] The principal can require one final report or a final report sequence after the agent has learned both signals. Formally, she may ask for one final report m^f or a vector of final reports $\vec{m}^f = (m_0, m_1)$ on the signals received. When one final report is required, the agent will be judged solely on its accuracy, which in turn determines his second stage wage. Clearly, the agent should report his best estimate of the state based on the signals. When a vector of final reports is required, however, the results are more subtle:

Proposition 4 *(4.1). If the principal requires one report m^f, then in equilibrium both types of agent report $m^f = i_1$, regardless of their first signal.*

(4.2). If the principal requires report $\vec{m}^f = (m_0, m_1)$ after the agent receives both signals, and if type H's signal improves faster (inequality (1) holds), then there exists an equilibrium in which both types of agent report $m_0 = m_1 = i_1$.

The first part of Proposition 4 shows that the principal receives the true final signal when she requires m^f only. The reason is straightforward: the agent should use his higher quality second signal and report $m^f = i_1$, which leads to a higher posterior estimate of his ability in expectation than reporting $m^f = i_0$. The final report m^f, however, is not a sufficient statistic of the agent's signals because both signal sequences (b, g) and (g, g) lead to the same report, whereas the principal forms very different opinions of the true state given these two signal sequences. A natural alternative is to require a vector of report at the end, $\vec{m}^f = (m_0, m_1)$, about the signals the agent received.[18] Despite some seeming similarity, however, the second part of Proposition 4 shows that requiring $\vec{m}^f$ at the end differs markedly from the sequential reports model in section 3.

The key difference is that inducing truthful initial report becomes more difficult in this case: the agent can send any reports that gives him highest reputational payoff in the next stage, knowing both signals.[19] Intuitively, the agent has little incentive to lie in his initial report in the sequential

[17] If the market is not perfectly competitive or when the principal can enter a multi-stage contract with the agent, the principal may face a tradeoff of the first and second stage profit. For example, she may choose a reporting protocol that does not lead to the highest first period profit, but gives a more precise estimate of the agent's type to increase her second stage expected profit.

[18] Note that if the principal asks the agent for a final report with probabilistic assessment of the state, there can only be four different distributions determined by the agent's signal structure. Thus requiring probabilistic final report is equivalent to requiring a vector of final reports.

[19] Formally, given a signal sequence, the agent can deliver any of the three report sequences other than the true signals. The incentive constraint in the sequential reports model is only one of them.

reports system, which to some extent, commits him to a given position. In the current case, without the early report, he can and will modify his report sequence in any way to increase the probability of being perceived smart.

As the second part of the above proposition shows, full revelation cannot be an equilibrium in this model, even for the parameter values such that a full revelation equilibrium exists in the sequential reports case ($p_0 < p_0^L$). Part of the reasoning is familiar: if there were a full revelation equilibrium, consistent reports would signal higher ability, and the agent would deviate and report more consistently. In the sequential reports system, type H agent reports inconsistent signals truthfully because his second report is more likely to be accurate and the benefit of a correct final report outweighs any cost of inconsistency. In the current case, there is no need for type H to send inconsistent reports: by reporting $m_0 = m_1 = i_1$, he is likely to be accurate as well as consistent. Since type H agent has more relative confidence in his second signal than type L, he is more confident that his consistent reports are likely to be correct and he will receive the highest posterior $Pr(H|m_0 = m_1 = s)$. As a result, type L agent gives more consistent reports too. Thus in equilibrium, both types report their second signal.

Proposition 4 shows that, in equilibrium, the final report system (requiring either m^f or $\vec{m}^f$) leads to truthful revelation of the second signal only. The following proposition compares the final report system with the sequential reports system, where the first report is truthful but an average agent may repeat his first uninformative report to appear smart.

Proposition 5 *(5.1) When type L improves faster (inequality (1) does not hold), and $p_0 \geq p_0^H$, the principal should require one final report m^f.*

(5.2) When η is sufficiently close to 0, $p_0 \geq p_0^L$ and type H improves faster (inequality (1) holds), the principal should require one final report if p_0 is sufficiently high and/or $w(\hat{\eta})$ is sufficiently convex.

(5.3) When p_0 is sufficiently close to $\frac{1}{2}$, the principal should require sequential reports and both types of agent report $m_0 = i_0, m_1 = i_1$.

(5.4) When type H improves faster (inequality (1) holds), $p_0 \geq p_0^L$, η is not too close to 0, and the principal needs highly precise reports ($Pr(s = g|m)$ sufficiently large), the sequential reports

system leads to better decisions in equilibrium than the final report system.

First, the final report system eliminates the average agent's incentive to appear consistent and thus elicits a true and higher quality signal i_1. Therefore it may seem that requiring one final report is always better than requiring sequential reports. Part one of Proposition 5 shows that this intuition is correct in the case when the average agent's signals improve faster, even though the smart agent receives signals of higher quality. In this case, as shown in Proposition 1, there exists a full pooling equilibrium in which the agent always reports $m_0 = m_1 = i_0$. Given the improvement in signal quality, requiring one final report elicits the true i_1, which is more informative than $m_0 = m_1 = i_0$. Thus when the average agent actually improves faster (when inequality (1) does not hold), the principal should only require a final report.

Moreover, the principal should choose a final report system when the inefficiency of the sequential reports system is sufficiently high. In the sequential reports system, when p_0 is very high or when the wage function $w(\hat{\eta})$ is sufficiently convex, an average agent repeats his first report with a high probability, as illustrated in Example 2. Thus even though a talented agent reports truthfully, the principal is much more likely to receive an uninformative report from the agent who is very likely to be average ($\eta \approx 0$). On average, she makes better decision with the final report system.

Next, Proposition 5 shows that the sequential reports are more valuable when the decision problem is very sensitive to accuracy of the reports, which in turn depends on the agent's type and the truthfulness of the reports. The reason is that the final report system does not, and cannot contain all the information revealed in the sequential reports system because it fails to convey how strongly the agent believes in the state he reported. That is, the final report system gives a less fine estimate of the true state. This poses a problem when the principal's optimal decision depends strongly on the agent's probabilistic estimate of the states.

To see this, suppose that the principal needs to be convinced that $s = g$ with such a high probability that she would not invest under the final report system even if $m^f = g$.[20] Is there any report sequence from the sequential reports system that leads her to invest? It can be shown that, on one hand, $Pr(s = g|g, g) > Pr(s = g|g)$ when the average agent does not lie with too

[20] Formally, when $[p_1\eta + r(1-\eta)]g + [(1-p_1)\eta + (1-r)(1-\eta)]b \leq 0$.

high a probability in his second report. Intuitively, the probability that the average agent lies against his true signal $i_1 = b$ is outweighed by the probability that the smart agent receives two positive signals. Thus the principal may invest instead. On the other hand, when p_0 is high and the mixing probability π^* very high, then $Pr(s = g|b, g) > Pr(s = g|g)$. This is the case described in Example 3: in the sequential reports system when state b is sufficiently bad, the principal may invest when she hears (b, g) and not invest otherwise.

The above proposition illustrates that the sequential reports system may offer more precise reports than the final report system and change the optimal decision of the principal. Therefore a principal may prefer sequential reports even when no action needs to be taken based on the agent's interim report. More generally, this section demonstrates that when truthful information revelation is the main concern, the principal may benefit from observing information gradually because the path of reports reveals finer information. With a final report system, the agent can fabricate early reports given his later information, and potentially lead the principal to believe too strongly in a certain direction. This is particularly important when the principal's decision is very sensitive to the exact accuracy of the reports.

5 Extension: Role of Additional Informative Report

The two signal model analyzed so far suggests that sequencing of reports provides information about the agent's ability. Sequencing of reports is also important in this type of multi-period incentive and learning model because the agent's incentive to report truthfully may depend very much on the path of his previous reports as well as all the possible future ones. Consider the scenario when the agent receives more than two signals, should the principal ask for an additional report? Does the agent become more entrenched in his position in the third report or more truthful in order to give a correct final report? This subsection adapts the previous model to a three signal setting to illustrate a new tradeoff the agent faces: to appear consistent early to show he has confidence in his early signals or to appear consistent late so that his reports are more likely to correct. It also shows that the principal should take into account this new tradeoff in deciding whether to require a new report, which may induce counterintuitive effect on the agent's overall incentives.

Since the agent's final signal is informative, the *direct* effect of an additional report necessarily improves the principal's investment decision. Moreover, the more informative the third signal is relative to the previous signals, the more valuable it becomes. However, due to reputational concerns, the *indirect* effect on the agent's incentives is the focus of this section. Namely, requiring the third report may change the agent's incentive to lie in the early reports, and in turn affect the informativeness of the third report. Intuitively, when a third report is required, the agent's second report can be thought of as an option. If the addition of a third report increases the option value of consistency in term of the agent's reputational stock $Pr(H|m_0, m_1)$, the probability that he is considered smart given his first two reports, then the agent wants to lie more in the second report; when it narrows down this option value, the agent would prefer to be more truthful in the second report.

Suppose that type H's final signal i_2 is perfect and m_2 is the last report the agent submits. Let $r_t \equiv Pr(i_t = s|L), t = 0, 1, 2$, and $r_0 < r_1 < r_2$. The parameters are restricted in the following way to simplify the analysis: 1). $p_0 \geq p_0^L$, $r_0 = 1/2$, $\frac{p_0(1-p_1)}{p_1(1-p_0)} < \frac{1-r_1}{r_1}$, and 2). $r_2 - r_1$ not too large.[21] These assumptions serve two purposes. First, it guarantees truthtelling by type H, thus makes it possible to focus on the average agent's incentives. Second, it allows comparison with a two signal model with partial revelation equilibrium. Relegating the detailed analysis to Appendix B, the following proposition shows how type L's equilibrium behaviors depend crucially on his report and signal history.

Proposition 6 *(6.1) Type H reports $m_t = i_t$ for all t. Type L reports the first signal truthfully. In the second report, L reports $m_1 = i_1$ if $i_1 = m_0$. When $i_1 \neq m_0$, there exists a cutoff value p_0' that type L repeats the first report with probability $\pi_1 > 0$ if $p_0 \geq p_0'$ and reports $m_1 = i_1$ otherwise. In the third report, type L agent reports $m_2 = i_2$ if $i_2 = m_1$.*

(6.2) When p_0, p_1 are sufficiently close to p_2, type L lies less in his second report than that in the partial revelation equilibrium of the two signal case ($\pi_1 < \pi^*$). In the third report, he repeats earlier report with probability one if $i_1 = m_0 = m_1$ and $i_2 \neq m_1$. He repeats m_1 with probability $\pi_2^3 > 0$ if $i_2 = i_1 \neq m_1, m_1 = m_0$.

[21] See Appendix B for further discussions on these assumptions.

(6.3) When p_0, p_1 are much smaller than p_2, type L lies more in his second report than that in the partial revelation equilibrium of the two signal case ($\pi_1 > \pi^*$). In the third report, he repeats earlier report with probability $\pi_2^2 > 0$ if $i_1 = m_0 = m_1$ and $i_2 \neq m_1$. He reports $m_2 = i_2$ truthfully if $i_2 = i_1 \neq m_1$ and $m_1 = m_0$.

At the center of Proposition 6 is the tradeoff between the truthfulness of the average agent's earlier report (m_1) and his later one (m_2). In the present model, when $i_1 \neq m_0$, an average agent decides whether to report $m_1 = i_1$ truthfully by looking forward, calculating how much reputation he gains from being consistent early vis-a-vis the potential reputation cost of changing his mind later. Similar to the two signal model before, type L can imitate H in two ways: consistency and (eventual) accuracy. Consistency is primarily driven by the high initial level of H's information while accuracy is driven by his big improvements. However, in a very short process such as the two signal model, an average agent has only one chance to balance the two attributes. With a longer report sequence, an average agent may be able to appear consistent and accurate at different times.

First, when type H's initial signals are highly accurate (p_0, p_1 sufficiently close to p_2), type H is unlikely to change his mind in his final report m_2. The second part of Proposition 6 shows that, if he has lied against his second signal, type L is more likely to repeat his second report ($m_2 = m_1$) even if both of his informative signals suggest the opposite ($i_1 = i_2 \neq m_2$). The reason is straightforward: the smart agent is very unlikely to change his mind at this point, therefore the average agent suffers a big loss in reputation from doing so. In other words, type L agent is too committed to the early reports to change his mind later.

This "escalation effect", however, improves the average agent's incentive to tell the true second signal ($\pi_1 < \pi^*$). Intuitively, an average agent wants to appear less consistent than he would have in a two signal model because if he lies in the second report, he may have little choice but to lie further in the next report. As a result, the average agent may risk a very low reputation due to the very likely loss in accuracy. In this case, the indirect effect of requiring the third report can improve type L's truthtelling and the principal's decisionmaking. If the principal gets rid of the third report because of the small signal quality improvement after the second signal, then type L lies more in his second report and the principal may have to rely on an uninformative signal with higher probability.

Second, when p_0, p_1 are not too close to p_2, even a smart agent may change his mind and report $m_2 \neq i_1$ with nonnegligible probability. The third part of Proposition 6 shows that an average agent reports his true final signal if he has lied in his second report to appear consistent. The reason is that, in this case, type L can afford to change his mind later without large loss in term of reputation. While improving accuracy of the final report, this "better late than never" effect worsens the truthtelling incentive in the second report. Intuitively, an average agent wants to appear more consistent than he would have in the two signal model because he can change his mind later and appear accurate. Therefore, counterintuitively, the principal may not want to require the last report when the smart agent's final signal is very accurate.

This subsection highlights the subtle yet important interaction between any new report required by a principal and its impact on the agent's *overall* incentives. It is also suggests that, when the agent receives multiple signals of increasing quality, the principal may want to choose both the optimal number of reports required and the optimal timing of each report.

6 Conclusion

This paper investigates the role sequential reports play when an agent improves in the ability to observe the state of the world. Since the agent receives multiple signals of ascending quality, both the sequencing and the accuracy of the reports may become a signal of ability.

Contrary to past work, this paper shows that an average agent may often repeat his initial report despite the fact that the market values mind changes more than consistency. Mind changes reflect confidence in signal quality improvement. Although consistency may be more valued in markets with reputational concerns and little signal quality improvement, mind changes can be the prized sign of the fast learners and the talented, at least when the agent knows how smart he is. In a similar model but with symmetric information, consistency again becomes more valued by the market because the agent infers-just like the market-that he is likely average if he receives inconsistent signals.

Inefficiency in the sequential reports model may be quite high when there are few smart types.

This paper shows that sequential reports can give the principal more, and finer, information about the agent's type before the accuracy of his reports can be checked. Therefore the principal may be able to make better first stage decision based on the sequencing of the reports, which is crucial when her decision problem is very sensitive to the accuracy of the reports. Requiring final report only is optimal when, ex ante, the agent is very unlikely to be smart, or when the average agent's signals exhibit larger improvement.

One question emerging from this paper is the structure of the reputational concerns itself. This paper shows that in a value of information model, even though the agent himself is risk neutral, the implicit incentive structure itself is generally convex and may encourage risk taking behavior, especially on the part of the average agent. How this type of implicit incentive structure evolves over time could be a question of interest. Another question is the optimal reporting protocol in a more general environment where the agent receives new and better signals over time. The incentives identified in this paper suggest that requiring very early report may cause the average agent to commit to an early opinion without knowing much just to appear smart. On the other hand, requiring late reports only is likely to make it difficult for the decisionmaker to learn the fineness of the agent's information. Therefore, given the agent's signals, the principal may want to choose both the optimal number of reports *and* the optimal timing of these reports to encourage truthful revelation. Optimality of such reporting protocols is a question of further research.

APPENDIX A: PROOFS

Proof of Lemma 1:

1. Recall that $w(\hat{\eta}) = \sum_{m^1} \Pi^1(m^1, \hat{\eta})|_{a=a^*} \equiv V(a^*(\hat{\eta}))$. Consider two posterior distributions of the agent's type θ_1 and θ_2 such that $Pr(\theta_i = H) = \eta_i, i = 1, 2$ and $\eta_2 > \eta_1$. Let $\theta = \gamma\theta_1 + (1-\gamma)\theta_2$ denote a convex combination of θ_1, θ_2 and let $V(a_1^*(\theta_1)), V(a_2^*(\theta_2)), V(a^*(\theta))$ denote the respective wages of the agent in the second stage given these posterior distributions. Then,

$$\begin{aligned}
\gamma V(a_1^*(\theta_1)) + (1-\gamma)V(a_2^*(\theta_2)) &\geq \gamma V(a^*(\theta_1)) + (1-\gamma)V(a^*(\theta_2)) \\
&= V(a^*(\gamma\theta_1)) + V(a^*((1-\gamma)\theta_2)) \\
&= V(a^*(\gamma\theta_1 + (1-\gamma)\theta_2)) \\
&= V(a^*(\theta)).
\end{aligned}$$

Thus the wage function $w(\hat{\eta})$ is convex.

Second, $w(\hat{\eta})$ is piecewise linear. Simple calculations show that the principal's profit in the second stage after each possible report sequence is a linear function of the posterior estimate of the agent's talent $\hat{\eta}$:

$$\Pi^1(g,g) = \frac{r}{2}g + \frac{1-r}{2}b + \hat{\eta}\left[g(p_0p_1 - \frac{r}{2}) + b((1-p_0)(1-p_1) - \frac{1-r}{2})\right]a^*;$$
$$\Pi^1(b,g) = \frac{r}{2}g + \frac{1-r}{2}b + \hat{\eta}\left[g((1-p_0)p_1 - \frac{r}{2}) + b(p_0(1-p_1) - \frac{1-r}{2})\right]a^*;$$
$$\Pi^1(g,b) = \frac{1-r}{2}g + \frac{r}{2}b + \hat{\eta}\left[g(p_0(1-p_1) - \frac{1-r}{2}) + b((1-p_0)p_1 - \frac{r}{2})\right]a^*;$$
$$\Pi^1(b,b) = \frac{1-r}{2}g + \frac{r}{2}b + \hat{\eta}\left[g((1-p_0)(1-p_1) - \frac{1-r}{2}) + b(p_0p_1 - \frac{r}{2})\right]a^*.$$

The constant part of Π^1 is the value of information provided by an average agent, and the slope part is the added value of a smart agent. The principal chooses $a^* = 1$ if $\Pi^1 \geq 0$ and $a^* = 0$ otherwise. Summing up the profit functions where $a = 1$, it is easy to see that the expected profit takes the form of a piecewise linear function where the slope varies for different ranges of $\hat{\eta}$. Moreover, the wage increases in $\hat{\eta}$ when type H learns faster than L as defined by inequality (1).

2. When the principal's optimal action $a^*(m^1; \hat{\eta})$ is independent of the agent's posterior ability $\hat{\eta}$, then her decision rule depends on the report sequence only. Observe the profit functions above, for any given report sequence and for any $\hat{\eta} \in [0,1]$, Π^1 is either strictly larger or smaller than zero. Therefore the principal chooses $a = 1$ for all the report sequences such that the respective profit is positive and $a = 0$ otherwise. Summing up the report sequences such at $a = 1$, the resulting wage function is strictly linear in $\hat{\eta}$. ‖

Proof of Proposition 1:

First, suppose that in equilibrium both types of agent report $m_0 = i_0$ truthfully. Then for the agent to report $m_1 = i_1$, the following four truthtelling incentive constraints need to be satisfied:

$$(IC_1^L) \qquad (w(CR) - w(W))Pr(b|b,g,L) < (w(R) - w(CW))Pr(g|b,g,L);$$
$$(IC_2^L) \qquad (w(CR) - w(W))Pr(g|g,g,L) > (w(R) - w(CW))Pr(g|g,b,L);$$
$$(IC_1^H) \qquad (w(CR) - w(W))Pr(b|b,g,H) < (w(R) - w(CW))Pr(g|b,g,H);$$
$$(IC_2^H) \qquad (w(CR) - w(W))Pr(g|g,g,H) > (w(R) - w(CW))Pr(b|g,g,H).$$

(1) Recall that w is a piecewise-linear, increasing function (the right derivative is used when w is not differentiable). Thus $w' \geq 0$, and for $\eta \leq \hat{\eta}$,

$$sign\left(\frac{\partial(w(CR) - w(W))}{\partial p_0}\right) = sign\left(-w'(p_1(1-p_1)\eta^2 - \frac{r(1-r)(1-\eta)^2}{4p_0^2}\right) \geq 0;$$
$$sign\left(\frac{\partial(w(R) - w(CW))}{\partial p_0}\right) = sign\left(w'(p_1(1-p_1)\eta^2 - \frac{r(1-r)(1-\eta)^2}{4(1-p_0)^2}\right) \leq 0.$$

At $p_0 = \frac{1}{2}$, it is easy to see that $w(CR) - w(W)) = w(R) - w(CW)$ and all the above ICs hold. At $p_0 = 1$, the agent always needs to be consistent and IC_1^L is clearly violated. Therefore there is a cutoff value p_0^L such that for $p_0 \leq p_0^L$, IC_1^L holds. When type H is assumed to improve faster, IC_1^H also holds. Therefore for $p_0 \leq p_0^L$, all four ICs hold strictly and there is a full revelation equilibrium.

(2) When type H improves faster (inequality (1) holds), and $p_0 > p_0^L$, then there is no full revelation equilibrium. Consider the mixing strategy that when $i_1 \neq m_0$, type L repeats m_0 with probability π and reports the true i_1 with probability $1 - \pi$. Observe that the number of constraints can be reduced by examining the agent's belief about the true state of the world after observing his second signals:

$$Pr(g|g,g,H) = \frac{p_0 p_1}{2p_0 p_1 + 1 - p_0 - p_1} \geq Pr(g|g,g,L) = r; \quad Pr(g|b,g,H) = \frac{p_1(1-p_0)}{p_0 + p_1 - 2p_0 p_1} \leq r.$$

Thus if IC_1^L binds, IC_2^L and IC_2^H hold automatically. Thus if L is willing to mix, both types of agents report the true second signal if it agrees with their initial report. Furthermore, since inequality (1) is assumed to hold, IC_1^H also holds. Moreover, note that $w(CR) - w(W)$ decreases with π while $w(R) - w(CW)$ increases with π. Thus the left hand side of IC_1^L decreases with π while the right hand side increases with π. At $\pi = 1$, i.e., when L always pools, $LHS < RHS = w(1)$. Thus when $LHS \geq RHS$ at $\pi = 0$, there exists a $\pi^* \in (0,1)$ such that $LHS = RHS$. The mixing probability π^* is implicitly defined by:

$$\begin{aligned} & \frac{p_0(1-r)(p_1 - r - \pi + r\pi)}{[p_0 p_1 \eta + \frac{1}{2}(r + \pi - r\pi)(1-\eta)][p_0(1-p_1)\eta + \frac{1}{2}(1-\pi)(1-r)(1-\eta)]} \\ = & \frac{(1-p_0)r(p_1 - r + r\pi)}{[(1-p_0)p_1\eta + \frac{1}{2}r(1-\pi)(1-\eta)][(1-p_0)(1-p_1)\eta + \frac{1}{2}(1 - r + r\pi)(1-\eta)]}. \end{aligned}$$

In order to see that π^* increases in p_0, consider $p_0 = p_0^1 > p_0^L$. Let type L's mixing probability at p_0^1 be π^1, then type L prefers to be consistent at $p_0 > p_0^1$ and $\pi = \pi^1$ because the left hand side of IC_1^L increases with p_0. In order to be indifferent, he needs to increase the mixing probability.

(3) When type L improves faster (inequality (1) does not hold), then similar to the proof for part (1), for p_0 sufficiently close to $\frac{1}{2}$, there still exists a full revelation equilibrium. However, in this case IC_H^1 becomes the binding constraint as p_0 increases. Thus for η small, there exists a cutoff p_0^H such that for all $p_0 \geq p_0^H$, IC_H^1 is violated even though IC_L^1 still holds strictly. Suppose that type H repeats his first report with probability y when $i_1 \neq m_0$. Simple calculation shows that the left hand side increases with y, thus type H will repeat his first report with probability 1. As a result, type L has to repeat his initial report because mind changes leads to zero posterior probability of being smart.

Finally, given the continuation equilibrium, we need to check whether the agent wants to report $m_0 = i_0$. Since type L's first signal is completely uninformative, he is indifferent between reporting his signal or its opposite. Assume that the agent reports truthfully when indifferent. Type H prefers to report $m_0 = i_0$ if the incentive constraint $w(CR)p_1 - w(R)p_1 + w(W)(1-p_1) - w(CW)(1-p_1) \geq 0$ is true. Simple algebra shows that the above holds if $CR \geq R$ and $W \geq CW$. Recall that $W \geq CW$ at $\pi = 0$, and W increases in π while CW decreases in π. Thus in a partial information revelation equilibrium, $W \geq CW$. Also, in equilibrium, $CR \geq R$, otherwise type L should deviate by reducing his mixing probability and receive higher payoff. Therefore the IC always holds and both types report $m_0 = i_0$ in equilibrium. ‖

Proof of Proposition 2:

First, compare the principal's beliefs about the agent's type before the state is observed in a partial information revelation equilibrium. Simple calculation shows that $Pr(H|m_0 = m_1) \geq Pr(H|m_0 \neq m_1)$ if $\pi \leq (2p_0 - 1)(2p_1 - 1)$ and $Pr(H|m_0 = m_1) \leq Pr(H|m_0 \neq m_1)$ if $\pi \geq (2p_0 - 1)(2p_1 - 1)$.

Second, note that if the agent reports truthfully, for any given p_0, p_1,

$$\frac{\partial(CR-R)}{\partial r} = \frac{\partial}{\partial r}\left(\frac{p_0 p_1 \eta}{p_0 p_1 \eta + \frac{1}{2}r(1-\eta)} - \frac{(1-p_0)p_1\eta}{(1-p_0)p_1\eta + \frac{1}{2}r(1-\eta)}\right) \leq K(p_1^2\eta^2 - r^2(1-\eta)^2) \leq 0$$

$$\frac{\partial(W-CW)}{\partial r} = \frac{\partial}{\partial r}\left(\frac{p_0(1-p_1)\eta}{p_0(1-p_1)\eta + \frac{1}{2}(1-r)(1-\eta)} - \frac{(1-p_0)p_1\eta}{(1-p_0)p_1\eta + \frac{1}{2}(1-r)(1-\eta)}\right) \leq K(p_1^2\eta^2 - r^2(1-\eta)^2) \leq 0$$

Because the expected value of giving consistent reports decreases with r and that of giving inconsistent reports increases with r, the mixing probability π decreases with r. Moreover, in a partial information revelation equilibrium, $CR \geq R, W \geq CW$. Otherwise the low type should reduce the mixing probability to appear more inconsistent. The highest mixing probability is obtained for any given p_0, p_1 at $r = \frac{1}{2}$. Substitute into formulas for CR and R, we have all $\pi^* \leq (2p_0 - 1)$. Denote $\overline{\pi} = (2p_0 - 1)$.

Third, the market's estimate of ability before state is realized can also be written as a linear function of the posteriors and the market's belief about the true state, i.e., $Pr(H|m_0, m_1; m_0 = m_1) = \beta(\pi) \times CR + (1-\beta(\pi)) \times CW$; $Pr(H|m_0, m_1; m_0 \neq m_1) = \gamma(\pi) \times R + (1-\gamma(\pi)) \times W$, where

$$\beta(\pi) \equiv \frac{p_1\eta + \frac{1+\pi}{4p_0}(1-\eta)}{p_1\eta + \frac{(1-p_1)(1-p_0)\eta}{p_0} + \frac{1+\pi}{2p_0}(1-\eta)}; \quad \gamma(\pi) \equiv \frac{p_1\eta + \frac{1-\pi}{4(1-p_0)}(1-\eta)}{p_1\eta + \frac{(1-p_1)p_0\eta}{1-p_0} + \frac{1-\pi}{2(1-p_0)}(1-\eta)}.$$

Moreover, for mixing probability $\pi^* \leq \overline{\pi}$, it is easy to see that $\beta(\overline{\pi}) \leq \beta(\pi^*)$ and $\gamma(\overline{\pi} \geq \gamma(\pi^*))$. Since $\beta(\overline{\pi}) \geq \gamma(\overline{\pi})$, it is easy to see that $\beta(\pi^*) \geq \gamma(\pi^*))$. Thus the market values consistency more for $\pi^* < \overline{\pi}$. But if the mixing probability $\pi^* \geq (2p_0-1)(2p_1-1)$, then the market values inconsistency more. Contradiction. Therefore when w is linear, all mixing probability $\pi^* < (2p_0-1)(2p_1-1)$ and consistency is more valued.

If $\pi^* \geq (2p_0-1)(2p_1-1)$, $Pr(H|m_0 = m_1) \leq Pr(H|m_0 \neq m_1)$. then the market values consistency more in equilibrium. The above shows that it would not occur when w is linear because the mixing probability does not exceed $(2p_0-1)(2p_1-1)$.

Now, suppose that the principal's decision problem is very type-dependent and w is a convex piecewise-linear function. This function, however, may take may different forms, depending on the principal's optimal decision rule. Consider a very convex case in which the principal only invests when type H reports g, g, then w is as demonstrated in Example 1. Then IC_1^L becomes: $w(CR)(1-r) < w(R)r$, which simplifies into $CR(1-r) < Rr$. Note that the gap between the left hand side and right hand side becomes larger and thus π^* increases. When p_0 is very high, $\pi^* > (2p_0-1)(2p_1-1)$. ‖

Proof of Proposition 3:

First, find the continuation equilibrium assuming that the agent reports truthfully in the first report. Suppose that $m_0 = g$, then there exist only two truth-telling IC conditions in the second report, depending on whether $i_0 = i_1$, rather than four in the asymmetric information model in proposition 1:

$$[Pr(H|g,g;g) - Pr(H|g,b;g)]Pr(g|g,g) \geq [Pr(H|g,b;b) - Pr(H|g,g;b)]Pr(b|g,g); \quad (IC_1)$$

$$[Pr(H|g,g;g) - Pr(H|g,b;g)]Pr(g|g,b) \leq [Pr(H|g,b;b) - Pr(H|g,g;b)]Pr(b|g,b). \quad (IC_2)$$

(1) Similar to proposition 1, if the agent reports truthfully, at $p_0 \approx \frac{1}{2}$, both incentive constraints hold strictly and there exists a full revelation equilibrium. Furthermore, since $\eta \leq \overline{\eta}$, the left hand side increases with p_0 and the right hand side decreases with it. Thus there exists a cutoff value of p_0, $\hat{p}_0$ such that the agent is indifferent between repeating his first report or reporting his second signal truthfully.

Moreover, the agent's own estimate of the state given his signals become:

$$Pr(g|g,g,L) < Pr(g|g,g) = \frac{p_0 p_1 \eta + \frac{r}{2}(1-\eta)}{[p_0 p_1 + (1-p_0)(1-p_1)]\eta + \frac{1}{2}(1-\eta)} < Pr(g|g,g,H)$$
$$Pr(g|g,b,H) < Pr(g|g,b) = \frac{p_0(1-p_1)\eta + \frac{1-r}{2}(1-\eta)}{[p_0(1-p_1) + (1-p_0)p_1]\eta + \frac{1}{2}(1-\eta)} < Pr(g|g,b,L)$$

Using the above inequalities, it is easy to see that $p_0 > p_0^L$ if type H improves faster, and $p_0 > p_0^H$ if type L improves faster.

(2) When $p_0 > \hat{p}_0$, IC_2 does not hold and the agent is tempted to appear more consistent by repeating his first report. Suppose that the agent repeats his initial report with probability τ, then the expected wage of giving consistent reports is $Ew(m_0 = m_1) = Pr(H|g,g;g)Pr(g|g,b) + Pr(H|g,g;b) + Pr(b|g,g)$. At $\tau = 0$, the expected value of consistent reports is higher than giving inconsistent reports. Simple calculation can show that $\frac{\partial Ew}{\partial \tau} > 0$, thus the expected wage of consistent reports is increasing, and the agent will repeat his initial report. Since there exists an uninformative equilibrium in which the agent always repeat his first report. ‖

Proof of Proposition 4:

(1) If only report m^f is required, then the principal's posterior estimates of the agent's ability become: $Pr(H|m^f = s) = \frac{p_1\eta}{p_1\eta + r(1-\eta)}$ and $Pr(H|m^f \neq s) = \frac{(1-p_1)\eta}{(1-p_1)\eta + (1-r)(1-\eta)}$. Because $w(\hat{\eta})$ is an increasing function, $Pr(H|m^f = s) > Pr(H|m^f \neq s)$ implies that $w(Pr(H|m^f = s)) > w(Pr(H|m^f \neq s))$ as well.

(2) When the principal requires $\vec{m}^f = (m_0, m_1)$ after the agent receives both signals, the agent will report both signals truthfully if he receives a higher expected wage in the second stage than any other three report sequences. Thus, for each signal history (i_0, i_1), truthtelling requires that, for all $m_0' \neq i_0, m_1' \neq i_1$:

$$\sum_s w(Pr(H|m_0 = i_0, m_1 = i_1; s))Pr(s|i_0, i_1, \eta) \geq \sum_s (Pr(H|m_0', m_1', s)Pr(s|i_0, i_1, \eta).$$

Step 1: suppose that the signal sequence is $(i_0 = b, i_1 = g)$, first consider the case when inequality (1) holds. For truthful revelation, the following three incentive constraints must hold:

$$\begin{aligned} & w(Pr(H|b,g;g))Pr(g|b,g,\eta) + w(Pr(H|b,g;b))Pr(b|b,g,\eta) \\ \geq\ & w(Pr(H|g,g;g))Pr(g|b,g,\eta) + w(Pr(H|g,g;b))Pr(b|b,g,\eta) \\ & w(Pr(H|b,g;g))Pr(g|b,g,\eta) + w(Pr(H|b,g;b))Pr(b|b,g,\eta) \\ \geq\ & w(Pr(H|g,b;g))Pr(g|b,g,\eta) + w(Pr(H|g,b;b))Pr(b|b,g,\eta) \\ & w(Pr(H|b,g;g))Pr(g|b,g,\eta) + w(Pr(H|b,g;b))Pr(b|b,g,\eta) \\ \geq\ & w(Pr(H|b,b;g))Pr(g|b,g,\eta) + w(Pr(H|b,b;b))Pr(b|b,g,\eta) \end{aligned} \tag{2}$$

First, for $p_0 > \frac{1}{2}$, there does not exist a full revelation equilibrium. Suppose so, recall that $CR - W \geq R - CW$, and the gap increases with p_0. Observe that $Pr(g|b,g,\eta) \geq \frac{1}{2}$, therefore both types receive better payoff reporting $m_0 = m_1 = g$.

Second, suppose that there exists a mixing equilibrium in which type L reports $m_0 = m_1$ with some probability $1 - \omega$ and $m_0 = b, m_1 = g$ with probability ω, then it must be that incentive constraint (2) binds for L and strictly holds for H. Given inequality 1, however, it is easy to observe that whenever type L is indifferent, type H strictly prefers to report g, g, contradiction.

Next, suppose that there exists a mixing equilibrium in which H mixes and type L always reports truthfully. This increases the reputational payoff of consistent messages, therefore type L can deviate and improves his payoff without reducing the probability of being correct ($Pr(g|i_0 = g) = \frac{1}{2}$), hence this cannot be an equilibrium. The only possible equilibrium in this case is for both type H and L to report $m_0 = m_1 = g$.

Step 2: suppose that the signals are consistent, e.g.,($i_0 = i_1 = g$):

$$\begin{aligned}
&W(Pr(H|g, g; g))Pr(g|g, g, \eta) + W(Pr(H|g, g; b))Pr(b|g, g, \eta)\\
\geq\ &W(Pr(H|b, b; g))Pr(g|g, g, \eta) + W(Pr(H|b, b; b))Pr(b|g, g, \eta)\\
&W(Pr(H|g, g; g))Pr(g|g, g, \eta) + W(Pr(H|g, g; b))Pr(b|g, g, \eta)\\
\geq\ &W(Pr(H|g, b; g))Pr(g|g, g, \eta) + W(Pr(H|g, b; b))Pr(b|g, g, \eta)\\
&W(Pr(H|g, g; g))Pr(g|g, g, \eta) + W(Pr(H|g, g; b))Pr(b|g, g, \eta)\\
\geq\ &W(Pr(H|b, g; g))Pr(g|g, g, \eta) + W(Pr(H|b, g; b))Pr(b|g, g, \eta)
\end{aligned}$$

It is easy to see that when the signals are consistent, from part 1, $Pr(H|g, g; g) > Pr(H|b, b; g)$. Therefore, the agent prefers to report truthfully when the signals agree, i.e., to report $m_0 = i_0 = g, m_1 = i_1 = g$.

Step 3: when inequality (1) does not hold. Then when the signals differ, it is possible to have a mixing equilibrium in which type L mixes and type H reports $m_0 = b, m_1 = g$. $\|$

Proof of Proposition 5

Let a^* be the optimal action given the message(s). Let $\Pi^0_f(g), \Pi^0_s(g), \Pi^0(g, g), \Pi^0(b, g)$ be, respectively, the expected first stage profit under the final reporting system given message g; that under the full pooling equilibrium of the sequential reporting system given message g, and the expected profits after report sequence g, g and b, g. Then,

$$\begin{aligned}
\Pi^0_f(g) &= \frac{r}{2}g + \frac{1-r}{2}b + \eta\left[g(\frac{p_1}{2} - \frac{r}{2}) + b(\frac{1-p_1}{2} - \frac{1-r}{2})\right]a^*;\\
\Pi^0_s(g) &= \frac{1}{4}g + \frac{1}{4}b + \eta\left[g(\frac{p_0}{2} - \frac{1}{4}) + b(\frac{1-p_0}{2} - \frac{1}{4})\right]a^*;\\
\Pi^0(g, g) &= \frac{r + (1-r)\pi}{2}g + \frac{1-r+r\pi}{2}b + \eta\left[g(p_0p_1 - \frac{r+(1-r)\pi}{2}) + b((1-p_0)(1-p_1) - \frac{1-r+r\pi}{2})\right]a^*;\\
\Pi^1(b, g) &= \frac{(1-\pi)r}{2}g + \frac{(1-\pi)(1-r)}{2}b + \eta\left[g((1-p_0)p_1 - \frac{(1-\pi)r}{2}) + b(p_0(1-p_1) - \frac{(1-r)(1-\pi)}{2})\right]a^*.
\end{aligned}$$

(1) When type L improves faster, proposition 1 shows that when $p_0 \geq p_0^L$, both type H and L report $m_0 = m_1 = i_0$ in equilibrium. When one final report is required, both types report $m_0 = m_1 = i_1$. Recall that $g + b < 0$, thus $a^* = 0$ if $m = b$. Now suppose that $m = g$, simple calculation shows that $\Pi^0_f(g) - \pi^0_s(g) = \frac{1}{2}[(p_1 - p_0)\eta + (r - \frac{1}{2})(1-\eta)](g - b) \geq 0$. Therefore the principal makes (weakly) better decision requiring one final report m^f.

(2) When $\eta \approx 0$, information quality provided by an average agent is crucial. When π is sufficiently large under a sequential reporting system, it is simple to show that $\Pi^0(g) > \Pi^0(g, g), \Pi^0(g) > \Pi^0(b, g)$. Thus the principal should use a final reporting system.

(3) Suppose p_0 is sufficiently close to $\frac{1}{2}$. More precisely, $p_0 \leq p_0^L$ when inequality (1) holds, and $p_0 \leq p_0^H$ when it does not hold, then proposition 1 shows that there exists a full revelation equilibrium. In a full revelation equilibrium, $\Pi^0(g,g) > \Pi^0(g) > \Pi^0(b,g)$, thus the principal can make better decision with sequential reports.

(4) When type H improves faster (inequality (1) holds), $p_0 > p_0^L$, there exists a partial revelation equilibrium with sequential reporting. The expected first profit $\Pi_f^0(g)$ can be broken down into two parts: the first part, $\frac{r}{2}g + \frac{1-r}{2}b$ is what the principal can get using a type L agent's information alone, and the second part $\eta\left[g(\frac{p_1}{2} - \frac{r}{2}) + b(\frac{1-p_1}{2} - \frac{1-r}{2})\right]$ is the additional information a type H brings. When η is not too close to 0, suppose that the principal needs high enough precision such that $Pr(s = g|m) > p_1\eta + \frac{r}{2}(1-\eta)$, then there are two cases. First, suppose that p_0 is not too high, the mixing probability π is high (which is likely to occur when the principal's decision in the second stage is more convex), then $Pr(s = g|gg) > Pr(s = g|g)$, and thus $\Pi^0(g,g) > \Pi^0(g)$. Second, when p_0 is high and π very high, then $\Pi^0(b,g) > \Pi^0(g)$. Therefore the sequential reporting system may offer higher precision than the final reporting system and changes the optimal decision of the principal. ‖

APPENDIX B: THE THREE SIGNAL MODEL

In order to study the changes in the low type's incentives as three reports are required, this section focuses on the case of a partial information revelation equilibrium in which smart agent always reports the truth if two signals are required, as shown in Proposition 1. Consider the case that the agent receives a third signal, let p_2 be the probability that H's third signal i_2 is correct and m_2 as the last report the agent submits. Let $r_t = Pr(i_t = s|L,s)$, $t = 0,1,2$. The parameters are restricted in the following way to simplify the analysis: 1). $p_0 \geq p_0^L$, $r_0 = 1/2$, $\frac{p_0(1-p_1)}{p_1(1-p_0)} < \frac{1-r_1}{r_1}$, and 2). $p_2 = 1$, $r_0 < r_1 < r_2$, $r_2 - r_1$ not too large.

Assumption 1) guarantees the existence of a partial revelation equilibrium when the principal only requires two reports. Assumption 2) describes the situation in which the low type improves at a slower pace than the high type, who improves faster until he learns the true state from his last signal. The simplifying requirement that $p_2 = 1$ makes it possible to focus on report sequences with an accurate final report-a wrong final report is a perfect signal of low ability because $Pr(H|m_0,m_1,m_2,s) = 0$ if $m_2 \neq s$.

Path-dependent Truthtelling in the Final Report

First, consider the change in the principal's posterior estimate of the agent's type, $\hat{\eta}$, when the agent of both type *always* report his true signals. Suppose that the true state $s = g$, and all the reports are accurate, then $\hat{\eta}$ after two reports and three reports are respectively:

$Pr(H|g,g;g) = \frac{p_0p_1\eta}{p_0p_1\eta + \frac{1}{2}r(1-\eta)}$, and

$Pr(H|g,g,g;g) = \frac{p_0p_1p_2\eta}{p_0p_1p_2\eta + \frac{1}{2}r_1r_2(1-\eta)}$.

Simple algebra can show that $Pr(H|g,g,g;g) > Pr(H|g,g;g)$. In a similar fashion, as long as $p_2 \geq r_2$ and the final report is correct, the posterior estimates are higher in the three signal case than that in the two signal case, given identical first reports. On the other hand, for the identical first two reports, the posterior after a wrong report decrease in the three signal case. For example, $Pr(H|b,b,b,g) < Pr(H|b,b,g)$. The reason is that a correct final report thus gives type H another chance to separate himself from type L. Consistent and correct message sequences thus gives higher reputational payoff in the three signal case than the two signal cases. With career concerns, however, the agent can appear more consistent in either the second or the final signal. From above, early consistency is more attractive if the agent tells the truth otherwise.

Assume that in the first two reports, both types report truthfully if their first two signals agree, but type L repeats his first report with some probability π_1 if his second signal differs from the first, which is shown later to be part of the equilibrium. Three key ICs constraints are as follows:

$$\begin{array}{ll} IC_1^i & Pr(H|g,b,g;g)Pr(g|g,b,g;\eta) \geq Pr(H|g,b,b;b)Pr(b|g,b,g,\eta) \\ IC_2^i & Pr(H|g,g,g;g)Pr(g|g,g,b;\eta) \leq Pr(H|g,g,b;b)Pr(b|g,g,b,\eta) \\ IC_3^i & Pr(H|g,g,g;g)Pr(g|g,b,b;\eta) \leq Pr(H|g,g,b;b)Pr(b|g,b,b,\eta) \end{array}$$

Of the above three ICs constraints, the first two describe L's incentive to repeat himself if he has reported truthfully before and his final signal $i_2 \neq m_1$. The last one studies his incentive to lie *if* he has lied against his true second signal to appear consistent, i.e., if his signals are (i_0^g, i_1^b, i_2^b) but he has reported (g,g) before. Since all these ICs are linear in the agent's posterior belief of the true state, simple algebra can show that when $i_2 \neq m_1$, there are two possibilities: 1). IC_2 binds or holds, IC_3 holds strictly; *or* 2). IC_3 binds then IC_2 does not hold.

All the incentive constraints consist of two components: reputational concerns and statistical concerns. Consider for instance the left hand side of IC_2^i: $\underbrace{\overbrace{Pr(H|g,g,g;g)}^{reputational\ concerns}\ \overbrace{Pr(g|g,g,b;\eta)}^{statistical\ concerns}}_{expected\ reputational\ payoff}$. Notice that H surely reports the true i_2 because $p_2 = 1$ by assumption: lying against his true signal i_2 yields zero probability of being a smart agent. More generally, one necessary condition akin to inequality (1) in the two signal model is needed to ensure that type H reports truthfully in the final report. This condition requires that type H has higher relative learning in the sense that his last signal is more likely to be correct:

$$\frac{\frac{p_2}{1-p_2}}{\frac{r_2}{1-r_2}} \geq \frac{\frac{1-r}{r}}{\frac{(1-p_0)(1-p_1)}{p_0 p_1}} \tag{3}$$

Because high type is far more likely to be correct and consistent in i_1 and i_2, type L's incentive depends on his history of private signals, specifically whether they agree with each other or not; *and* his history of reports, specifically whether he has lied before or not.

Note that IC_2 and IC_3 never bind simultaneously. The reason is that the incentive to appear consistent and defy the last signal after private signals (g,g,b) differs from that after private signals (g,b,b), even if the report history is the same. Suppose that type L repeats his first report if $i_1 \neq m_0$ with probability π_1, the following lemma describes the continuation equilibrium:

Lemma 2 (Continuation Equilibrium in the Third Report) *When $\Delta r = r_2 - r_1$ not too large, inequality 1 holds, $w(\hat{\eta}) = \hat{\eta}$ and $p_2 = 1$, there exists a partial revelation equilibrium in the third report in which type H always reports the true i_2. The low type will always report the true signal if $i_2 = m_1$. When $i_2 \neq m_1$, there are two cases:*

(1) IC3 holds at π_1, then if $i_2 \neq m_1, i_1 = m_1, m_1 \neq m_0$, type L repeats m_1 with probability $\pi_2^1 \geq 0$, π_2^1 decreases with π_1. If $i_2 \neq m_1, i_1 = m_1 = m_0$, type L repeats m_1 with probability $\pi_2^2 \geq \pi_2^1, \pi_2^2 = min\{1, \pi_2^2(\pi_1)\}$ and π_2^2 increases with π_1. If $i_2 = i_1 \neq m_1, m_1 = m_0$, type L reports $m_2 = i_2$.

(2) IC3 does not hold at π_1, then if $i_2 \neq m_1, i_1 = m_1, m_1 \neq m_0$, type L repeats m_1 with probability $\pi_2^1 \geq 0$, π_2^1 decreases with π_1. If $i_2 \neq m_1, i_1 = m_1 = m_0$, type L repeats m_1 with probability $\pi_2^2 = 1$. If $i_2 = i_1 \neq m_1, m_1 = m_0$, type L repeats m_1 with probability $\pi_2^3 > 0$, π_2^3 decreases with π_1.

In case (1), when IC_3 holds, or the gap between $Pr(H|g,g,g;g)$ and $Pr(H|g,g,b;b)$ is not too large, the low type mixes more in the continuation equilibrium when his first two true reports agree than they disagree.[22] The mixing probabilities in the third stage, π_2^1 and π_2^2, depend crucially on the level of improvement p_0 and p_1. They both increase in p_1 because the smaller is the gap between p_2 and p_1, the more consistent H becomes in the last two reports. π_2^1 decreases in p_0 while π_2^2 increases in p_0. The intuition is that the higher is p_0, the more important overall consistency becomes.

The reason that $\pi_2^1 < \pi_2^2$ is that having given two truthful and consistent reports, L looks more like a H by repeating himself again. The relative gain in accuracy is small if L changes his mind after two consistent signals because Δr is small; but the relative reputational cost of mind changes may be large because H is likely to have reported correctly in m_0 and m_1. In comparison, if the first two true signals/reports differ and p_0 is high, then type L knows that he is unlikely to appear like type H in term of consistency. Thus he should trust his last, and most informative, signal more to be accurate, which is also a sign of high ability.

Moreover, type L tells the true third signal if he has lied to appear consistent before. The reason is that in this case, if $i_1 \neq i_2, i_2 = m_1$, then the true state is more likely to be $s = m_1 = i_2$, type L would appear more consistent and accurate. If $i_1 = i_2 \neq m_1$ and the gap between $Pr(H|g,g,g;g)$ and $Pr(H|g,g,b;b)$ is not too large, repeating the second message is very likely to be wrong because it is against both of L's informative signals. Here the desire for accuracy outweighs the reputational concerns, and the low type tells the true i_2 if he has lied before. Intuitively, this is the "better late than never" effect.

In case (2), IC_3 does not hold because the gap between $Pr(H|g,g,g;g)$ and $Pr(H|g,g,b;b)$ is very large. Type L would always repeat his previous report when his first two truthful messages agree, but his last signal $i_2 \neq m_1$. The reason is that in this case, type H is very accurate and is unlikely to change his mind later. The low type would rather appear consistent since it is possible that his last signal is wrong. When he has lied before and $i_2 = i_1, i_2 \neq m_1$, type L still wants to repeat his second report with probability π_2^3 against both his true signals and hope for the big prize $Pr(H|g,g,g;g)$. Intuitively, the low type is too committed to his earlier reports and this is the "escalation" effect.

Comparing these two types of continuation equilibria, it is clear that whether type L has reported early consistent report against his true signal i_1 has an impact on his third report. In case (1), he would report the true final signal, because the expected reputational gain of overall consistency cannot outweigh the expected accuracy cost of defying both of his informative signals. In case (2), he would lie against both of his informative signals if his final signal disagrees with his early report. The reason is that even though the expected cost in accuracy is high, the expected cost in reputation due to a late change of mind is even higher.

Equilibrium of the Three Signal Model

Given the above continuation equilibrium, type L needs to decide whether to report $m_1 = i_1$ by calculating how much "reputational stock", $[Pr(H|m_0, m_1)]$, he possesses after the second report. Such calculations include type L's optimal strategy in the third report for any possible third signal. As an example, suppose

[22] Their incentives are history dependent as the result of the nonstationary nature of the model. In a typical model with normality and uncertainty only over means of the parameter of interest, incentives are relatively simpler because they do not depend on previous actions.

that type L receives a second signal different from his initial report g, and the continuation equilibrium is of case (2), in which type L repeats his second report with probability $\pi_2^3 > 0$ if he has lied before, then for him to be willing to report $m_1 = i_1$, the following incentive constraint is needed:

$$[Pr(H|g,g,g;g)\overbrace{[r_2 + (1-r_2)\pi_2^3]}^{future\ strategy} - Pr(H|g,b,g;g)\overbrace{r_2(1-\pi_2^1)}^{future\ strategy}]\overbrace{(1-r_1)}^{dir.\ accuracy\ loss}$$
$$\leq [Pr(H|g,b,b;b)[r_2 + (1-r_2)\pi_2^1] - Pr(H|g,g,b;b)r_2(1-\pi_2^3)]r_1$$

The left hand side describes the net benefit of repeating the first report g, conditional on L is lucky and his second signal is mistaken. Observe that if type L lies and repeats his first report, he suffers a direct accuracy loss because his first signal is useless. Indirectly, lying affects his future strategy, as marked in the left hand side of the IC. Similarly, the right hand side says that if the second signal is indeed correct, what is the net benefit of giving the true message sooner than later. L's tradeoff when his first two signals differ is whether to bet on his first signal or on his second: if he repeats g, he receives potentially big gains $Pr(H|g,g,g;g)$ or big loss $Pr(H|g,g,b;b)$ in his reputations. If he reports i_1^b truthfully, he receives gains $Pr(H|g,b,b;b)$ or suffers losses $Pr(H|g,b,g;g)$.

Proposition 6: See text.

Proof: First, see Lemma 2 in Appendix B for the continuation equilibrium in the third report. Then, given Lemma 2, two key truthtelling ICs conditions for type L in the second report are:

$$\begin{aligned} & (Pr(H|g,g,g;g)(r_2 + (1-r_2)\pi_2^2) - Pr(H|g,b,g;g)r_2)r_1 \\ \geq\ & (Pr(H|g,b,b;b)r_2 - Pr(H|g,g,b;b)r_2(1-\pi_2^2))(1-r_1)) \\ & (Pr(H|g,g,g;g)(r_2 + (1-r_2)\pi_2^3) - Pr(H|g,b,g;g)r_2(1-\pi_2^1))(1-r_1) \\ \leq\ & (Pr(H|g,b,b;b)(r_2 + (1-r_2)\pi_2^1) - Pr(H|g,g,b;b)r_2(1-\pi_2^3))r_1 \end{aligned}$$

When $p_0 \approx 1/2$, then it is easy to see that both types are willing to separate since the first report is not informative and thus does not reflect ability. Taking derivatives, it is shown that there exists a boundary above which type L mixes. The future strategy feeds back into the incentive constraints in the second report. It is important to use the mixing constraints in the third report, or $CRr_1(1-r_2) = CWr_2(1-r_1)$ at π_2^2 and $Rr_1(1-r_2) = Wr_2(1-r_1)$ at π_2^1. There are four cases:

Case 1: $\pi_2^1 = 0, \pi_2^2 \leq 1, \pi_2^3 = 0$: First, note that since IC_2^L binds, at $\pi_2^1 = 0$, $CR + CW > R + W$. From which both IC_1^H and IC_1^L hold because after receiving $m_1 = i_0$, more weight is placed on getting CR, the best posterior. By inequality (1), IC_2^H holds. So we have type L pools with π_1 when he receives $i_1 \neq m_0$ but reports truthfully otherwise.

Case 2: $\pi_2^1 > 0, \pi_2^2 \leq 1, \pi_2^3 = 0$: Here since IC_1^L binds, we have $(CRr_2 - Wr_2(1-\pi_2^1))(1-r_1) = (R(r_2 + (1-r_2)\pi_2^1)) - CWr_2)r_1$. Rearrange and we have $CR + CW > R + W)$, then it is similar to case 1.

Case 3: $\pi_2^1 = 0, \pi_2^2 = 1, \pi_2^3 > 0$: When $\pi_2^1 = 0$, we can see that $(CR(r_2 + (1-r_2)\pi_2^3) - Wr_2)(1-r_1) = (Rr_2 - CWr_2(1-\pi_2^3))r_1$. It is easy to check IC_1^L holds. To check IC_2^H, the difference between the expected reputational payoff of consistency for type L and type H is:

$CR(r_2 + (1-r_2)\pi_2^3) + CWr_2(1-\pi_2^3)) - CR - CW = CR(1-r_2) - CWr_2 > 0$, and that of inconsistency is the same. Therefore IC_2^H holds strictly when IC_2^L binds.

Case 4: $\pi_2^1 > 0, \pi_2^2 = 1, \pi_2^3 > 0$: Since IC_2^L binds, we have $(CR(r_2 + (1-r_2)\pi_2^3) - W(r_2))(1-r_1) = (Rr_2 - CWr_2(1-\pi_2^3))r_1$. Substitute in the IC constraints from the last report, $Rr_1(1-r_2) = Wr_2(1-r_1)$ since $\pi_2^1 > 0$. Then we have $CR(1-r_1) = Rr_1$. It is easy to check that IC_1^L holds strictly:

$CRr_1 = R\frac{r^2}{1-r} > Wr_1r_2 + R(1-r_1)r_2 = R(\frac{r_2^2(1-r_2)}{1-r_1} + (1-r_1)r_2)$.

As for type H, we can show that $(CR - W)(1 - r_1) = (R - CW)r_1$, and since inequality (1) holds, it is easy to see that both type H's ICs hold strictly as in the two signal model. Also use the mixing constraints in the third report, we can have the following requirement on $\pi_2^2(\pi_1)$ and $\pi_2^1(\pi_1)$: if $\frac{\partial}{\partial p_0}(\frac{\pi_2^2}{p_0}) < \frac{r_2}{1-r_2}\frac{1}{p_0^2}$ and $\frac{\partial}{\partial p_0}(\frac{\pi_2^1}{1-p_0}) > -\frac{r_2}{1-r_2}\frac{1}{(1-p_0)^2}$, then there exists a p_0' such that if $p_0 > p_0'$, type L wants to pool with probability π_1. The cutoff value is calculated at the mixing constraint with $\pi_1 = 0, \pi_2^2(0), \pi_2^1(0)$. ‖

Notice that if $p_0 \approx 1/2$, the above equilibrium is obvious because the first signal is uninformative for both types. The principal should not update his belief about managerial ability from the first report. After full separation in the second report, the above equilibrium is similar to the partial revelation equilibrium in the two signal case.

In the three signal game, however, importance of early consistency is not solely determined by p_0. It also depends on how much type L may have to contradict his most informative signal i_2 later. When the premium on $Pr(H|g, g, g; g)$ is not too large, early consistency is more valuable because type L reports i_2 truthfully if he has lied in the second report. Repeating his first report with a high probability thus gives him the possibility of getting the best reputation with relatively little accuracy loss.

The premium on $Pr(H|g, g, g; g)$ can be very large, for example, when p_0 and/or p_1 are much higher than r. Then early consistency, though highly desirable if obtained, tends to come at a potentially high cost. The reason is that not all mind changes are equal with a longer sequence. Changing one's mind early is more likely to be a result of initial lower quality signal, which is common for both types. But changing mind later indicates that the earlier reports are less reliable *and* the agent's ability may have improved very slowly. In this case, reporting the true second signal and then trying to be consistent enables the low type to use his better signal, i_2 more and to give more accurate final report.

The intuition is that for H, if he makes a wrong report, it is much more likely to be early; and if he receives consistent signals about the state of the world, he is unlikely to change his mind later, especially if p_1 is high. Therefore if L repeats his first signal with a high probability, he has to repeat those signals in the third signal with some probability to *continue* to imitate H. Recall that π_2^3 increases with π_1: the more he lies against his second signal, the more he has to lie in the last. This results in type L pools more towards his first uninformative signal and give consistently wrong reports. If he reports true i_1 instead, he may look less consistent to begin with, which reduces his updated probability of being a high type modestly. But he is more correct in his later reports, which is a more reliable sign of being H to the principal. Therefore the low type is willing to tell the true second signal, even when it disagrees with his initial report.

Equilibrium of the three signal game shows that when the H type's first two signals are not too superior, the principal could learn more from the last report, whereas when the first two signals of H type are very good, then she could learn more from the second report, when the mixing is relatively low. The principal knows the equilibrium behaviors of the agents. Thus she may value early inconsistency for two reasons: first if $m_1 \neq m_0$, the agent's second report is true; second, it also means that the agent lies less in his third report.

Change of Equilibrium without the Final Report

This subsection compares the truthtelling incentive in the two signal model with that of the three signal model. In equilibrium, would the principal receive more truthful reports with three reports than two? The answer depends on whether the possibility of a third report opens up or narrows down the expected value of such an option, if appearing consistent in the second report is considered as an option.

First, consider the case when the three signal game has an equilibrium in which type L reports the true final signal if he has lied before,which tends to occur if the quality of type H's second signal is much lower than the quality of his final signal. In this equilibrium, an average manager changes his mind in the final report if $i_1 = i_2 \neq m_1$, and fully utilizes his third and most accurate signal, therefore he is still likely to be accurate eventually despite the initial lying. Because the expected cost in term of accuracy of early consistency is relatively low, he is likely to lie with a higher probability in the second report. Intuitively, the option value of early consistency increases: repeating his first report may lead to the high reward of a consistent and correct signal sequence. If his third signal indicates that his second report is wrong, then he follows his two true signals that is likely to be accurate. Moreover, since the high type has relatively large improvement in the third signal, a late change of mind has a relatively low reputational cost. If the low type did report the true second signal, however, he loses the possible high reward of overall consistency for sure, and risks a lower accuracy: if his third signal disagrees with his second signal, he repeats second signal with probability $(1 - \pi_2^1)$. Therefore the low type prefers to commit to early consistency and then possibly change his mind later after he observes his final signal.

Recall that π^* is the equilibrium mixing probability for type L in a two signal partial information revelation equilibrium, when he receives signal $i_1 \neq m_0$. Similarly, π_1 is the type L's equilibrium mixing probability in the second report when $i_1 \neq m_0$ in the three signal model.

Result 1 *When $\pi_2^2 \leq 1, \pi_2^3 = 0$, and $\frac{(r_2(1-r_1)^2+r_1(1-r_2))}{1-r_1} Pr(H|g,g,g;g) > r_1 Pr(H|b,g,g,g)$ at $\pi^*, \pi_2(\pi^*)$, and $\pi_2^1(\pi^*)$, then the mixing probability in the second message $\pi_1 \geq \pi^*$. One necessary condition for the above result to hold is when $p_0(\frac{r_1 r_2(1-r_1)^2+r^2(1-r_2)}{1-r_1})(1-\pi^*)(r_2+(1-r_2)\pi_2^1) - (1-p_0)(r_1(r_2+(1-r_2)\pi_2^2) + (1-r_1)\pi^* r_2) \geq 0$.*

The above result tends to hold when the original mixing probability π^* is small, since π_2^1 decreases with π^* while π_2^3 increases with it. Note that the indirect effect makes the third signal less desirable for the principal: the final report may reveal more truth, but the second report she receives is less truthful.

Second, consider the case when the three signal game has an equilibrium in which type L repeats his second report with some probability if he has lied before,which tends to occur if p_0, p_1 is very high. Moreover, the mixing probability in the third report increases with that in the second (π_2^3 increases in π_1). In this case, the high type's gain in signal informativeness is small and H is quite likely to find the true state of the world in the second report already. The direct effect suggests that the principal is unlikely to gain more new information. The high type is unlikely to change mind, and the low type is pooling with very high probabilities. In fact, after L reports his first two true consistent signals, he always repeats the report. However, concern for the accuracy of the third report narrows down the option value of consistency and the low type tends to repeat his first report with smaller probability than that in the two signal case.

Result 2 *When $\pi_2^2 = 1, \pi_2^3 \leq 1$, and $(1-r)Pr(H|g,g,g;g) \leq rPr(H|g,g,g;g)$ at $\pi^*, \pi_2^3(\pi^*)$, and $\pi_2^1(\pi^*)$, then the mixing probability in the second message $\pi_1 \leq \pi^*$. One necessary condition for the above result is when: $p_0(1-r_1)(1-\pi^*)(r_2+(1-r_2)\pi_2^1) - (1-p_0)(r_1+(1-r_1)\pi^*(r_2+(1-r_2)\pi_2^3)) \leq 0$.*

This condition is true when the original mixing probability is large. Intuitively, the distance between $(1-r)Pr(H|g,g,g;g)$ and $rPr(H|g,g,g;g)$ at $\pi^*, \pi_2^3(\pi^*)$ measures the reputational gain of early commitment, *holding the future equilibrium level of commitment constant.* The intuition is that type L knows that the more he lies in the second message, he would have to deviate from his third signal more to appear like high type. But the high type is relatively unlikely to change mind later, the more L lies in the second report, the more likely for a late mind change to signal low ability. If the low type repeats his second report, it means that he has to deviate against both of his informative signals and to give a completely wrong report sequence. Thus he lies less than in the two signal case to avoid the overcommitment to consistency.

In this case the requirement of a third signal gives the principal more information and increase her decisionmaking. If the principal gets rid of the third report because the learning curve for both types are relatively flat after the second signal and the low type is pooling with some probability with his second report, depending on p_1, L has less incentive to tell the truth earlier. Without the third signal, the principal's information may deteriorate significantly because she may have to rely on an uninformative signal with high probability. Therefore even though the third signal itself reveals little new information, it provides right incentives for the low type to tell the truth in the second period.

When the principal's decision-making is type dependent, then the reputational payoff in the second stage is convex. With a convex payoff function, the distance between posteriors increases. Therefore, for given parameter values, completely consistent and correct message sequence becomes more valuable than the linear case. Therefore in the third signal, the low type agent's behavior is more likely to fall into the second category of equilibrium. In such a continuation equilibrium, the low type would lie against his true signals ($\pi_2^3 \leq 0$) and fully pool if his first two signal/reports are consistent. As in the linear case, the intuition is that changing one's mind in the third report is too costly in term of reputation.

References

Aghion, P., and J. Tirole (1997): "Formal and Real Authority in Organizations," *Journal of Political Economy*, 105(1), 1–29.

Arya, Anil, J. G., and K. Sivaramakrishnan (1997): "Commitment Issues in Budgeting," *Journal of Accounting Research*, 35(2), 273–278.

Banerjee, A. V. (1992): "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, 107, 797–819.

Blackwell, D. (1953): "Equivalent Comparison of Experiments," *Annals of Mathematics and Statistics*, (24), 265–272.

Blume, Andreas, K. Y.-G., and J. Sobel (1993): "Evolutionary Stability in Games of Communication," *Games and Economic Behavior*, 5(3), 547–575.

Crawford, V. P., and J. Sobel (1982): "Strategic Information Transmission," *Econometrica*, 50(6), 1431–1451.

Dye, A. R., and R. E. Verrecchia (1995): "Discretion vs. Uniformity: Choices among GAAP," *The Accounting Review*, 70(3), 389–415.

Farrell, J. (1993): "Meaning and Credibility in Cheap-Talk Games," *Games and Economic Behavior*, 5(3), 514–531.

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

Myerson, R. (1989): "Credible Negotiation Statements and Coherent Plans," *Journal of Economic Theory*, 48, 264–303.

Ottaviani, M., and P. Sorensen (2001): "Professional Advice: The Theory of Reputational Cheap Talk," Mimeo, University College London.

Penno, M. (1985): "Information Issues in the Financial Reporting Process," *Journal of Accounting Research*, 23(1), 240–255.

Prendergast, C. (1993): "A Theory of "Yes Men"," *American Economic Review*, pp. 757–770.

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

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

Staw, B. M. (1976): "Knee-Deep in the Big Muddy: A Study of Escalating Commitment to a Chosen Course of Action," *Organizational Behavior and Human Performance*, 16, 27–44.

——— (1981): "The Escalation of Commitment to a Course of Action," *Academy of Management Review*, 6, 577–587.

——— (1992): "Deescalation Strategies: a Comparison of Techniques for Reducing Commitment to Losing Courses of Action," *Journal of Applied Psychology*, 77, 419–426.

Wicklund, R. A., and O. Braun (1987): "Incompetence and the Concern with Human Categories," *Journal of Personality and Social Psychology*, 53(2), 373–382.