Mind Changes in Designing Reporting Systems*

Wei Li

Massachusetts Institute of Technology

June, 2003

Abstract

This paper presents models in which a principal makes a decision based on a sequence of reports from an agent with improving ability to observe the state of the world. Agent of different types receives signals of different initial quality as well as varying quality improvements, thus both the accuracy and the sequencing of the reports affect the principal's perception of the agent's ability, which determines the agent's future wage. This paper finds that (1). in equilibrium, mind changes (or inconsistent reports) may signal high ability, yet mediocre agents still repeat their early report with higher frequency because they have less *relative* confidence in their signal quality improvement; (2). sequential reporting can be superior to a final reporting system when first, the sequencing of reports offers the principal more information about the agent's ability; and/or second, it helps the principal gauge how strongly the agent believes in his reports. On the other hand, requiring a final report is optimal when there is no improvement in signal quality, or when the mediocre agents improves relatively faster than the smart ones; (3). request of additional informative report may lead to further distortions of the agent's earlier reports and is thus undesirable to the principal.

Keywords: sequential information transmission, reputational concerns, cheap talk games.

JEL classification: D82 (Asymmetric and Private Information), C70 (Game Theory and Bargaining Theory), M50 (Personnel Economics).

---

*I thank George Akerlof, Abhijit Banerjee, Mathias Dewatripont, Glenn Ellison, Bengt holmström, Botond Kőszegi, Jean Tirole and Muhamet Yildız for insightful comments and encouragement.

1 Introduction

In many economic models, important information, once collected, is transmitted once and for all from a sender to a receiver who may act upon such information (Crawford and Sobel 1982, Aghion and Tirole 1997, Morris 2001). In certain situations, however, a sender's ability to observe the underlying true state of the world may improve with time as he becomes more familiar with the task at hand.[1] As a result, progress reports are frequently employed to transmit information sequentially and in a piecemeal fashion.

Progress reports are frequently observed in reality. Examples include accounting and auditing (Dye and Verrecchia 1995), budgeting (Arya and Sivaramakrishnan 1997), and information disclosure in the financial market (Penno 1985). When an agent is asked to give an initial report and more accurate ones later, however, accuracy of an agent's initial assessment frequently reflects his ability to discover the truth quickly. Thus he 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.[2]

This paper investigates how agents of different abilities react strategically to improvements in the quality of their information. In the model, an agent delivers a sequence of reports about the state of the world, based on his sequence of private signals. To reflect improvement in the agent's ability to observe the state of the world, the agent receives multiple signals of *increasing* accuracy. After each signal, the agent sends a report to the principal, who makes a decision only after the agent's *final* report. Then the state realizes and becomes observable to all. The same reporting 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 their signal quality, but also in its *slope* of improvement. A smart agent learns about the true state of the world with higher initial accuracy than an average one, but his slope of improvement may be higher *or* lower than that of an average agent. The agent only cares about his future reputation,

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

[2] There exists strong experimental and sociological evidence for increasing commitment to a wrong project. See Staw (1976, 1981, 1992), and the references within. Wicklund and Braun (1987), however, show that people who are more confident in their ability seem to be less committed to their early positions than the less confident ones.

i.e., how smart he is perceived to be in the second stage, and cannot be paid contingent on the accuracy of his reports.

First, a two signal model is analyzed to examine whether and to what extent the difference in signal quality improvement matters for the principal's optimal decisions. The first main insight emerging from this model is that mind changes, or inconsistent reports, may signal high ability in equilibrium. Consider the case when a smart agent learns relatively faster than an average one. Ex ante, the smart type is more likely to deliver consistent and accurate reports. Thus an average agent is more likely to defend an early report to appear consistent when he receives conflicting signals. However, unlike some existing models (Scharfstein and Stein 1990), an average agent is more likely to give consistent reports *even though such consistency per se may signal low ability in equilibrium.* The reason for this paradoxical result is that both consistency and accuracy matter in the agent's future reputation. And type L's mind change is more likely to be in the wrong direction than type H's, due to differences in the improvement of their signal qualities. Thus, in expectation, the reputational reward for consistent reports may still exceed that for inconsistent ones for type L. Intuitively, it takes a lot of confidence in one's information *improvement* to change one's mind and to admit an early mistake.

Relatedly, the two signal model studies what the principal can learn about the agent's type from the *sequencing* of the reports alone. That is, whether consistent reports or mind changes signal high ability in equilibrium. In the second stage, the agent is paid the expected value of his reports as a function of his perceived ability. The expected value of his reports, however, depends on the type of optimal second stage decision problem the principal faces. The second main insight of the model is that: in cases when her second stage optimal decision is very sensitive to the accuracy of the reports-which in turn depends on the agent's perceived ability-mind changes may be more valued by the principal because consistent reports very likely result from average agent's lying. In cases when the principal's optimal decision is independent of type, i.e., her decision only depends on the reports she receives regardless of type, then in equilibrium consistency is (weakly) more valued by the market.

The two signal model demonstrates that truthful revelation may not be in the interest of an average agent due to career concerns. The principal, however, is only concerned with the accuracy of the reports and may want to design a reporting system to encourage truthful reporting. One natural question is whether the principal should require sequential reports at all, given that the average agent may repeat his initial, low quality signal just to appear smart. It may be better to require a report after the agent has received all his signals, thereby eliminating the average agent's incentive to appear consistent.

An answer to the above question, and the third main insight of this paper is that both requiring one final report and requiring sequential reports can be optimal, depending on the dependence of the principal's optimal decision on the accuracy of the agent's information, as well as the scale of the agent's signal improvement. Sequential reporting may be optimal when the sequencing of reports (whether the reports are consistent or not) offers the principal a way to generate more information about the agent's type in the first stage and thereby improve her first stage decisionmaking *before* the true state is revealed to all. In contrast, requiring a final report (or vector of reports) will not provide the sequencing information to help the principal better distinguish the agents. Sequential reporting may also be optimal when it helps the principal gauge the strength of the agent's belief in his reports and thus make better decisions. When only a final report is required, the principal cannot determine how likely the state is the one suggested by the report, which may be important for her decision in the first stage if it is sensitive to the report's accuracy. On the other hand, when the average agent improves faster in term of signal quality, and/or when the principal's first stage decision is type-independent, and/or when the fraction of the average agent is sufficiently high, requiring a final report (vector of reports) is more preferred by the principal.

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, counterintuitively, the principal may *not* want to require the last report when the smart agent's signal quality improves a lot in the final signal. 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, this "better late than never" effect 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. On the other hand, the principal may want to request the final report when the improvement in the smart agent's signal quality *levels off*. The reason is that in equilibrium, an average agent may lie against his true third signal 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.

There exists a large literature on herding examining why economic agents may want to appear consistent with others or some existing consensus to increase the market's perception of their talent. In this literature, herding is driven by a statistical externality (Banerjee 1992) and/or reputational externality exerted by the early movers (Scharfstein and Stein 1990). The key difference of the interpersonal herding models and intra-personal models such as the current one is that reputational linkage between the early movers and the late movers is exogenously stipulated in the herding models. If smart agents' observations are independent conditional on the state, then the information derived from the early mover's action merely changes the *priors* about the state. Such updated priors are common knowledge between the market and the late mover. Controlling for the statistical externality of early movers, consistency is valued by the market in herding models only due to the exogenously stipulated reputation linkage between agents. In contrast, in the intra-personal model, reputational concerns are always present and distortions arise even though only the agent's own information is used.

Ottaviani and Sorensen (2001) analyze a static reputational cheap-talk game with very general type distribution where the agent/sender only cares about the market's perception of their ability, following the seminal model of Holmström (1999). Ottaviani and Sorensen (2001) find that when experts do not know their ability type and when the prior of state distribution is sufficiently concentrated, there does not exist an informative equilibrium; when experts do know their ability, there exists a binary informative equilibrium despite a rich message space. The current model focuses on the dynamic aspect of the agent's incentive problem. Due to the improvement in signal qualities, both the sequencing and the accuracy of the reports reflect how talented an agent is. In equilibrium, the sequencing, or the history of the reports, can be quite important in both the principal's decision problem and in the agent's determination of his future report.

Most closely related to this paper, Prendergast and Stole (1996) consider a career concerns' model featuring individual agents with privately known abilities. In their model, agents receive noisy signals about the true profitability of their investments, and the more capable agents have higher precision. They establish a full separating equilibrium in which agents of different abilities choose different amounts of investment in *each* period, but each agent exaggerates their private information initially and become too conservative later. Their model differs from the current one in a key assumption, namely, what reputation the agent is concerned about. In Prendergast and Stole, the agent only cares about his *period-by-period* reputation, which renders it possible to study the effect of each period's action on that period's incentives. This paper aims to focus on the agent's concerns for his overall reputations, and to study explicitly the linkage of the agent's future incentives and the history of his past reports.

The paper proceeds as follows: Section 2 introduces the two signal model and discusses important assumptions. Section 3 conducts equilibrium analysis and shows that, with learning and asymmetric information, mind changes may signal high ability in equilibrium. Section 4 explores theoretically when the principal may make better decisions using progress reports. Section 5 introduces some extensions. Section 6 concludes. All proofs are collected in the Appendix.

2 The Two Signal Model

In the basic model of sequential reporting, a principal needs to make a decision based on an agent's information. Although the model is clearly more general, this paper will couch it in a simple example: the owner of a company needs to make an investment decision after reviewing an initial report (m_0) and a final report (m_1) from an agent 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.[3] Moreover, it is assumed that the principal does not invest without further information, i.e., $g + b \leq 0$.

The agent may be 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 signals about the underlying true state, but his signals' quality depends on his type in a way specified below.

2.1 Environment and Information

The agent works in two stages $N = 0, 1$.[4] In each stage, the true state of the world s is, independently, either good (g) or bad (b). Events *within the first stage* proceed as follows:

- At $t = 0$: the agent gets a fixed wage w_0 and then receives his first signal $i_0 \in \{i_0^g, i_0^b\}$;
- At $t = 0.5$: the agent sends an initial report $m_0 \in \{m_0^g, m_0^b\}$ as to which state his initial signal indicates;
- At $t = 1$: the agent receives his second signal $i_1 \in \{i_1^g, i_1^b\}$;
- At $t = 1.5$: the agent sends a final report $m_1 \in \{m_1^g, m_1^b\}$ 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.

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

[4] 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 truth-telling incentives in the first period. See Lemma 1 for characterization of the wage function in the second stage.

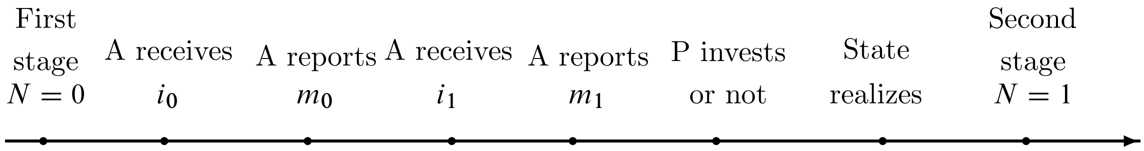

Figure 1: Timeline

The second stage repeats the above process: the agent receives a fixed wage at the beginning. Later, he delivers two sequential reports on the profitability of another project. Timing of this game is illustrated in Figure 1.

Signals an agent receives are independent conditional on the state. Qualities of these signals given an agent's type are:

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

This specification means that, first, agents of both types learn more about the true state from their second signal.[5] Second, the initial signal i_0 itself is not informative about ability, i.e., $Pr(i_0^g|H) = Pr(i_0^g|L) = \frac{1}{2}$. This assumption enables the analysis to focus on the dynamic incentive problems due to improvement in signal quality, because in equilibrium the agent is not tempted to lie in his first report. However, this assumption restricts the state distribution to be symmetric.[6]

Progress reports before decisionmaking are not uncommon in reality. Sometimes, certain actions have to be taken based on the available early information.[7] When the agent's signal quality improves gradually, as in the current model, there are two reasons for the sequential reporting process to occur. First, the principal may require multiple reports to fathom how strongly the agent believes in his reports, rather than just which state is more likely. Second, even if requiring one report at the end is optimal, it is *not* always possible for the principal to forbid voluntary flow of unsolicited information. A smart agent may want to signal his ability by giving an early assessment through

[5] The assumption that type L's first signal is completely uninformative simplifies the analysis. The results apply when L's first signal is of very low quality.

[6] 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.

[7] For example, a bank regulator in charge of both licensing banks and monitoring problematic banks has to decide whether to open a bank based on his initial information, before new and better quality information on the bank's profitability is revealed (Dewatripont and Tirole 1994).

either formal or informal channels if his first signal is quite accurate. Section 4 studies formally when sequential reporting is optimal and when one final report is optimal.

2.2 Payoffs

Both 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 and $\hat{\theta}$ be the agent's type conditional on his reports and the observed state in the first stage: $\hat{\theta} = H$ with probability $\hat{\eta}$ and $\hat{\theta} = L$ with probability $1 - \hat{\eta}$, $\hat{\eta} \equiv Pr(\theta = H|m^0, s)$. Let $I(a^N) = 1$ if $a^N = 1$ and $I(a^N) = 0$ if $a^N = 0$. The principal chooses action $a^N \in \{0, 1\}$ to maximize stage N's profit Π^N:

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

The agent cannot be paid conditional on the accuracy of reports because the true state of the world is assumed to be eventually observable but unverifiable, thus the agent cares only about his wage in the second stage. Furthermore, this paper makes the standard assumption in the cheap talk literature that no contract can be written on messages. Assume perfect competition for the agent's services among the principals, his second stage wage is simply the expected value of his information conditional on the principal's updated belief of him being type H. Let $a^*(\hat{\eta}, m)$ denote the principal's optimal action given $\hat{\eta}, m$, then the wage of the agent is $w(\hat{\eta}) = \sum_{m^1} \Pi^1(\hat{\eta}, m^1)|_{a=a^*} \equiv V(a^*(\hat{\eta}))$.[8] 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 $a^*(\hat{\eta}, m)$ is independent of the agent's posterior ability.*

The agent's estimated type and his reports in the second stage jointly affect the belief of the principal in the true state of the world. The value function of his information is convex as shown by Blackwell (1953), because the principal can make better (and potentially different) decisions given

[8] 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.

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. To see this, 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 not do worse than if she has to choose an action knowing only the distribution of the agent's type. 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 the types.

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 $s = g$ with very high probability.[9] In this case, the payoff function takes the form of a kinked function: $w = 0$ when $\hat{\eta} \le \eta_0$ and w increases linearly with $\hat{\eta}$ when $\hat{\eta} \in [\eta_0, \eta_1)$, $w(\eta_0) \equiv 0$. If the probability that the agent is smart is sufficently high, say $\hat{\eta} \in [\eta_1, 1]$, then $w(\hat{\eta})$ increases linearly with $\hat{\eta}$ but at a steeper slope. Intuitively, an agent likely to be average is not worth hiring again ($w = 0$) because no reports from him can lead the principal to change her decision from the no investment default option to investment. On the other hand, an agent very likely to be smart may correctly change the principal's decision and raise her profit. The *implicit* incentive system here is straightforward: the agent is fired at the end of the first stage if $\hat{\eta} < \eta_0$. If $\hat{\eta} \ge \eta_0$, the he is retained and his wage depends on which segment $\hat{\eta}$ falls in: he gets either a good or a star wage.

The second part of the lemma shows that $w(\hat{\eta})$ is strictly linear when the principal's optimal action depends only on the reports she receives regardless of type. In this case, the expected value of an agent's information in the second stage can be shown to be $w = (g-b)(p_1-r)\hat{\eta}+$ constant. In the above example where state b is sufficiently bad, the cutoff value is determined by the usefulness of the average agent's report sequences, namely, $\bar{\eta} \propto (gr + b(1-r))$. When $gr + b(1-r) = 0$, or report sequences (m_0^g, m_1^g) or (m_0^b, m_1^g) from an average agent yield expected profit of exactly zero, then the principal will invest if there is *any* probability that the agent is smart. In this case $\bar{\eta} = 0$

[9] Formally, this requires $\frac{r}{2}g + \frac{1-r}{2}b \le 0$, and $(1-p_0)p_1 g + p_0(1-p_1)b > 0$. The first inequality means that report sequence (m_0^b, m_1^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.

occurs and $w(\hat{\eta})$ becomes a linear function.

Albeit simple, this lemma shows that the reduced form approach used in many reputational concerns models where the agents maximize the posterior probability that he is smart because his future wage is linear in such probability, is a special case (Prendergast and Stole 1996, Scharfstein and Stein 1990). 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 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

In this model, 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).

Let $I = (I_0, I_1)$ and $M = (M_0, M_1)$ denote respectively the set of the agents' signals and his reports in the first stage, then the strategy of the agent is $\Sigma = (\Sigma^0, \Sigma^1)$, where $\Sigma^0 : \Theta \times I_0 \to \Delta(M_0)$ and $\Sigma^1 : \Theta \times I_0 \times M_0 \times I_1 \to \Delta(M_1)$. The agent's strategy in the second stage is similarly defined. Thus the equilibrium consists of a triple $(\sigma^*, a^*, \hat{\eta})$ such that:

$$\sigma^*(\theta, I) = \text{argmax}_m \; w(\hat{\eta}); \text{ and } a^* = \text{argmax}_{a \in \{0,1\}} \; \Pi(a, m),$$

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.

3 Equilibrium Information Revelation

Section 3.1-3.2 categorize equilibrium strategy of the principal and the agent, focusing on how information revelation in equilibrium depends on both the initial difference and the improvement in the agent's 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 his true signals and the principal chooses the profit maximizing action. The reason is that the agent's wage $w(\hat{\eta})$ does not depend on his second stage performance and he has no further career concerns because the second stage is the end of his career. Thus it is assumed that, from now on, this truthtelling equilibrium is always played in the second stage. The more interesting question is the agent's equilibrium behavior in the first stage.[10]

3.1 Partial Information Revelation: When Type H's Signals Improve Faster

This section focuses on how much information the agent reveals truthfully in equilibrium when type H agent's signal quality improves faster than that of type L's, in a fashion defined below. In particular, whether and how the existence of smart agents affects the average agent's reports in equilibrium.

In the first stage, assume that agent of both types report his first signal i_0 truthfully (shown later to be part of an equilibrium strategy). Without loss of generality, the agent's continuation strategy after receiving i_1 is either: always report true i_1; or always repeat $m_0 = i_0$ again.[11]

[10] In cheap talk games, there always exists a plethora of uninformative, "babbling" equilibria in which the agent randomizes messages in a way unrelated to his type and the principal ignores these messages. 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.

[11] Agents can use other strategies; for example, they can always report 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, since the meaning of messages in cheap talk games is endogenously determined in equilibrium. Thus a message "my signal indicates a good state" may mean many different things in different equilibria. For example, 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 tells the truth. Thus limiting attention to the aforementioned two strategies is without loss of generality.

Note 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 his second signal truthfully; a "full pooling equilibrium" in which both types simply repeat their initial report, and finally, a "partial information revelation equilibrium" in which the agent plays a convex combination of A) and B), with possibly different weight across types.

This subsection focuses on the case when the smart agent's signal quality improves faster than average one. Since both type H and type L agent receive signals of increasing quality, it is necessary to define a measure of signal quality improvement. A smart agent is considered to improve faster than an average one if:

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

while an average agent is considered to improve faster if inequality (1) does not hold. The above inequality compares the confidence of agents in their second signal *relative* to the first when the two signals disagree.[12] 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.

Suppose that both types of agent report truthfully, it is simple to compare the posterior probabilities that the agent is smart given his reports and the observed true state:

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

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. That is, *if the agent reports truthfully*: 1). given

[12] Note that inequality (1) implies that type H 's second signal is also better than that of type L: or $p_1 \geq r$. In the more general case when type L's first signal is informative, i.e., $r_0 \geq \frac{1}{2}$, then inequality (1) is easily modified to $\frac{r_0(1-r_1)}{r_1(1-r_0)} \geq \frac{p_0(1-p_1)}{p_1(1-p_0)}$.

the correct final report, a change of mind is bad for agent's reputation because it means that he is wrong at the beginning; 2). being consistently wrong is worse than a wrong change of mind because L is more likely to get two wrong signals in a row; 3). given consistent reports, being correct indicates good ability. Comparison of these posteriors suggests that both *accuracy* and *consistency* of reports indicate high ability.

Define p_0^L as the cutoff value of p_0 such that type L is indifferent between repeating his first report or sending his second report truthfully when $i_0 \neq i_1$. In the current model with career concerns, the following proposition characterizes the equilibria with sequential reporting:

Proposition 1 *(1.1). A full revelation equilibrium exists where both type H and L report $m_0 = i_0, m_1 = i_1$ if p_0 is sufficiently close to $\frac{1}{2}$.*
(1.2). When type H learns relatively faster, there exists a partial information revelation equilibrium if $p_0 \geq p_0^L$. In this equilibrium, type H always reports truthfully, i.e., $m_0 = i_0, m_1 = i_1$. Type L always reports $m_0 = i_0$. Type L reports $m_1 = i_1$ if his signals agree, i.e., $i_0 = i_1$. If $i_1 \neq i_0$, L repeats m_0 with probability $\pi^* > 0$ and reports $m_1 = i_1$ with probability $1 - \pi^*$.
(1.3). When type L learns relatively faster, and $p_1 \geq r$, there exists a pooling PBE in which both type H and L always report $m_0 = m_1 = i_0$ when $p_0 \geq p_0^L$.

Proposition 1 shows that when type H learns faster, he always reports his true signals, while type L may repeat his first report (which is his first signal) with some probability when his signals disagree, but report truthfully otherwise. Two key factors determine the agent's equilibrium strategy: accuracy of the final report and consistency of reports.

First, the agent is naturally concerned about the accuracy of the final report, which influences his truthtelling incentives through his *private* estimate of the true state, i.e., $Pr(s|i_0, i_1, \theta)$. When signals disagree, the more the agent is convinced about the true state of the world, the less attractive repeating his first report becomes. The reason is that, in this case, lying and repeating the first report is very likely to lead to a consistently wrong sequence of reports, which gives the lowest reputational payoff. For example, when r is close to $\frac{1}{2}$, L repeats his first report with higher probability because his second signal is not very informative, thus repeating his first report may give him the highest posterior $w(CR)$ half of the time. On the other hand, if r is sufficiently high, then the more L attempts to report consistently when he receives conflicting signals, the more he will appear consistently wrong. *Ceteris paribus*, the faster the agent learns, 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 report but who finds out about the truth after all.

Second, the higher quality the smart agent's first signal is, the more likely the average agent prefers repeating his first report because consistently correct reports are increasingly likely to signal type H. When $p_0 \approx \frac{1}{2}$, both H and L have almost uninformative signals i_0 and the final report is more indicative of ability. Thus inconsistency is not a bad signal while giving a wrong final report is. As p_0 increases, a correct first signal is more likely to reflect high ability. As a result, type L is tempted to repeat his first report with probability π^* after receiving conflicting signals to appear smart when type H's first signal is quite accurate.

More subtly, type H agent reports more truthfully than type L not only because type H receives better signals than type L in *absolute* terms. That is, $p_0 \geq \frac{1}{2}, p_1 \geq r$ is not enough. In fact, after receiving conflicting signals, the net gain in the quality of H's information may be *smaller* than that of type L. For example, when $p_0 \approx p_1$ 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. This shows that always receiving better signals is not enough for truthful reporting, rather, a smart agent is more truthful only when he improves faster as defined by inequality (1).

Models without improvement in signal quality is a special case. Suppose that the agent's initial signal is as accurate as his second one, then both types of agent have the same estimate of the true state after receiving conflicting signals because his 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 $\frac{1}{2}$. For p_0 accurate enough, this would lead to a pooling equilibrium in which both types simply reiterate their first report, as in the last part of Proposition 1.[13] As a result, the agent's informative second signal is unused.

This partial revelation equilibrium may be highly inefficient: the principal's information may deteriorate significantly even if there are very few smart agents in the population. The reason is that an average agent may repeat his first uninformative report with a high probability to appear consistent despite a high quality second signal to try to signal he is type H. The following example demonstrates the phenomenon:

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, almost all agents are average, and their second signal is very accurate. In equilibrium, however, type L agent repeats his first report with probability $\pi^*(p_0) = 9.9822p_0 + 1.7796 \times 10^{-2} p_0^2 - 9.0$. It is obvious that π^* increases in p_0, and when $p_0 = 0.95, \pi^* = \frac{1}{2}$. Hence a type L agent lies against his highly informative signal i_1 half of the time even if only one out of a thousand agents is smart.

3.2 Value of Inconsistent Reports

This section focuses on whether the sequencing of the reports alone may signal different ability in models with improving signal quality and implicit incentives, before the agent can be judged based on the accuracy of his report, 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 $Pr(\theta = H|m_0 = m_1) > Pr(\theta = H|m_0 \neq m_1)$ and inconsistency more otherwise.

Proposition 1 shows that the principal may receive consistent reports because the agent receives consistent signals, which increases the probability that $s = m_0 = m_1$; or because the average agent pretends to be consistent. The following proposition describes when consistent reports signal higher ability and when mind changes do.

[13] When both types of agent report true first signal, the expected future wage after reporting consistent signals is $w(CR) + w(CW)$, which is always larger than that of reporting truthfully $w(R) + w(W)$. This means that both H and L have an incentive to deviate and repeat the first report when they receive conflicting signals. As a result, there can not exist a full revelation equilibrium. Among pooling equilibria, L is indifferent as to which i_1 to report regardless of i_0 and H is indifferent when he receives conflicting signals. Therefore it is assumed that they repeat their first report when they are indifferent and thus the principal does not learn anything new from the second report.

Proposition 2 *In the first stage game, if there exists a partial information 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:*

(2.1) the market values consistent reports more than mind changes before observing the state when $\pi^* \leq (2p_0 - 1)(2p_1 - 1)$, which occurs when the principal's optimal decision problem is type-independent, i.e., when $w(\hat{\eta})$ is linear.

(2.2) the market values mind changes more than consistent reports before observing the state when $\pi^* > (2p_0 - 1)(2p_1 - 1)$. This occurs when the payoff function $w(\hat{\eta})$ is sufficiently convex.

Many existing models show that agents prefer to appear consistent because it signals that they are smart. The second part of the above proposition thus 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, in this model with improvement in signal quality, the market may value accurate reports disproportionally more than somewhat accurate ones when the wage function is sufficiently convex. As can be seen from Figure 2, for the type L agent who receives inconsistent signals, repeating his first report may lead either to the best future wage $w(CR)$ or the worst future wage $w(CW)$ with probability $1 - r$ and r respectively. If he follows his true second signal and gives inconsistent reports, he receives $w(R)$ or $w(W)$ with probability r and $1 - r$ respectively. In other words, consistent reports are riskier than inconsistent ones and therefore tends to give higher reputational payoff in expectation. The tradeoff, of course, is that type L has a smaller probability of getting the high reward when he lies.

In equilibrium, a type L agent therefore reports consistently with positive probability such that he is indifferent between consistency and inconsistency in term of expected payoffs. The smart agent, on the other hand, reports his second signal truthfully because his second signal is sufficiently superior to his first signal that repeating his first signal is very likely to lead to the worst payoff of being consistently wrong.

Example 3: Market values mind changes more. The following example illustrates the changes in mixing probabilities, when the payoff function becomes increasingly convex. Let $\theta = 0.1$, $p_0 =$

0.8, $p_1 = 0.9$; and $r = 0.6$. And as an approximation, the payoff function is of the type $w(\eta) = \eta^{\alpha}$, then the following table can show that when $\alpha = 5$, the market values mind changes more:

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

Proposition 2 shows that sequencing of the reports itself, e.g. whether it is consistent or not, signal different levels of abilities, even *before* the true state is observed. Therefore requiring sequential reports provides the principal with a valuable tool before the agent can be evaluated on his accuracy. In fact:

Remark: Sequential reporting may improve the principal's first stage decisionmaking by improving her estimate of the agent's type before the state is observed.

Whether the principal's improved knowledge about the agent's type, based on the report history, can improve her first stage decisionmaking depends on the particular problem: it will not occur when her first stage decision problem is type-independent. In this case the only thing matters to the principal's optimal action is the reports themselves, and even though she forms different opinions based on the report sequence, it bears no impact on her action. On the other hand, when her first stage decision problem is very type-dependent, then the principal may take different actions based on the content of the reports, as well as her belief of the agent's type, inferred from the sequencing of the reports. Therefore the *sequencing of the reports* may lead the principal to choose different actions.

Example 1, continued: Recall that 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 the average agent is very tempted by the star wage $w(CR)$ in the second stage and thus repeats his first report with a high mixing probability π^* such that mind changes in the first stage signal higher ability, the principal may invest when she hears reports (m_0^b, m_1^g), which is more likely to come from a smart manager and are true signals in equilibrium, but not to invest otherwise.

Proposition 2 and the above remark show that the sequencing of reports may improve the principal's first stage decisionmaking by offering her more information about the agent's type. In contrast, suppose that the principal only asks for one report at the end of the first stage, since the states are symmetrically distributed, the principal cannot gain more information about the agent than her priors, that is, $Pr(H|m^g) = Pr(H|m^b) = \eta$. Thus her first stage expected profit may be lower that that after sequential reports.

3.3 Symmetric Information in a Two Signal Model

The two signal model shows that mind changes can be a good sign of the agent's ability in equilibrium, because a smart agent is more confident about the relative accuracy of his second signal and thus changes his mind more. This conclusion differs from many existing models which assume that the agent's type is symmetric information. This subsection turns to a two signal model with symmetric information to show that the premium on inconsistent reports is not only driven by type H's faster improvement in signal quality, but also by his private knowledge of his type.

To study the equilibrium with symmetric information, observe that this differs from the model with asymmetric information in that the agent does not know his type and his first signal is uninformative about type by assumption. Thus his private belief of the state, $Pr(s|i_0, i_1)$, depends on his updated belief of his ability given the signal sequence. The main results of the symmetric information case are the following:

Proposition 3 *With symmetric information, in the first stage,*

(3.1) If $p_0 < \hat{p}_0$, there exists a full revelation equilibrium in which the agent always reports $m_0 = i_0, m_1 = i_1$;

(3.2) If $p_0 > \hat{p}_0$, $p_1 > r$, and the principal's second stage decisionmaking is type independent, then there exists a pooling equilibrium in which the agent always repeats his first report.

(3.3) Before the state is observed, the market always values consistent reports more than mind changes, i.e., $Pr(\theta = H|m_0 = m_1) > Pr(\theta = H|m_0 \neq m_1)$.

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 revelation equilibrium in Section 3: when $p_0 \approx \frac{1}{2}$, type H's first signal is not very accurate and mind changes 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 p_0, p_1, η, r, the cutoff point $\hat{p}_0$ is larger than p_0^L, the cutoff point in the asymmetric information case when type H improves faster than type L. The reason is that if the agent receives inconsistent signals in the symmetric information case, he is more confident in the relative quality of the second signal vis-a-vis the first than type L, but less confident than type H in the asymmetric case. The uncertainty of one's type renders type H less likely to "go out on a limb" and report i_1 and type L to behave more boldly than they would have if their types were known.

Second, with symmetric information, if the agent repeats his first report with some probability π, the updated probabilities that he is smart differ from those in the asymmetric information case. Recall that in the asymmetric information case, mixing by type L dilutes the expected reputational payoff of consistency simply because consistent reports may result from lying of the low type. Here since both types of agent mix when $p_0 > p_0^L$, the expected value of consistent messages tends to increase with π. In other words, when p_0 is large, H is much more likely to be correct in the first signal and consistent reports indicate high ability. Thus the more the agent is lying, the higher the value of consistency becomes and the agent prefers to repeat his first report.

As a result, the most important insight from the symmetric information case is that inconsistent reports fail to be a good sign of ability in equilibrium. The reason 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 $p_0 \geq \frac{1}{2}$. This is because H is relatively rare in the population and H is less likely to be inconsistent. Thus inconsistent reports indicate that the agent is less likely to be H than the principal's prior η. On the other hand, the minimum probability that the agent is smart, conditional on consistent reports, is η because when agents always repeat their first reports, the probability that they are smart is the population average. The probability of being H only increases when the agent repeats his first report with higher probability, because reporting consistent signals are a good sign, and relatively more H reports consistent signals.

This section reconciles the different conclusions of this paper with some of its predecessors showing that the market values consistent reports in yet a different way. That is, in professions where one's talent is commonly known, consistency is more valued while in professions where talent is crucial to success and is private information, mind changes may be more valued because they signal relative confidence in one's later, improved assessments.

4 Optimal Reporting Systems

Section 3 shows that when the agent's signal quality improves at different speed, requiring sequential reports may induce an average agent to lie and report more consistently to appear smart. As illustrated in Example 2, the inefficiency due to type L's career concerns can be quite significant. One natural alternative is for the principal to require report(s) only *after* the agent has received all of his signals. That is, the principal may adopt a hands-off approach: asking the agent for a final assessment and acting on it.

This section therefore addresses two normative questions. First, it investigates when the principal receives higher expected profits with progress reports and when she is better off with a final report. Second, it asks whether the agent is willing to remain silent initially in equilibrium if the principal requires only one final report, but the agent can *voluntarily* give an initial report via either formal or informal channels.

4.1 Optimality of Sequential Reporting

The principal can require one final report or a final report sequence after the agent has learned his both signals. Formally, she may ask for one final report m^f about what the state is in stage $N = 0$. Or she could ask for a vector of reports $\vec{m}^f = (m_0, m_1)$ on what signals he has received. This section compares the sequential reporting system with these alternatives in order to understand when and under what conditions one particular reporting system is more preferred by the principal.

When one final report is required, the agent will be judged solely on its accuracy, which in turn determines his second stage wage. It is easy to see that the agent should report his best estimate of the state based on his signals. When a vector of final reports is required, however, the results are more subtle:

Lemma 2 *(1). In the first stage, if the principal only requires one report m^f, then there exists an* *equilibrium in which both types of agent report* $m^f = i_1$*, regardless of their first signal. (2). If the principal requires report* $\vec{m}^f = (m_0, m_1)$ *after the agent receives both reports, and if inequality (1) holds, then there exists an equilibrium in which both types of agent report* $m_0 = m_1 = i_1$*, regardless of* i_0*.*

Lemma 2 shows that, in this system, the principal always receives one truthful final report from her agent, who no longer needs to worry whether his reports are consistent or not. In comparison with the sequential reporting system, where the first report is truthful but an average agent may repeat his first uninformative signal with some probability to appear smart, it may seem that requiring one final report is always better. Recall Proposition 1, when average agent learns relatively faster even though the smart agent receives signals of higher quality in absolute terms, there exists a 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 useful for the principal than that in a pooling equilibrium of the sequential reporting process. Therefore, one insight is that when the average agent actually improves faster, the principal should ask for a final report, rather than focusing on the path of reports.

The final report m^f, however, is not a sufficient statistic for the two reports in the sequential reporting model because it fails to convey how strongly the agent believes in the state he reported. That is, under this reporting system, the agent only reveals which state he thinks is more likely, but not how much more likely. For example, a report of $m^f = g$ may result from two signal sequences (m_0^b, m_1^g) and (m_0^g, m_1^g), which can lead to quite different estimates on the likelihood of the good state. In the sequential reporting case, the principal may invest when she hears (m_0^b, m_1^g) and not invest otherwise, as discussed in Example 1. If she faces a single report $m = g$ instead, then the expected profit if she invests becomes: $[p_1\eta + r(1-\eta)]g + [(1-p_1)\eta + (1-r)(1-\eta)]b$, which may be negative for η small and the principal will never invest in the first stage.[14]

The alternative reporting system is to require a vector of report at the end, $\vec{m}^f = (m_0, m_1)$, about the signals the agent received.[15] This reporting system may seem to overcome the shortcom-

[14] The example assumes that $rg + (1-r)b < 0$. Therefore for small η, the expected profit is negative.

[15] Note that if the principal asks the agent for one final report with probabilistic assessment of the state, there can only be four different distributions determined by the signals the agent receives. Thus requiring probabilistic final

ing of the above one: it may be able to convey the strength of the agent's belief in the true state. Therefore it warrants a closer look at the amount of information the agent reveals in equilibrium of this reporting system.

Despite the seeming similarity, the model where the principal requires a sequence of reports at the end differs markedly from the sequential reporting process in Section 3. The key difference is that the agent has a larger choice set: that is, he can send any report that gives him highest reputational payoff in the next stage.[16]

The second part of the above lemma shows that full revelation cannot be an equilibrium in this model. The reasoning is familiar: consistent reports will signal higher ability and both types of agent would deviate and give more consistent reports, except that now they repeat their second signal, as opposed to the first in the sequential reporting case studied in Section 3. Moreover, this observation holds even for the parameter values such that a full revelation equilibrium exists in the sequential reporting case, say when $p_0 \leq p_0^L$. Because now there is no previous commitment to defend, the agent can choose his best estimate to appear consistent on. Intuitively, it is harder for the principal to elicit truthful reports because now the agent does not have to commit to an early position from his early report and he can modify his reports in any way to increase the probability of being perceived smart.

Given the above results, the following proposition compares the optimality of the two reporting systems, one with sequential reporting as given in section 3, and one with a final report (vector of reports) only.

Proposition 4 *(4.3) requiring only final report(s) is optimal if 1) the principal's decision problem in both stages are type-independent; or 2) when type L learns faster as defined by inequality (1) and $p_1 > r$; 3) when type H learns faster, but η, fraction of type H is sufficiently small and $p_0 > p_0^L$. (4.4) if type H improves faster and the principal's decision problem is type dependent, then the principal may make better decision with sequential reporting than with final report(s).*

report is equivalent to requiring a vector of final reports.

[16] Formally, given a signal sequence, the agent can deliver any of the three report sequences other than the true signals. In total, each type agent has twelve incentive constraints to satisfy to tell the truth, much more than the four in the sequential reporting case.

Second, distinctly different from the previous case, now type H agent is more likely to report consistently when he learns faster. The reason is exactly the same: he has more relative confidence in his second signal, therefore he is more confident that his consistent reports are more likely to be correct. Therefore type L agent appears more consistent too and the result is that both simply report their second signal. Therefore in equilibrium, requiring a final report is quite similar to requiring a vector of reports on the agent's signals. The agent reveals his second signal truthfully or with some probability.

Combine the current section and Remark 1, the sequential reporting is valuable in the first stage when the decision problem is quite type dependent because of its screening role and its benefit on the first stage decision.

4.2 Voluntary Disclosure of Initial Signal

Suppose that the proportion of type H is very small, type H improves faster, and there exists a partial information revelation equilibrium as shown in Example 1. Then it is optimal for the principal to require one final report, as discussed in the previous subsection. The principal, however, may not be able to forbid voluntary information transmission. For instance, the agent may casually mention what he thinks in an informal gathering. Then the second question arises: would anyone have an incentive to report his first signal anyway? Formally, the message space after the first signal becomes $M_0 = \{m_0^g, m_0^b, \emptyset\}$: the initial report $m_0 = \emptyset$ if the agent does not report voluntarily. For simplicity, let $p_1 = 1$ and $w(\hat{\eta}) = \hat{\eta}$ throughout this subsection.

Proposition 5 *With voluntary disclosure,*
(5.1) If $p_0 < p_0^L$*, then there does not exist a PBE in which the agent reports* $m_0 = \emptyset$*,* $m_1 = i_1$ *unless the principal holds extreme belief* $\kappa \equiv Pr(H|m_0 \neq \emptyset) \leq \kappa^*, \kappa^* < \eta$*.*
(5.2) If $p_0 \approx 1$*, there does not exist a PBE in which the agent reports* $m_0 = \emptyset$ *if the principal holds belief* κ *close to* η*.*

Intuitively, if the high type's first signal is much more accurate than the low type's, a correct first report is more likely to suggest high ability. By giving two reports, type H can benefit from his higher likelihood of giving consistently correct reports. Thus unsolicited revelation of early information may signal high ability in itself. If type H prefers to reveal his initial signal, the absence of an initial report implies low ability and type L has to give early report as well, even though his first signal is uninformative.

Note that reporting $m_0 = \emptyset$ can always be part of a PBE profile. The reason is similar to the previous subsection: the agent's future wage is the posterior of his ability. However, if the principal is extremely pessimistic about anyone who sends her an initial report, the agent will never report his initial signal. Thus if the principal is not too extreme in her beliefs, both types of agent report their first signal. As a result, sequential reporting can emerge endogenously in equilibrium with voluntary revelation.

5 Extension: Role of Additional Informative Report

Consider the scenario when the agent receives a longer sequence of signals and the principal is not under time pressure to make the decision. Should the principal ask for reports after each of the agent's signal? What is the impact of the agent's earlier reports on his later incentive to report truthfully? This section studies the change in the agent's truthtelling incentive when the principal requires one additional signal before taking action. In particular, it examines the relationship between his *earlier* commitment and *later* commitment to a position. For instance, suppose that he has lied to appear consistent in the second report, would the agent lie more in the third report to defend his early reports or lie less to try to give a correct final report?

Since the agent's final signal is informative, the additional report necessarily improves the principal's investment decision if he reports at least partially truthfully. Moreover, the more informative the third signal is relative to other signals, the more valuable the additional information tends to be. However, due to reputational concerns, there are other, indirect effects in additional to the direct beneficial effect, which 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 benefit of extra information. Intuitively, in the two signal case, type L takes a bet if he lies to appear consistent. The addition of a third report may increase the expected value of such a bet, in which case type L wants to lie more in the second report; or it may decrease the value and type L would prefer to be more truthful in the second report.

5.1 Early Consistency and Late Consistency

A simple three signal model is used to facilitate comparison with the two signal model. 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.

It is shown in Lemma 3, Appendix B that there exists a continuation equilibrium in the third report such that type H always reports the true i_2. The low type always reports the true signal if $i_2 = m_1$. When $i_2 \neq m_1$, the average agent mixes with different probabilities depending on the path of his past reports details). Then in the three signal model:

Proposition 6 *(6.1) Whenever type H agent finishes improving, i.e, $p_t = 1$, there is no need to require further reports because both type will give consistent reports from t on;*
(6.2) Given the semi-pooling continuation equilibrium in the third report, type H reports the first and the second signal truthfully. Type L reports the first signal truthfully. In the second report, L reports true i_1 if $i_1 = m_0$. When $i_1 \neq m_0$, L repeats the first report with π_1 if $p_0 \geq p_0'$.

Observe that if $p_1 = 1$, then L has to report $i_1 = i_2$ with probability one to imitate H. In this case the principal does not gain any new information from the third signal, even if it is the best signal type L would ever receive. This holds more generally: if the smart agent learns the true state s at signal i_t, then the principal cannot gather new information because type L has to repeat m_t until the last report in order to look like type H with some positive probability. Thus the principal should not require further reports when type H has finished learning, even if there are few smart agents in the population and the signal L would have received otherwise is highly accurate.

In the present model, type L can imitate H in two ways: consistency and (eventual) accuracy. Consistency is primarily driven by the high level of H 's information while correctness is driven by their high slope. In a very short process such as the two signal game, an average manager has only one chance to balance the two attributes. With a longer sequence, if the type H's final signal is relatively more superior to the early ones, type L can afford to change his mind and become accurate later without large loss in term of reputation. Then the low type is apt to separate the two attributes: appear consistent first, and if necessary, change mind and become accurate. This effect dampens the use of the third signal to the principal. On the other hand, if the H type's learning is slow, even though the gain in term of additional information is small, the agent becomes more accurate because late change of mind is costly both in reputation and in accuracy. Thus the interest of the agent becomes more aligned with that of the principal, who should require the third signal.

On one hand, the three signal model shows that, counterintuitively, the principal may *not* want to require the last report when the smart agent's signal quality improves a lot in the final signal. 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, 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.

On the other hand, the principal may want to request the final report when the improvement in the smart agent's signal quality *levels off*. The reason is that in equilibrium, an average agent may lie against his true third signal 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. Moreover, early mind changes are valuable with a longer sequence of reports in another way: if the agent has reported truthfully in his initial reports, early inconsistency tends to imply more truthful reports later in equilibrium for the low type.

6 Conclusion

In many situations, progress reports are required before important actions are taken. In professions where implicit incentives are important, however, agents who provide these progress reports may behave too consistently. As shown in earlier work, they may repeat their early assessment (or other early mover's assessment) to appear capable even when they have information to the contrary as time goes by (Scharfstein and Stein 1990). This poses two questions: first, are consistent reports always more valued by the decisionmakers even then the agents may improve in their ability to observe the true state of the world? Second, should decisionmakers ask for the agents' opinion only after he finishes receiving signals or should they require progress reports?

This paper investigates the role sequential reports play when agents of different ability improve in their ability to observe a state of the world. Agents receive multiple signals of ascending quality, thus both the *sequencing* of reports and the accuracy of the reports may become a signal of ability. Contrary to the past work, this paper shows that even though an average agent may appear more consistent by repeating their initial reports, the market may value mind changes more than consistency: mind changes reflect confidence in improvement in signal quality.

This result hinges on two key factors. First, the smart agent reports more truthfully in equilibrium not because they receive signals of higher qualities than the average ones in *absolute* terms, but because they improve faster and thus have more confidence in their later signals relative to their early ones than an average one. The speed of signal quality improvement is crucial: consistency may be more valued in markets with career concerns and little signal quality improvement, but inconsistency can be the prized sign of the fast learners and the talented. Second, the agent's type is private information, thus they can base their reports on their signals and how good they are. 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 they tend to be not the smart type if they receive inconsistent signals.

Inefficiency may be quite high when there are few smart types in the population but all average ones may discard their second, more informative signal with some probability to appear smart. Therefore, whether and when the principal prefers sequential reporting is the next set of results.

Basically, sequential reporting can give the principal more 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. Also, it helps when the principal's decision problem is very sensitive to the accuracy of the reports. Sequential reporting helps the principal to gauge how strongly the agent believes in the state he reports, rather than just the direction of his belief, which can improve her first stage decision problem.

Moreover, the agent's incentive to report truthfully may be history dependent. A decisionmaker may not desire additional information when the quality improvement is large because it increases average agent's incentive to lie in the earlier periods. When the quality improvement of additional information is small, however, the market benefits from asking for one more report: it offers potentially new information and reduces the average agents' lying in earlier assessments. The principal needs to be careful about the indirect effect of asking for more reports on agent's earlier reports.

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 agents themselves are risk neutral, the implicit incentive structure itself is generally convex and may thus encourage risk taking behavior, especially on the part of the average agents. How this type of implicit incentive structure evolves over time could be a question of interest. Another question is that when the principal should ask for reports and how many. Requiring very early report may cause the average agents to commit to an early opinion without knowing much just to appear smart; requiring late reports only is likely to make it difficult for the decisionmaker to learn the strength of the agent's belief. 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. How to design such reporting systems is a question of further research.

APPENDIX A: PROOFS

Proof of Lemma 1:

1. First, recall that $w(\hat{\eta}) = \sum_{m^1} \Pi^1(\hat{\eta}, m^1)|_{a=a^*} \equiv V(a^*(\hat{\eta}))$. First, $w(\hat{\eta})$ is convex. 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 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 $\hat{\eta}$:

$$\begin{aligned}
\Pi^1(i_0^g, i_1^g) &= \frac{r}{2}g + \frac{1-r}{2}b + \hat{\eta}\left[g(p_0 p_1 - \frac{r}{2}) + b((1-p_0)(1-p_1) - \frac{1-r}{2})\right] I(a); \\
\Pi^1(i_0^b, i_1^g) &= \frac{r}{2}g + \frac{1-r}{2}b + \hat{\eta}\left[g((1-p_0)p_1 - \frac{r}{2}) + b(p_0(1-p_1) - \frac{1-r}{2})\right] I(a); \\
\Pi^1(i_0^g, i_1^b) &= \frac{1-r}{2}g + \frac{r}{2}b + \hat{\eta}\left[g(p_0(1-p_1) - \frac{1-r}{2}) + b((1-p_0)p_1 - \frac{r}{2})\right] I(a); \\
\Pi^1(i_0^b, i_1^b) &= \frac{1-r}{2}g + \frac{r}{2}b + \hat{\eta}\left[g((1-p_0)(1-p_1) - \frac{1-r}{2}) + b(p_0 p_1 - \frac{r}{2})\right] I(a),
\end{aligned}$$

where $I(a)$ is an indicator function defined in the text. 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. In the current model, it is assumed that type H learns faster than L and inequalit) holds. For each of the profit function, 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 $\hat{\eta}$.

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

Proof of Proposition 1:

To simplify notation, the following shorthand expressions denote different posterior probabilities that the agent is smart, given his reports and the observed true state. The posterior probability when the reports are consistent and right is denoted by CR, i.e., $m_0 = m_1 = s$. Similarly, CW denotes the posterior when the reports are consistent and wrong; R denotes the posterior when there is a correct change of mind; and finally, W when there is a wrong change of mind. Furthermore, these posterior beliefs of the principal are updated by Bayes' rule. Formally,

$$
\begin{aligned}
CR &\equiv Pr(H|m_0^g, m_1^g, g) = \frac{p_0 p_1 \eta}{p_0 p_1 \eta + \frac{1}{2}[r + (1-r)\pi](1-\eta)} \\
CW &\equiv Pr(H|m_0^g, m_1^g, b) = \frac{(1-p_0)(1-p_1)\eta}{(1-p_0)(1-p_1)\eta + \frac{1}{2}[(1-r) + r\pi](1-\eta)} \\
W &\equiv Pr(H|m_0^g, m_1^b, g) = \frac{p_0(1-p_1)\eta}{p_0(1-p_1)\eta + \frac{1}{2}(1-r)(1-\pi)(1-\eta)} \\
R &\equiv Pr(H|m_0^g, m_1^b, g) = \frac{(1-p_0)p_1\eta}{(1-p_0)p_1\eta + \frac{1}{2}(1-\pi)r(1-\eta)}.
\end{aligned}
$$

First, suppose that in equilibrium both types of agent report $m_0 = i_0$ truthfully. There are four incentive constraints in the second period when the agent decides what to report:

$$
\begin{aligned}
&(IC_1^L) \quad (w(CR) - w(W))Pr(g|i_0^b, i_1^g, L) < (w(R) - w(CW))Pr(b|i_0^b, i_1^g, L); \\
&(IC_2^L) \quad (w(CR) - w(W))Pr(g|i_0^g, i_1^g, L) > (w(R) - w(CW))Pr(g|i_0^g, i_1^b, L); \\
&(IC_1^H) \quad (w(CR) - w(W))Pr(g|i_0^b, i_1^g, H) < (w(R) - w(CW))Pr(b|i_0^b, i_1^g, H); \\
&(IC_2^H) \quad (w(CR) - w(W))Pr(g|i_0^g, i_1^g, H) > (w(R) - w(CW))Pr(b|i_0^g, i_1^g, H).
\end{aligned}
$$

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:

$$
Pr(g|i_0^g, i_1^g, H) = \frac{p_0 p_1}{2p_0 p_1 + 1 - p_0 - p_1} \geq Pr(g|i_0^g, i_1^g, L) = r; \quad \Pr(g|i_0^b, i_1^g, H) = \frac{p_1(1-p_0)}{p_0 + p_1 - 2p_0 p_1} \leq \frac{1}{2} \leq r.
$$

Observe that 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. That is, when type L is indifferent between reporting the true i_1 and repeating $i_0 = m_0$, type H prefers to tell the truth.

Next, consider IC_1^L: $LHS \leq RHS$ at $p_0 = \frac{1}{2}, \pi = 0$ because $CR = R, W = CW$. Moreover, at $\pi = 0$, the left hand side of IC_1 increases with p_0 while the right hand side decreases with p_0.***revise in term of the restrction on η***. Because w is a piecewise-linear function, the right derivative is used when w is not differentiable. ***need to show the increasing part*** Since $w' \geq 0$:

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

Thus for all $p_0 \in [\frac{1}{2}, p_0^L)$, IC_1^L holds strictly at $\pi = 0$. For all $p_0 \geq p_0^L$, IC_1^L cannot hold because its left hand side strictly exceeds its right hand side. 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 $\pi^* > 0$ such that $LHS = RHS$.[17] When IC_L^1 binds, inequality (1) implies that IC_H^1 holds strictly, i.e., type H still prefers to report truthfully if he receives inconsistent signals. Thus we have a partial pooling continuation equilibrium when $p_0 \geq p_0^L$.

Second, 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. As in the text, we require that the agent reports truthfully when indifferent. Type H prefers to report $m_0 = i_0$ if the following incentive constraint is true:

$$w(CR)p_1 - w(R)p_1 + w(W)(1-p_1) - w(CW)(1-p_1) \geq 0 \tag{2}$$

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 2 always holds and both types report $m_0 = i_0$ in equilibrium. ‖

Proof of Proposition 2:
Step 1: in a partial information revelation equilibrium, for any given p_0, p_1, the mixing probability π decreases in r. It is obvious that if the agent reports truthfully, both CR and R decrease in r while CW and W increase in r. Moreover,

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

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

Therefore the expected value of giving consistent reports decreases with r and that of giving inconsistent reports increases with r. Intuitively, the larger is r, the smaller is the low type's incentive to lie and give consistent reports. Therefore the mixing probability π decreases with r.
Step 2: In a partial information revelation equilibrium, $CR \geq R$. Otherwise the low type should reduce the mixing probability to appear more inconsistent. From step 1, the highest mixing probability is obtained for any given p_0, p_1 at $r = \frac{1}{2}, CR = R$. Substitute into formulas for CR and R, we have all $\pi^* \leq (2p_0 - 1)$. Denote $\overline{\pi} = (2p_0 - 1)$.
Step 3: 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:

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

[17] The mixing probability π^* solves:

$$\frac{p_0(1-r)(p_1 - r - \pi + r\pi)}{[p_0p_1\eta + \frac{1}{2}(r+\pi-r\pi)(1-\eta)][p_0(1-p_1)\eta + \frac{1}{2}(1-\pi)(1-r)(1-\eta)]}$$
$$= \frac{(1-p_0)r(p_1 - r + r\pi)}{[(1-p_0)p_1\eta + \frac{1}{2}r(1-\pi)(1-\eta)][(1-p_0)(1-p_1)\eta + \frac{1}{2}(1-r+r\pi)(1-\eta)]}$$

Moreover, it is easy to show that β decreases in π while γ increases in it.

Step 4: In step 3, we have decomposed the market's estimate of ability before state realizes into convex combinations of the agent's posteriors. Recall that in a partial information revelation continuation equilibrium, IC_1^L binds at π^*, therefore $CR(1-r)+CWr = Rr + W(1-r)$. Since $r \geq \frac{1}{2}$, it is obvious that at $\pi^*, CR + CW \geq R + W$. Hence when $\beta \geq \gamma$, simple calculation can show that consistency is valued more by the market.

At the highest mixing probability possible, $\overline{\pi}$, we can show that:

$\beta(\overline{\pi}) \geq \gamma(\overline{\pi})$, so $Pr(H|m_0, m_1 = m_0) \geq Pr(H|m_0, m_1 \neq m_0)$ for all π^*.

Suppose that in equilibrium, $\pi^* \geq (2p_0-1)(2p_1-1)$, then the market values inconsistency more. However, from step 1-4, market values consistency more when $\pi^* \in [(2p_0-1)(2p_1-1), (2p_0-1)]$, Contradiction. Therefore when w is linear, all mixing probability $\pi^* < (2p_0-1)(2p_1-1)$ and consistency is more valued.

Second, if $D(\pi^*) \leq 0$, or equivalently 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 proof 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 not type-independent, w is a convex piecewise-linear function, as discussed in the text. 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 m_0^g, m_1^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, $\pi^* > (2p_0-1)(2p_1-1)$.

We can also approximate w by a continuous convex function $w(\hat{\eta}) = (\hat{\eta})^{\alpha}$, then it is easy to see that π increases with α: the more convex the payoff function is, the larger is π. ‖

Proof of Proposition 3:

First, find the continuation equilibria assuming that the agent would tell the truth in the first report. Suppose that the first report is 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:

$$[w(Pr(H|m_0^g, m_1^g; g)) - w(Pr(H|m_0^g, m_1^b; g))]Pr(g|i_0^g, i_1^g) \geq [w(Pr(H|m_0^g, m_1^b; b)) - w(Pr(H|m_0^g, m_1^g; b))]Pr(b|i_0^g, i_1^g)$$
$$[w(Pr(H|m_0^g, m_1^g; g)) - w(Pr(H|m_0^g, m_1^b; g))]Pr(g|i_0^g, i_1^b) \leq [w(Pr(H|m_0^g, m_1^b; b)) - w(Pr(H|m_0^g, m_1^g; b))]Pr(b|i_0^g, i_1^b)$$

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

$$Pr(g|i_0^g, i_1^g, L) < Pr(g|i_0^g, i_1^g) = \frac{p_0 p_1 \eta + \frac{r}{2}(1-\eta)}{[p_0 p_1 + (1-p_0)(1-p_1)]\eta + \frac{1}{2}(1-\eta)} < Pr(g|i_0^g, i_1^g, H)$$
$$Pr(g|i_0^g, i_1^b, H) < Pr(g|i_0^g, i_1^b) = \frac{p_0(1-p_1)\eta + \frac{1-r}{2}(1-\eta)}{[p_0(1-p_1) + (1-p_0)p_1]\eta + \frac{1}{2}(1-\eta)} < Pr(g|i_0^g, i_1^b, L)$$

From 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, 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.

Second, when $p_0 > \hat{p}_0$, the second IC does not hold and the agent is tempted to appear more consistent by repeating his first report. However, since $Pr(H|m_0^g, m_1^g, g) = \frac{p_0[p_1+(1-p_1)\pi]\eta}{p_0[p_1+(1-p_1)\pi)\eta + \frac{1}{2}[r+(1-r)\pi](1-\eta)}$, the left hand increases with the mixing probability π, thus there exists a pooling equilibrium in which the agent always repeat his first report. ‖

Proof of Lemma 2:

First, if only report m^0 is required, then the principal's posterior estimates of the agent's ability become:

$$Pr(H|m^0 = s) = \frac{p_1\eta}{p_1\eta + r(1-\eta)}; \quad Pr(H|m^0 \neq s) = \frac{(1-p_1)\eta}{(1-p_1)\eta + (1-r)(1-\eta)}.$$

As discussed in the text, in the second stage the agent will report both signals truthfully. Therefore the wage function is still a convex, increasing and piece-wise linear function as shown in Lemma 1. Since $Pr(H|m^0 = s) > Pr(H|m^0 \neq s)$, then $w(Pr(H|m^0 = s)) > w(Pr(H|m^0 \neq s))$ as well.

Second, given that the agent receives higher wage if his final report is accurate than otherwise, and the agent's second signal is more accurate than his first, the agent's expected wage after reporting $m^0 = i_1$ is:

$$\begin{aligned} & w(Pr(H|m = i_1 = s))Pr(s = i_1, \theta) + w(Pr(H|m^0 = i_1, i_1 \neq s))Pr(s \neq i_1, \theta), \\ \geq \quad & w(Pr(H|m = i_1 = s))Pr(s \neq i_1, \theta) + w(Pr(H|m^0 = i_1, i_1 \neq s))Pr(s = i_1, \theta) \quad \text{for both} \quad \theta \in \{H, L\} \end{aligned}$$

Therefore there exists an equilibrium in which both agents report $m^0 = i_1$. ‖

Proof of Proposition 4:

When the principal requires $m = (m_0, m_1)$ after the agent receives both signals, then for an agent to tell the truth, he needs to receive a higher wage in the second stage than all other three report sequences can bring him. For example, after receiving consistent signals i_0^g, i_1^g:

$$\begin{aligned} & W(Pr(H|(g,g);g)Pr(g|i_0^g, i_1^g, \theta) + W(Pr(H|(g,g);b)Pr(b|i_0^g, i_1^g, \theta) \\ \geq\ & W(Pr(H|(b,b);b)Pr(b|i_0^g, i_1^g, \theta) + W(Pr(H|(b,b);g)Pr(g|i_0^g, i_1^g, \theta) \\ & W(Pr(H|(g,g);g)Pr(g|i_0^g, i_1^g, \theta) + W(Pr(H|(g,g);b)Pr(b|i_0^g, i_1^g, \theta) \\ \geq\ & W(Pr(H|(g,b);g)Pr(b|i_0^g, i_1^g, \theta) + W(Pr(H|(g,b);b)Pr(b|i_0^g, i_1^g, \theta) \\ & W(Pr(H|(g,g);g)Pr(g|i_0^g, i_1^g, \theta) + W(Pr(H|(g,g);b)Pr(b|i_0^g, i_1^g, \theta) \\ \geq\ & W(Pr(H|(b,g);g)Pr(g|i_0^g, i_1^g, \theta) + W(Pr(H|(b,g);b)Pr(b|i_0^g, i_1^g, \theta) \end{aligned}$$

And after receiving inconsistent signals, say i_0^b, i_1^g,

$$\begin{aligned} & W(Pr(H|(b,g);g)Pr(g|i_0^b, i_1^g, \theta) + W(Pr(H|(b,g);b)Pr(b|i_0^b, i_1^g, \theta) \\ \geq\ & W(Pr(H|(g,g);g)Pr(g|i_0^b, i_1^g, \theta) + W(Pr(H|(g,g);b)Pr(b|i_0^b, i_1^g, \theta) \\ & W(Pr(H|(b,g);g)Pr(g|i_0^b, i_1^g, \theta) + W(Pr(H|(b,g);b)Pr(b|i_0^b, i_1^g, \theta) \\ \geq\ & W(Pr(H|(g,b);g)Pr(g|i_0^g, i_1^g, \theta) + W(Pr(H|(g,b);b)Pr(b|i_0^b, i_1^g, \theta) \\ & W(Pr(H|(b,g);g)Pr(g|i_0^b, i_1^g, \theta) + W(Pr(H|(b,g);b)Pr(b|i_0^b, i_1^g, \theta) \\ \geq\ & W(Pr(H|(b,b);g)Pr(g|i_0^g, i_1^g, \theta) + W(Pr(H|(b,b);b)Pr(b|i_0^b, i_1^g, \theta) \end{aligned}$$

Simple calculations can show that $CR + CW \geq R + W$ at truthtelling.

Then use the proof of Lemma 2. ‖

Proof of Proposition 5:

(5.1) Suppose that there exists a PBE in which the agent reports $m_0 = \emptyset$ and $m_1 = i_1$. The expected reputational payoff under the putative equilibrium for type H and type L are respectively:

$$\eta_h^1 = \frac{p_1^2\eta}{p_1\eta + r(1-\eta)}, \quad \eta_l^1 = \frac{p_1 r\eta}{p_1\eta + r(1-\eta)}.$$

If the agent deviates and reports his first signal, then since $p_0 < p_0^L$, he adopts the full revelation equilibrium strategy as analyzed in Proposition 1. The expected payoff of type H and L becomes:

$$\kappa_h^1 = \frac{p_0^2\kappa}{p_0\kappa + \frac{1}{2}r(1-\kappa)} + \frac{(1-p_0)^2\kappa}{(1-p_0)\kappa + \frac{1}{2}r(1-\kappa)},$$
$$\kappa_l^1 = \frac{\frac{1}{2}p_0 r\kappa}{p_0\kappa + \frac{1}{2}r(1-\kappa)} + \frac{\frac{1}{2}r(1-p_0)\kappa}{(1-p_0)\kappa + \frac{1}{2}r(1-\kappa)}.$$

Simple algebra can show that for $\kappa > \kappa^*, \kappa^* < \eta$, $\eta_h^1 < \kappa_h^1$ and $\eta_l^1 > \kappa_l^1$. In other words, type H agent strictly prefers to report his first signal and the L type strictly prefers not to. Therefore all type H will deviate and report his first signal. Type L cannot report $m_0 = \emptyset$ either, because then the principal knows he is average for sure. In this way we break the putative equilibrium.

Proof to claim (5.2) is similar to that of (5.1)except that in this case the agent will play the partial information revelation equilibrium as in Section 3. Hence we also have to take into account the mixing probability π^*. ‖

Proof of Proposition 6:

First, see Lemma 3 in Appendix B for the continuation equilibrium in the third report. Then, given Lemma 3, two key truthtelling IC conditions for type L in the second report are:

$$\begin{aligned}
& (Pr(H|i_0^g, i_1^g, i_2^g, g)(r_2 + (1-r_2)\pi_2^2) - Pr(H|i_0^g, i_1^b, i_2^g, g)r_2)r_1 \\
\geq\ & (Pr(H|i_0^g, i_1^b, i_2^b, b)r_2 - Pr(H|i_0^g, i_1^g, i_2^b, b)r_2(1-\pi_2^2))(1-r_1)) \\
& (Pr(H|i_0^g, i_1^g, i_2^g, g)(r_2 + (1-r_2)\pi_2^3) - Pr(H|i_0^g, i_1^b, i_2^g, g)r_2(1-\pi_2^1))(1-r_1) \\
\leq\ & (Pr(H|i_0^g, i_1^b, i_2^b, b)(r_2 + (1-r_2)\pi_2^1) - Pr(H|i_0^g, i_1^g, i_2^b, b)r_2(1-\pi_2^3))r_1
\end{aligned}$$

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 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 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 we can show that 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 $\pi_1=0, \pi_2^2(0), \pi_2^1(0)$. ‖

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.

The first part guarantees the existence of a partial revelation equilibrium when the principal only requires two signals. The second part 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:

$Pr(H|m_0^g,m_1^g,g)=\frac{p_0p_1\eta}{p_0p_1\eta+\frac{1}{2}r(1-\eta)}$, and

$Pr(H|m_0^g,m_1^g,m_2^g,g)=\frac{p_0p_1p_2\eta}{p_0p_1p_2\eta+\frac{1}{2}r_1r_2(1-\eta)}$

Simple algebra can show that $Pr(H|m_0^g,m_1^g,m_2^g,g)>Pr(H|m_0^g,m_1^g,g)$. 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 posteriors after a wrong report decrease in the three signal case. For example, $Pr(H|m_0^b,m_1^b,m_2^b,g)<Pr(H|m_0^b,m_1^b,g)$. 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 IC constraints are as follows:

$$\begin{array}{ll} IC_1^i & Pr(H|m_0^g, m_1^b, m_2^g, g)Pr(g|i_0^g, i_1^b, i_2^g, \theta) \geq Pr(H|m_0^g, m_1^b, m_2^b, b)Pr(b|i_0^g, i_1^b, i_2^g, \theta) \\ IC_2^i & Pr(H|m_0^g, m_1^g, m_2^g, g)Pr(g|i_0^g, i_1^g, i_2^b; \theta) \leq Pr(H|m_0^g, m_1^g, m_2^b, b)Pr(b|i_0^g, i_1^g, i_2^b, \theta) \\ IC_3^i & Pr(H|m_0^g, m_1^g, m_2^g, g)Pr(g|i_0^g, i_1^b, i_2^b; \theta) \leq Pr(H|m_0^g, m_1^g, m_2^b, b)Pr(b|i_0^g, i_1^b, i_2^b, \theta) \end{array}$$

Of the above three IC 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 (m_0^g, m_1^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: $\underbrace{\overbrace{Pr(H|m_0^g, m_1^g, m_2^g, g)}^{reputational\ concerns}\overbrace{Pr(g|i_0^g, i_1^g, i_2^b, \theta)}^{statistical\ concerns}}_{expected\ reputational\ payoff}$. 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:

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

The picture is more complicated for L: because high type is far more likely to be correct and consistent in i_1 and i_2, low type'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 (i_0^g, i_1^g, i_2^b) differs from that after private signals (i_0^g, i_1^b, i_1^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 3 (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 $Pr(H|m_0^g, m_1^g, m_2^g, g)$ and $Pr(H|m_0^g, m_1^g, m_2^b, b)$ is not too large, the low type mixes more in the continuation equilibrium when his first two true reports agree than they disagree.[18] 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 $Pr(H|m_0^g, m_1^g, m_2^g, g)$ and $Pr(H|m_0^g, m_1^g, m_2^b, b)$ 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 $Pr(H|m_0^g, m_1^g, m_2^g, g)$ and $Pr(H|m_0^g, m_1^g, m_2^b, b)$ 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 $Pr(H|m_0^g, m_1^g, m_2^g, g)$. 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 that type L receives a second signal different from his initial report m_0^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 IC constraint is needed:

[18] 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.

$$[Pr(H|m_0^g, m_1^g, m_2^g; g)\overbrace{[r_2 + (1-r_2)\pi_2^3]}^{future\ strategy} - Pr(H|m_0^g, m_1^b, m_2^g; g)\overbrace{r_2(1-\pi_2^1)}^{future\ strategy}\ ]\quad \overbrace{(1-r_1)}^{dir.\ accuracy\ loss}$$
$$\leq [Pr(H|m_0^g, m_1^b, m_2^b; b)[r_2 + (1-r_2)\pi_2^1] - Pr(H|m_0^g, m_1^g, m_2^b; b)r_2(1-\pi_2^3)]r_1$$

The left hand side describes the net benefit of repeating the first report m_0^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 m_0^g, he receives potentially big gains $Pr(H|m_0^g, m_1^g, m_2^g; g)$ or big loss $Pr(H|m_0^g, m_1^g, m_2^b; b)$ in his reputations. If he reports i_1^b truthfully, he receives gains $Pr(H|m_0^g, m_1^b, m_2^b; b)$ or suffers losses $Pr(H|m_0^g, m_1^b, m_2^g; g)$.

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 $Pr(H|m_0^g, m_1^g, m_2^g; g)$ 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 $Pr(H|m_0^g, m_1^g, m_2^g; g)$ 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 *), 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 $(1-r)Pr(H|m_0^g, m_1^g, m_2^g, g) \leq rPr(H|m_0^g, m_1^g, m_2^g, g)$ 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 $(1 - r)Pr(H|m_0^g, m_1^g, m_2^g, g)$ and $rPr(H|m_0^g, m_1^g, m_2^g, g)$ 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.

Dewatripont, M., and J. Tirole (1994): *The Prudential Regulation of Banks*. MIT Press, Cambridge, MA.

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.

Holmström, B. (1999): "Managerial Incentive Problems: A Dynamic Perspective," *Review of Economic Studies*, 66(1), 169–182, first published in 1982.

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

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., 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.