

Feedback on Ideas

Zeinab Aboutalebi Ayush Pant
Department of Economics, The University of Warwick *

October 2, 2019

Abstract

Employees are often assigned tasks comprising two distinct phases: in the first phase, ideas are generated; in the second phase, the best idea is implemented. Furthermore, it is common for supervisors to give feedback to their employees during this process. This paper studies the supervisors problem. Supervisors face the following tradeoff: while honest feedback encourages employees to discard bad ideas, it can also be demotivating. We obtain three main results. First, the supervisor only gives honest feedback to agents who believe in their ability to succeed. Second, receiving honest feedback leads such high self-opinion agents to exert more effort. Third, overconfidence is potentially welfare improving.

Keywords: Experimentation; Feedback; Dynamic cheap-talk

JEL Classification Numbers: C73, D82, D83, O32, M5

*Email: A.Pant@warwick.ac.uk and Z.Aboutalebi@warwick.ac.uk. Telephone: +44(0) 7473372277, +44 (0) 7787203822. Address: Department of Economics, The University of Warwick, Coventry CV4 7AL, UK. This work was supported by the Economic and Social Research Council [grant number ES/J500203/1]. We would like to thank Robert Akerlof, Motty Perry, Kobi Glazer and Ilan Kremer for their time and insightful comments. We are also grateful to Dan Bernhardt, Costas Cavounidis, Daniel Habermacher, Qingmin Liu, Phil Reny, Federico Trombetta, and seminar participants at Warwick Microeconomic Theory Work-in-Progress seminars, European Winter Meeting of Econometric Society 2017 and Israel Game Theory Conference 2018 for additional comments and clarifications. All errors are our own.

1 Introduction

Employees are often assigned tasks with two distinct phases: in the first phase, ideas are generated; in the second phase, the best idea is implemented. Furthermore, it is common for supervisors to give feedback to their employees during this process. For instance, a partner in a law firm supervises an associate developing a litigation strategy, a project manager in a technology firm supervises an engineer solving a bug in app development, and a senior designer in an architecture firm supervises a junior designer looking for a design solution. One can trace such examples of feedback and supervision outside of corporate organizations as well; for instance, a professor supervising her grad student in a university.

This paper studies the supervisor’s problem. Supervisors face the following trade-off. On the one hand, honest feedback encourages employees to discard bad ideas. On the other hand, such feedback can be demoralizing and discourage both idea generation and effort implementation. We build a model to describe how this trade-off shapes the supervisor’s feedback, the employee’s effort, and the employee’s trust in the supervisor.

We consider a supervisor-agent model with two phases: experimentation and implementation. In the experimentation phase, the agent sequentially generates ideas at a cost, receives feedback from the supervisor on her ideas, and selects an idea to implement. In the implementation phase, the agent decides how much effort to put into completing her chosen idea. The agent’s ability is initially unknown, and the agent and supervisor share a common prior. Importantly, we assume the supervisor does not internalize the agent’s cost of effort. This misalignment of preferences means that dishonesty is a possibility.¹

Ability plays a central role in our model. We assume a high-ability agent both generates and implements ideas better than a low-ability agent. As a result, the agent’s self-opinion (belief about her ability) affects both the agent’s decision regarding how much to experiment and her choice of implementation’s effort. Both of these effects, in turn, impact the supervisor’s feedback.

There are three key findings of our model. First, the supervisor never gives a low self-opinion agent honest feedback because doing so is demotivating: it discourages effort in both the experimentation and implementation phases. When negative feedback discourages further experimentation, the supervisor prefers to falsely encourage the agent to induce her to put a higher effort in implementation instead. Therefore, negative feedback is only forthcoming for a high self-opinion agent. Moreover, a high self-

¹Note that if providing feedback is costly to the supervisor (such as time costs) this could realign the principal’s and agent’s interests, thereby restoring honesty. We show that the supervisor is more (less) honest when he is more (less) time constrained, and therefore less (more) willing to supervise.

opinion agent, independent of her actual ability, is repeatedly informed about her bad ideas and can end up being “treated more harshly”.

Second, receiving supervisor feedback magnifies performance differences between high and low self-opinion agents. Because high self-opinion agents receive honest feedback, they have confidence both in their ability and in the quality of their ideas, which leads to high effort. Low self-opinion agents, in contrast, lack confidence, which leads to low effort. Receiving more honest feedback with a higher self-opinion allows the agent not only to experiment more but also to exert an optimal effort in implementing her chosen idea. Such an opportunity might not be available to a slightly lower self-opinion agent because she does not receive honest feedback as often. As a result, she has lower confidence in her idea. Therefore she might end up exerting too much effort on a bad idea, and too little effort on a good idea.

Third, overconfidence can be welfare improving. The discontinuous change in the supervisor’s feedback strategy as the agent incorrectly goes from a low self-opinion to a high self-opinion gives rise to this possibility. The cost of overconfidence in ability is that it leads to too much effort exertion. However, the benefit of overconfidence is that it can lead to honest feedback. This benefit may outweigh the cost.

Our results find support in The Sensitivity to Criticism Test from PsychTests which collected responses from more than 3,600 participants.² The study revealed that those who tended to be defensive about negative feedback had lower performance ratings and lower self-esteem. Moreover, managers were skeptical to give feedback to workers who get defensive. “If there was an esteem problem, both men and women seemed to block out the constructive part of the equation and only focus on the criticism”, revealed a manager. This further meant that the manager would rather “develop the more (talented and) mature employee,” instead of spending time counselling him or her. These ideas also appear in the situational leadership theory developed by Paul Hersey and Ken Blanchard in mid-1970s. According to Ken Blanchard, “Feedback is the breakfast of champions.”

Related Literature. Our paper relates to two distinct strands of literature: experimentation and dynamic communication games. Within experimentation, our work falls under models of motivating experimentation. Previous research has looked at how information can be optimally delivered to the agents arriving sequentially to experiment (such as Kremer, Mansour, and Perry (2014) and Che and Horner (2015)) or at how information should be designed for a single agent to motivate her to experiment (such as Renault, Solan, and Vieille (2017) and Ely (2017)).³ Among the two, our setting

²<https://eu.usatoday.com/story/money/columnist/kay/2013/02/15/at-work-criticism-sensitivity/1921903/>

³See Hörner and Skrzypacz (2016) for a survey on the recent advancements in experimentation and information design.

falls in the latter category. Ely and Szydlowski (2017), Smolin (2017) and Ali (2017) are the closest in this respect.⁴ In each of these papers, a principal must reveal information by balancing the positive effect of good news with the discouraging effect of no or bad news. Nonetheless, none of these papers can address the situation in many examples where ex-ante commitment to a disclosure rule is not possible, and feedback can improve the result of the experiment. How the same tradeoff shapes the honesty in strategic feedback with no commitment is our point of departure from these papers. Thus, our model is one of communication rather than information design. To the best of our knowledge, we are the first to study such settings without commitment.

In our setting, the supervisor tries to motivate the agent to exert effort in both the experimentation and implementation phases. As a consequence, honest feedback can discourage the agent at two levels. The first is stopping experimentation too early, and the second is exerting low effort in implementation. Introducing this novel objective makes our setting unique in feedback and experimentation literature.

Some older papers like Lizzeri, Meyer, and Persico (2002) and Fuchs (2007) have looked at feedback in dynamic settings without experimentation and show that often it is not optimal to provide feedback. Orlov (2013) considers a setting in which providing information to the agent might benefit the principal in the short-run but may lead to long term agency costs. There the principal designs an optimal information sharing rule along with a compensation scheme. Boleslavsky and Lewis (2016) also study dynamic settings with commitment in which the agent has new information every period. The principal makes sequential decisions, after which he observes a private signal of the state. These works consider the effect of feedback in settings with commitment but no experimentation. Our paper connects these two types of literature in a no-commitment setting.⁵

The other strand of literature related to our work is dynamic communication games. A few papers like Aumann and Hart (2003), Krishna and Morgan (2004), Forges and Koessler (2008) and Goltsman, Hörner, Pavlov, and Squintani (2009) look at repeated communication with an action at the end. Our setup is different in that the receiver should decide after each round whether she wants to experiment again. Golosov, Skreta, Tsyvinski, and Wilson (2014) and Renault, Solan, and Vieille (2013) are closer in this sense. They look at situations where the receiver decides after every round of communication. However, neither has the above-stated feature of persuasion in two phases.

In this respect, our work relates to dynamic persuasion games. Morris (2001), Honryo (2018) and Henry and Ottaviani (2019) are a few papers that do not assume commitment by the sender of

⁴Some other related papers have looked at settings in which a sender commits to dynamic information design to influence a receiver. See, for example, Bizzotto, Rüdiger, and Vigier (2018).

⁵Orlov, Skrzypacz, and Zryumov (2018) is an exception. They look at commitment and no commitment case in a setting in which an agent tries to convince the principal to wait before exercising a real option. Again, however, their model does not have experimentation.

information. The seminal paper by Morris (2001) deals with a potentially biased advisor persuading a decision-maker to choose actions dynamically when reputation matters. Honryo (2018) and Henry and Ottaviani (2019), however, are closer to our setting. In these papers, a sender (entrepreneur or researcher) tries to persuade a receiver (venture capitalist or publisher) to take a favorable action by sequentially disclosing some verifiable or costly information. However, we can generate a tradeoff for the sender without assuming verifiability or costly information transmission. In our model, when the supervisor persuades the agent to experiment again, he inadvertently also persuades her to exert lower effort in implementation. It is this feature that creates the main honesty/dishonesty tradeoff in our model.

Finally, our result on the importance of beliefs in final performance is related to some of the older research starting with Bénabou and Tirole (2002). This vast line of economics research is itself based on the original psychology research of Bandura (1977). However, such research usually looks at the importance of belief absent any external supervision. The presence of a supervisor drives our results on the effect of higher self-opinion and overconfidence.⁶

The rest of the paper is structured as follows. In Section 2, we describe the basic model. In Section 3, we solve two benchmark cases of the model without supervision, which help us build intuition and solve the complete game. Then, in Section 4, we present the main analysis of the game with a supervisor without commitment. We move onto presenting how our results are qualitatively the same in a few extensions and offer new interpretations of our model in Section 5. We conclude in Section 6 by revising the main results and identifying the further scope of research.

2 The Model

We consider a setting in which an *agent* (she) works on a project and a *supervisor* (he) is responsible for providing feedback. The project involves two distinct stages that proceed sequentially. The first stage is *planning* or *experimenting with ideas*, and the second stage is *execution* or *implementation of a chosen idea*. The agent is responsible for both experimenting with and implementing ideas for the completion of the project. The supervisor has no commitment power and provides feedback based on what he observes.

Stage 1: Idea generation. The process of idea generation involves multiple rounds $t = 1, 2, \dots$. In each round t , the agent decides whether she wants to draw a new idea. The quality of an idea

⁶Koellinger, Minniti, and Schade (2007) and Hirshleifer, Low, and Teoh (2012) are two papers that empirically show the importance of overconfidence in the context of innovation and creativity.

is determined by its *ex-ante potential to succeed* θ_t which could be either high (\hbar) or low (ℓ). The distribution of θ_t is given by

$$\theta_t = \begin{cases} \hbar & \text{with probability } \alpha, \\ \ell & \text{otherwise} \end{cases}$$

where α is the *ability* of the agent. $\alpha \in \{0, q\}$ where zero is “low”, and $q \in [0, 1]$ is “high”. Therefore, only a high-ability agent can come up with a high potential idea, which happens with probability q . The ability (unlike the idea) remains persistent throughout the play. The agent and the supervisor only know the distribution of the ability; neither observes it. The belief that the agent is high-ability at the beginning of round t is denoted by β_t , with a common prior $\beta_1 \in (0, 1)$ at the beginning of the game in round 1. For much of the text, we use belief and self-opinion interchangeably. We assume that the agent possesses a low potential outside option idea at the beginning in round 1 denoted by $\bar{\theta} = \ell$.

Actions and timing: In each round of experimentation the agent chooses $I_t \in \{0, 1\}$. $I_t = 0$ denotes the agent’s decision to stay in Stage 1 and experiment with another idea in round t (i.e., not implement). c is the *cost of experimentation* which could arise from searching the Internet, looking up for data, reading material, previous works, and seeking inspiration.

Importantly, we assume that only the supervisor can see the potential of the idea generated. The supervisor privately observes θ_t and chooses an announcement about its observed potential, $m_t \in \{\ell, \hbar\}$.⁷ We initially assume limited recall of the agent and the supervisor so that they only talk about the last idea produced (and not the entire history of past ideas). We present the analysis of perfect recall in which the supervisor is allowed to make backdated messages in Section 5.2.

Alternately, the agent could decide to implement the last idea after the supervisor’s message. This is denoted by $I_t = 1$.

Stage 2: Idea implementation. If the agent decides to move to the idea implementation stage in $t + 1$ following the last message of the supervisor m_t , then her idea gets fixed at $\theta \equiv \theta_t$.

Actions and timing: The agent chooses effort $e \in [0, 1]$ at cost $\frac{e^2}{2}$ to complete the project. The final outcome of the project, success or failure, is determined according to the following distribution

⁷We can also start with an arbitrary message space M but since we consider a game of cheap talk with binary types and we focus on pure strategy equilibria, what matters are the equilibrium mappings from the supervisor type (what potential of the ideas he observes) to the message space, i.e. what is the meaning of the messages. Here, messages ℓ and \hbar have their natural meaning and are understood as the potential of the idea developed.

function

$$\Pr(\text{success}) = \begin{cases} e & \text{if } \theta = \hbar, \\ ke & \text{if } \theta = \ell \text{ and agent is high-ability where } k \in (0, 1), \\ 0 & \text{otherwise.} \end{cases}$$

The probability of success is a function of the potential of the chosen idea θ , effort exerted by the agent e and the ability of the agent α . It must be noted that only the high-ability agent is capable of successful completion of the project, but success may still be obtained with a low potential idea. Therefore, when the ability is unknown there is an incentive to implement a low potential idea instead of experimenting again.

We will make the following assumption for mathematical convenience.

$$q \geq (q + (1 - q)k)^2 \geq k \tag{A}$$

Intuitively, this assumption implies that in case the agent has a low potential idea, the supervisor finds it beneficial for the agent to experiment than to implement that idea (with the maximum possible effort of 1). Further, an additional round of experimentation with feedback is preferred to an additional round of experimentation without feedback. We explain these ideas further when presenting the main analysis in Section 4.⁸

Payoffs: Completion of the project yields V . If the completed project is successful, it yields a normalized value of 1, and zero otherwise. The payoff of the agent is given by

$$u_A = V - Tc - \frac{e^2}{2}$$

where T is the number of rounds for which the agent has experimented. The payoff of the supervisor is given by

$$u_S = V.$$

The payoffs highlight the incentive misalignment between the agent and the supervisor. While both players prefer success over failure, the agent alone bears the cost of experimentation and implementation.

Once the payoffs are realized, the game ends. A summary of the timing of the game is provided in Figure 1. We provide an alternate interpretation of the model and additional examples in Section 5.3.

⁸This assumption helps simplify the proofs by providing sufficient conditions. In the absence of this assumption, all our proofs go through but will be belief dependent, which makes them less obvious and more cumbersome.

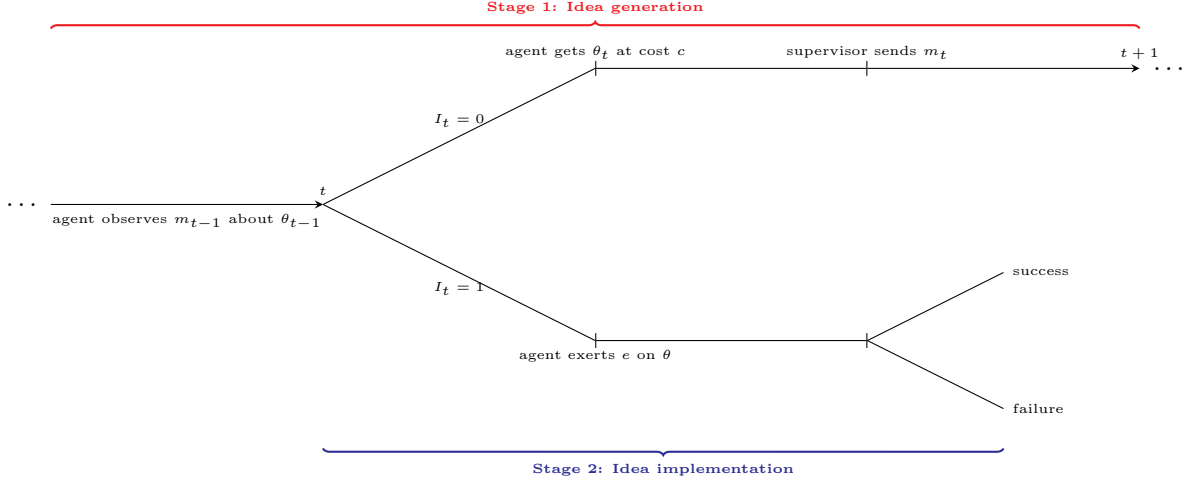


Figure 1: Summary of the timing of the game

We now turn to the analysis of the game. Before we describe the behaviour of a strategic supervisor, we describe the benchmark case in the following section without the supervisor. We then introduce the supervisor in Section 4 and search for honest equilibrium feedback strategies.

3 Benchmark: Single Agent Problem

In this section we look at a setting in which an agent works on the project without any supervision. This preliminary analysis helps us put bounds on the behaviour of the agent and supervisor when they interact with each other as shown in Section 4. Two cases are possible – the agent does not observe the potential of her idea, or she does so perfectly.

3.1 No information (NI) about θ

If the agent does not observe the potential of her idea θ from attempting experimentation at belief β and there is no outside support, then the two alternatives available to her are as follows:

1. The agent can choose to not experiment and directly implement the project using the outside option idea. In this case, the agent $\max_e \beta k e - \frac{e^2}{2}$, which yields a maximized payoff of $\frac{(\beta k)^2}{2}$.
2. The agent can choose to experiment once and then execute the resulting idea. In this case, the agent $\max_e \beta(q + (1 - q)k)e - \frac{e^2}{2} - c$, which gives a maximized payoff of $\frac{\beta^2(q + (1 - q)k)^2}{2} - c$.

Observe that the agent does not want to try experimenting more than once in this setting because experimenting is an additional cost without any added benefit. She will not learn the quality of the

new idea and the odds of coming up with a high potential idea remain unchanged. The only reason she might want to experiment once is to take the gamble of coming up with a high potential idea. She will do so if her belief is high enough. This is illustrated in the following condition:

$$\overbrace{\frac{\beta^2(q+(1-q)k)^2}{2}}^{\text{expected benefit of experimentation}} \geq \overbrace{\frac{(\beta k)^2}{2}}^{\text{cost of experimentation}} + \underbrace{c}_{\text{actual cost}}, \quad (\text{C1})$$

which leads to the following lemma:⁹

Lemma 1 *Let $c < \frac{(q+(1-q)k)^2 - k^2}{2}$. If there is no information about θ , there exists a unique threshold $\beta_0^{NI} := \left(\frac{2c}{(q+(1-q)k)^2 - k^2}\right)^{\frac{1}{2}}$ such that*

1. *if the prior belief $\beta_1 \geq \beta_0^{NI}$ then the agent experiments once before finishing the project by exerting effort $\beta_1(q+(1-q)k)$, and*
2. *if the prior belief $\beta_1 < \beta_0^{NI}$, the agent uses the outside option idea $\bar{\theta} = \ell$ to finish the project by exerting effort $\beta_1 k$.*

In the text we will also be interested in how β_0^{NI} responds to changes in the cost of experimentation c . It is easy to see that a higher cost of experimentation raises this threshold as it reduces the incentives to experiment *ceteris paribus* (see Appendix B for other comparative statics result).

3.2 Full information (FI) about θ

When the agent can perfectly observe the outcome of each round of experimentation, then she potentially wants to experiment at least once. This, as before, depends on her belief about her ability. But now she uses Bayes' rule sequentially to update her belief after observing the potential of the idea from the latest round of experimentation such that

$$\beta_t = \begin{cases} \frac{(1-q)\beta_{t-1}}{1-\beta_{t-1}q} & \text{if } \theta_{t-1} = \ell, \\ 1 & \text{otherwise.} \end{cases}$$

As is standard in good-news models, the agent revises her belief downwards each time she generates a low potential idea, but her belief jumps to 1 if she generates a high potential one. The agent enters

⁹A similar lemma with a belief threshold condition can also be obtained if the agent has no outside option idea. Denote such a cutoff by β_ϕ^{NI} . Then it can be shown that such a cutoff exists and is given by $\beta_\phi^{NI} = \frac{(2c)^{1/2}}{q+(1-q)k}$. Obviously, $\beta_\phi^{NI} < \beta_0^{NI}$. However, we make use of β_0^{NI} in the main analysis – we assume away the possibility of quitting when there is no support from a supervisor.

the implementation phase and finishes the project upon observing $\theta_{t-1} = \hbar$. At this point, she does not have an incentive to experiment further as she only bears an additional cost without any extra benefit. She finalizes the project with the maximum effort of 1 which leads to the project being successful with certainty, and yields a maximized payoff of $\frac{1}{2}$ (the previous cost of experimentation is sunk). Thus, independent of which round of experimentation she is at if $\theta_{t-1} = \hbar$ then $I_t^{FI}(\beta_t = 1) = 1$ with $e^{FI}(\beta_t = 1) = 1$.

On the other hand, after observing $\theta_{t-1} = \ell$ (with the agent observing low potential ideas $\theta_{t'} = \ell$ for all the previous rounds $t' < t - 1$ as well) the agent holds a belief $\beta_t < 1$ about her ability. The agent again faces two choices – to implement the low potential idea or to continue experimenting. If she chooses to implement her low potential idea then she chooses the optimal effort to $\max_e \beta_t k e - \frac{e^2}{2}$. This yields a maximized payoff of $\frac{(\beta_t k)^2}{2}$ where she exerts effort $\beta_t k$ according to her belief β_t . Depending on her belief β_t she might be a high-ability agent with a positive probability of success. If she chooses to experiment once more, then with probability $\beta_t q$ she comes up with a high potential idea and exerts maximal effort of 1 thereafter to finish the project (from above). With probability $1 - \beta_t q$ she comes up with a low potential idea and she faces the same decision problem but with a lower belief $\beta_{t+1} < \beta_t < 1$. Denote the value function of the agent at the beginning of round t with belief β_t when her last observed outcome is $\theta_{t-1} = \ell$ by $\mathcal{V}^\ell(\beta_t)$, such that

$$\mathcal{V}^\ell(\beta_t) = \max \left\{ \frac{(\beta_t k)^2}{2}, -c + \frac{\beta_t q}{2} + (1 - \beta_t q) \mathcal{V}^\ell(\beta_{t+1}) \right\}.$$

Assuming that the agent wants to start experimenting (the condition for which we will outline below), we are interested in if and when the agent stops experimenting with repeated low potential ideas. To do so, let the maximum number of rounds the agent experiments be T . The agent at belief $\beta_T \equiv \beta$ after $T - 1$ rounds will attempt another *final* round of experimentation knowing that irrespective of the outcome she will move to implementing her idea in the following round. So

$$\begin{aligned} \mathcal{V}^\ell(\beta) &= \max \left\{ \frac{(\beta k)^2}{2}, -c + \frac{\beta q}{2} + (1 - \beta q) \mathcal{V}^\ell(\beta') \right\} \\ &= -c + \frac{\beta q}{2} + (1 - \beta q) \mathcal{V}^\ell(\beta') \geq \frac{(\beta k)^2}{2} \end{aligned}$$

where

$$\beta' = \frac{(1 - q)\beta}{1 - \beta q} \text{ and } \mathcal{V}^\ell(\beta') = \frac{(\beta' k)^2}{2},$$

which can be rearranged to

$$\overbrace{\frac{\beta q}{2} + (1 - \beta q) \frac{(\beta' k)^2}{2}}^{\text{expected benefit of experimentation}} \geq \overbrace{\frac{(\beta k)^2}{2} + c}_{\substack{\text{cost of experimentation} \\ \text{opportunity cost} \quad \text{actual cost}}} \quad (\text{C2})$$

Lemma 2 follows from condition C2 and captures the optimal behaviour of the agent under full information about θ . (All proofs are presented in Appendix A.)

Lemma 2 *If there is full information about θ , the optimal decision rule of the agent I_t^{FI} is a unique belief threshold rule such that*

$$I_t^{FI} = \begin{cases} 0 & \text{if } \theta_{t-1} = \ell \text{ and } \beta_t \geq \beta_0^{FI}, \\ 1 & \text{otherwise.} \end{cases}$$

for $c < \frac{q(1-k^2)}{2}$. Further, the optimal effort that the agent exerts to implement her idea is given by

$$e^{FI} = \begin{cases} \beta_{T+1} k & \text{if } \theta_T = \ell, \\ 1 & \text{otherwise.} \end{cases}$$

When $c \geq \frac{q(1-k^2)}{2}$ the agent does not experiment for any belief, and implements her outside option idea with effort $\beta_1 k$.

Figure 2 plots the expected benefit from experimentation (LHS plotted in green) and the cost of experimentation (RHS plotted in red) from condition C2 for different levels of beliefs β . It illustrates the uniqueness result of Lemma 2 under the cost condition $c < \frac{q(1-k^2)}{2}$. Note that both the benefit and the costs are declining in belief about ability. A lower belief in ability means that the agent is less likely to get a high potential idea, which reduces the expected benefit of experimentation. At the same time, for the same reason, it induces the agent to exert lower effort when implementing the outside option idea, thereby reducing the opportunity cost of experimentation. However, the fixed component c of the total costs of experimentation ensures that the costs never go down to zero, which in turn guarantees the existence of the unique threshold.

Observe that the optimal decision rule does not depend on t but only on the belief β , which is a function of the potential of the last observed idea. For a given set of parameters, the maximum number of rounds the agent experiments T is only defined by the prior belief β_1 . The agent wants to




Figure 2: The optimal belief threshold β_0^{FI} for the complete information about θ case

start experimenting with ideas if $\beta_1 \geq \beta_0^{FI}$, and goes on doing so with repeated low potential ideas as long as the belief hits β_0^{FI} . T is therefore determined by how far β_1 is from β_0^{FI} .

It only remains to show how β_0^{FI} varies with a change in parameters. Again, we'll be interested in how β_0^{FI} responds to a change in the cost of experimentation. As expected, an increase in the cost of experimentation raises the threshold belief β_0^{FI} as the agent wants to experiment fewer rounds now (for any prior).

3.3 Comparing β_0^{NI} and β_0^{FI}

Lemma 3 *If $c < \frac{(q+(1-q)k)^2 - k^2}{2}$, then both β_0^{NI} and β_0^{FI} exist and are unique with $\beta_0^{NI} > \beta_0^{FI}$.*

Figure 3 illustrates why $\beta_0^{NI} > \beta_0^{FI}$. It shows that for any belief β the value of experimenting is always lower in the case when the agent has no information about her output of experimentation. Experimentation is merely a gamble to try luck without any learning. This makes the threshold for experimentation higher under the no information case.




Figure 3: Comparing β_0^{NI} and β_0^{FI}

3.4 An important definition

Before moving to the main analysis, we introduce some additional terminology that we will use extensively in the following sections.

Given the no information and the full information belief thresholds β_0^x for $x \in \{NI, FI\}$, define recursively a sequence of belief thresholds $\{\beta_i^x\}_{i=0}^\infty$ such that $0 < \beta_i^x < 1$ and $\beta_{i+1}^x = \frac{\beta_i^x}{1-q(1-\beta_i^x)}$. Starting with the threshold β_0^x the sequence identifies β_1^x , the belief that leads to β_0^x when the agent correctly finds out that her idea has a low potential to succeed, and so on. β_{i+1}^x is the belief which when updated with the correct information about a low potential outcome leads to the belief β_i^x , and this is recursively defined all the way down to the belief β_0^x .

4 Strategic supervisor

4.1 Preliminaries

The game between a strategic supervisor and an agent in Stage 1 is one of dynamic cheap talk. The supervisor can costlessly send either of the two messages independent of the true potential of the idea. Our solution concept is (perfect) Bayesian Equilibrium.

To define the strategies of the agent and the supervisor at any time, we would need to define the history for each player when they are called upon to make a decision. Round t begins for the agent after having observed the last message sent by the supervisor m_{t-1} . Accordingly, a realized history for the agent includes the set of all previous messages sent by the supervisor until and including the last message m_{t-1} and the sequence of past decisions made. Round t begins for the supervisor after observing the last idea of the agent θ_t . Accordingly, a realized history for the supervisor includes, in addition to the history viewed by the agent, the sequence of all the realized idea potential from the past experimentation.¹⁰

For most of the paper, we focus on pure strategy equilibria and limited recall, i.e. we are interested in whether the supervisor is honest with the agent when he can only send a message about the last idea generated. A pure strategy for the supervisor in round t is a mapping from the realized history to the message space $\{\ell, \hat{\ell}\}$. The supervisor is honest with the agent if for any realization of the history the supervisor sends a message that matches the observed potential of the idea. If the supervisor reveals to the agent the outcome of her last experimentation in round t starting from a prior β_t the agent's updated posterior in round $t + 1$ is as in the full information case:

$$\beta_{t+1}^\ell = \frac{(1-q)\beta_t}{1-q\beta_t} \text{ if } m_t = \ell, \text{ and} \quad (1)$$

$$\beta_{t+\tau}^{\hat{\ell}} = 1 \text{ otherwise.} \quad (2)$$

If the supervisor uses the same message independent of the realized history the supervisor is said to lie or babble (see footnote 11). In this case the agent's posterior belief is the same as her prior belief. We will assume that when the supervisor is expected to lie the agent does not consult the supervisor. This rules out the possibility of the supervisor privately learning and not revealing to the agent the outcome, and the arising deviations.

Given our focus on pure strategies and that the two players share a common prior, the agent and the supervisor symmetrically update their belief on the agent's ability. If the agent stops experimenting (and implements her last idea) because the supervisor is babbling, neither the agent nor the supervisor

¹⁰Let $I^t := (I_1, \dots, I_t)$ and $m^t := (m_1, \dots, m_t)$ be the sequence of decisions made by the agent and the public messages given by the supervisor until round t . Define the set of histories for the agent and the supervisor at the beginning of round t by H_t^A and H_t^S respectively. The history for the agent at the beginning of round t is

$$h_t^A = (I^{t-1}, m^{t-1}) \in H_t^A \subset (\{0\}^{t-1} \times \{\ell, \hat{\ell}\}^{t-1}).$$

This is also the public history of the play of the game up to round t . In addition to the public history, the supervisor observes $\theta^t := (\theta_t, \dots, \theta_t)$ and an extra decision of the agent to experiment $I_t = 0$. The history for the supervisor at the beginning of round t is

$$h_t^S = (\theta^t, I_t, h_t^A) \in H_t^S \subset (\{\ell, \hat{\ell}\}^t \times \{0\}^t \times \{\ell, \hat{\ell}\}^{t-1}).$$

have any new information. There is learning only insofar as the supervisor is honest.

4.2 Analysis

What feedback strategy the supervisor employs will depend on how he expects the agent will respond to it, both in the experimentation phase and the implementation phase. We begin by discussing the obvious babbling equilibria. Babbling is always an equilibrium for any prior β_1 in the first stage of the game. The agent does not learn about the true potential of the last idea as the supervisor is always expected to send the same message. This is equivalent to the single agent decision-making problem without advice and Lemma 1 applies. Thus, the agent experiments once before finishing the project if $\beta_1 \geq \beta_0^{NI}$, otherwise she uses the outside option idea to finish the project. Neither supervisor type can profitably deviate from such an equilibrium given the beliefs. The supervisor sends meaningless messages, the agent correctly believes that there is no information content in the recommendations and she makes her decision only on the basis of her prior belief.¹¹

In what follows we determine if pure strategy equilibria in which the supervisor is honest exist, and under what conditions. The approach will be to determine if honest equilibria exist (in addition to babbling) for different ranges of beliefs starting with low ones.¹²


Proposition 1 *For any belief $\beta < \beta_0^{FI}$, any communication strategy is an equilibrium and none induces the agent to experiment.*

From Lemma 3, we know that $\beta_0^{FI} < \beta_0^{NI}$. The region of beliefs $\beta < \beta_0^{FI} < \beta_0^{NI}$ is the one in which the agent does not want to experiment with ideas independent of how much information is provided to her. So all communication strategies are equally informative to the agent and are an equilibrium. The agent does not consult the supervisor in any equilibria as she is very pessimistic about her ability to come up with a high potential idea. She does not want to bear the cost of experimentation at such low beliefs. She simply implements her low potential outside option idea $\bar{\theta} = \ell$ with an effort βk .


¹¹ When the supervisor babbles, it might be useful to think of babbling in mixed strategies rather than in pure strategies (see description of mixed strategies in Appendix A). A supervisor babbling in mixed strategies makes use of both the messages in equilibrium, and the posterior β_n after either message remains unchanged. There are also babbling equilibria in pure strategies. Say the agent conjectures that the supervisor only says $m = \bar{k}$ on-the-equilibrium path. We have that $\Pr(m = \bar{k} | \theta = \bar{k}) = 1 - \Pr(m = \ell | \theta = \ell)$ and a potential babbling equilibrium. While there is no update of beliefs on path, the message $m = \ell$ is off path and we would need to specify beliefs in the information set following this message. Such an equilibrium is supported by any belief $\beta^{\text{offpath}} \in [0, \beta_1)$.

¹²The proofs will be presented in terms of a generic belief β wherever possible. The intuition is the same – whether the agent starts out in the given range with a low potential outside option idea or whether she lands there after continued experimentation (and ending up with a low potential idea that she is aware of), if she finds herself there her behaviour is the same. If she finds herself in any of the ranges with the knowledge that her idea was definitely a high potential idea, then she will always immediately implement her idea by exerting effort 1.

Step 1: Babbling is unique for $\beta_0^{FI} \leq \beta_1 < \beta_1^{FI}$



Step 2: Babbling is unique for $\beta_1^{FI} \leq \beta_1 < \beta_0^{NI}$



Step 3: Babbling is unique for $\beta_0^{NI} \leq \beta_1 < \beta_1^{NI}$




Figure 4: Uniqueness of babbling equilibria for priors $\beta_1 < \beta_1^{NI}$

A concern when evaluating whether the supervisor can be honest for higher beliefs will be what he thinks is the possibility of the agent experimenting again after a negative message. As in the full information case outlined in Section 3.2, the agent experiences a decline in both the benefit and cost of coming up with a new idea after receiving a truthful negative message. With continued discouragement the agent must stop experimenting at belief β_0^{FI} . However, the supervisor's payoff is contingent on the agent's success. This implies that he faces a discontinuous drop in the benefit of being honest at β_0^{FI} , while the cost is that the agent exerts a lower effort in implementation. Our first main result and proposition builds on this intuition. It defines the range of beliefs for which the cost of being honest are higher than the benefits.

Proposition 2 *For any belief $\beta_0^{FI} \leq \beta < \beta_1^{NI}$, babbling is the unique equilibrium strategy.*

The intuition for this proposition is illustrated in steps using Figure 4.¹³

We begin by showing that babbling must be a unique equilibrium strategy of the supervisor in the range of priors $\beta_0^{FI} \leq \beta_1 < \beta_1^{FI}$ (see Step 1 of Figure 4). In this range of priors, a message about the idea being low potential if expected in equilibrium must lead to a posterior about ability $\beta_2^\ell < \beta_0^{FI}$. At this point the agent does not want to experiment any more (from Proposition 1). Moreover, after

¹³Here we discuss the intuition of why honesty cannot be an equilibrium strategy but the proposition is stronger. The argument will also hold to prove that no informative equilibria will survive in this range of beliefs. Our proof in Appendix A presents a general proof that allows for mixed strategies as well.

experimenting and learning that her idea had a low potential to succeed she reduces her effort when implementing the idea. As a result, the expected probability of success further reduces with the low potential idea. This leads the supervisor observing a low potential idea to deviate from honesty and always send a positive message instead.

A positive message is believed by the agent pushing up the posterior of the agent to 1. The agent best responds by implementing the chosen idea with the maximal effort of 1, which increases the expected probability of success with a low potential idea. The supervisor is at the very least able to extract a higher effort on a low potential idea by deviating. Thus, no equilibria in which the supervisor is honest will survive – babbling is unique in this range of priors. In such a babbling equilibrium, the agent best responds by not experimenting because this is identical to a situation with no supervisor and $\beta_1 < \beta_0^{NI}$ (from Lemma 1).

Now, in Step 2 consider the range of priors which when updated with negative messages lead to posteriors below β_1^{FI} . The same argument as the one highlighted above holds because such low posteriors lead the agent to implementing the low potential idea with a lower effort. This time because the supervisor is expected to babble if updated with an honest discouraging message. Therefore, an agent expecting information can be taken advantage of by supervisor type who has only observed low potential ideas. This kills honesty and only the babbling equilibria survive. The same logic can now be extended all the way up to all the prior beliefs which when updated with a discouraging message about the idea lead to posteriors below β_0^{NI} . Below β_0^{NI} the agent does not want to experiment when no information is provided by the supervisor. Such is the case for all prior beliefs $\beta_1 < \beta_1^{NI}$ (illustrated in Step 3).

The total communication breakdown between the supervisor and the agent in this range of beliefs is driven by the fear of the supervisor to discourage the agent to the point of no further experimentation. This is why we call this region of beliefs as those in which the agent has a *low self-opinion*. When he sees that the agent has produced a low potential idea the supervisor finds it beneficial to cajole the agent by calling it a high one, so that at the very least the agent exerts a high effort to implement a low potential idea. But lying is counter-productive as the agent expects the supervisor to only provide fake encouragement; neither does she consult the supervisor nor does she experiment.

This region of beliefs $\beta_0^{FI} \leq \beta < \beta_1^{NI}$ where the agent has a low self-opinion reflect pure inefficiencies in the supervisor-agent relationship. From Lemma 2 we know that the agent would continue experimenting with ideas until she produces a high potential idea for beliefs $\beta \geq \beta_0^{FI}$ if she receives honest feedback. At the same time, the supervisor is also (always) better off with repeated experimentation until a high potential idea is produced. But neither can achieve this better outcome because

the supervisor is unable to commit to honestly revealing the result of the agent's experimentation. Even though the agent is willing to listen to honest feedback, her reaction to negative feedback is too extreme from the supervisor's point of view. If the agent must give up, he prefers she exert the maximum effort instead. Such inefficiency will be a feature of any communication equilibrium we can construct since babbling is unique. The supervisor cannot offer any information in equilibrium. The extent of babbling and that of the resulting inefficiency is determined by the gap between β_0^{FI} and β_1^{NI} , which is a function of the parameters. An increase in the cost of experimentation (c) increases both these thresholds and causes babbling for even higher beliefs (and also no experimentation for higher beliefs). An increase in the probability of generating a high potential idea (q) reduces the region of babbling. An increase in the success rate from implementing a bad idea (k) can *decrease* the inefficiency by reducing the babbling region as it makes the agent want to experiment more without supervision by reducing β_0^{NI} .

Note, however, the difference in the agent's best response to such an uninformative strategy of the supervisor. Since the supervisor babbles in the entire region of beliefs below β_1^{NI} , from Lemma 1 the agent best responds by not experimenting in the region below β_0^{NI} and by experimenting once in the region between β_0^{NI} and β_1^{NI} . This produces an added source of inefficiency when she experiments in this region i.e. when the belief is above β_0^{NI} but below β_1^{NI} . In this case, the agent exerts an inefficient level of effort to implement the idea as she is unable to observe the potential of her idea without honest supervision. She exerts more effort on a low potential idea and a lower effort on a high potential idea.

We are now in a position to determine if there are any honest equilibria. The possibility of honesty opens up for beliefs $\beta > \beta_1^{NI}$ because the agent is now willing to experiment at least once without the supervisor's support. This happens in the region of beliefs between β_0^{NI} and β_1^{NI} . The previous threat point for the supervisor now potentially disappears as the supervisor can guarantee that the agent will experiment even when she is discouraged. In this sense, we call this the region of *high self-opinion*. We are now in a position to analyse whether this one extra round of experimentation (without the consultation of the supervisor) and a high self-opinion is sufficient for the supervisor to be honest.

Proposition 3 For $c \geq \frac{\kappa k - (\kappa k)^2}{2}$ where $\kappa \equiv \frac{k}{(q + (1-q)k)^2}$ and for all $t \geq 1$,

1. truth-telling is an equilibrium strategy for the supervisor for $\beta_t \geq \beta_1^{NI}$, and
2. babbling is the unique equilibrium strategy for the supervisor for $\beta_t < \beta_1^{NI}$.

The agent's equilibrium strategy is given by

$$I_t^* = \begin{cases} 0 & \text{if } m_{t-1} = \ell \text{ and } \beta_t \geq \beta_1^{NI}, \text{ or } \beta_0^{NI} \leq \beta_t < \beta_1^{NI}, \\ 1 & \text{otherwise.} \end{cases}$$

The agent's optimal effort is given by

$$e^* = \begin{cases} 1 & \text{if } m_{t-1} = \hbar, \\ \beta_t(q + (1 - q)k) & \text{otherwise.} \end{cases}$$

Proposition 3 identifies the necessary and sufficient condition for an honest equilibrium to arise in the *entire* region above babbling equilibria, i.e. one of high self-opinion. This is shown to be when the agent's cost of experimentation is sufficiently *high*. To see this, let us first look at the supervisor's incentives to be honest in the region of priors $\beta_1^{NI} \leq \beta_1 < \beta_2^{NI}$. Here the agent experiments once even when discouraged. At most the agent's belief can fall down to β_0^{NI} after a negative message. The supervisor is then willing to discourage the agent with a negative message only if he can ensure that even after discouragement the agent does not reduce her effort significantly. In the absence of further supervision, he can only expect a higher expected probability of success if she exerts a high enough effort in implementation.

A supervisor who has observed a low potential idea expects the project to be successful with probability $(\beta_2^\ell(q + (1 - q)k))^2$ from being honest. After receiving a message $m_1 = \ell$, the agent correctly believes her current idea has a low potential to succeed and experiments once again but does not seek supervision because the supervisor is expected to babble. In this case, the agent then implements the next idea with effort $e = \beta_2^\ell(q + (1 - q)k)$. On the other hand, if such a supervisor deviates from honesty and announces $m_1 = \hbar$, then he expects the probability of success to be $\beta_2^\ell k$. The agent incorrectly believes that her idea had a high potential to succeed and exerts effort of 1 in implementing a low potential idea. For such a conjectured strategy to be an equilibrium, we must have that

$$\begin{aligned} & (\beta_2^\ell)^2(q + (1 - q)k)^2 \geq \beta_2^\ell k \\ \implies \beta_1 & \geq \frac{k}{qk + (1 - q)(q + (1 - q)k)^2} := \beta^{\text{truth}} \end{aligned}$$

Thus, the supervisor requires agent's belief to be sufficiently high even after discouragement, which

in turn requires the prior to be large enough. This ensures that the agent exerts a higher effort in implementing her idea of unknown potential. We call this truth-telling threshold on prior β^{truth} .

The truth-telling threshold β^{truth} is a conditional threshold. It identifies how high the prior should be such that the supervisor has an incentive to reveal the truth about the agent's negative outcome *if the agent experiments again without supervision following the negative message*. The supervisor does not directly care about the agent's cost of experimentation in so far as she attempts to experiment again with an idea. So β^{truth} does not depend on c .

Now all we need to do is identify whether the range of priors we are considering delivers honesty by the supervisor, that is we are interested in if $\beta^{\text{truth}} < \beta_2^{NI}$. Specifically, if $\beta^{\text{truth}} \leq \beta_1^{NI}$ then truth-telling is an equilibrium for the full range of beliefs above β_1^{NI} and up to β_2^{NI} . If this condition is satisfied, the supervisor has an incentive to be honest because the prior is sufficiently high given the parameters. As outlined above, β^{truth} does not depend on the cost of experimentation c while β_1^{NI} does. The one free parameter can be used to determine if truth-telling is an equilibrium. The condition $\beta^{\text{truth}} \leq \beta_1^{NI}$ can then be rearranged to

$$c \geq \frac{\kappa k - (\kappa k)^2}{2} \quad \text{where } \kappa \equiv \frac{k}{(q + (1 - q)k)^2} < 1.$$

Intuitively, a lower bound on the cost of experimentation ensures that the agent's no information thresholds β_0^{NI} and β_1^{NI} are high enough. Thus, when the agent decides to experiment and consult the supervisor her belief in her ability is already high. The supervisor can then be content with revealing the truth about low potential ideas to the agent. Discouragement does not lead to quitting with low effort; the agent still experiments once more and does so by exerting a sufficiently high effort. While the conditional truth-telling threshold β^{truth} is not a function of the cost of experimentation c , whether truth-telling is an equilibrium depends on it. An increase in the cost of experimentation raises the threshold β_0^{NI} (increasing the region of babbling) but has no effect on β^{truth} , making it easier to satisfy the condition $\beta^{\text{truth}} \leq \beta_1^{NI}$ and ensuring truth-telling above β_1^{NI} .

We are now only left with determining why if the supervisor is honest in the range of beliefs $\beta_1^{NI} \leq \beta_1 < \beta_2^{NI}$, then he should be honest in the range of beliefs above β_2^{NI} . For expositional convenience start now with the range of beliefs $\beta_2^{NI} \leq \beta_1 < \beta_3^{NI}$ when it is an equilibrium for the supervisor to be honest in the next lower range of beliefs. Consider whether a conjectured strategy of honesty is an equilibrium for the supervisor. A supervisor who observes a low potential idea can induce another two rounds of experimentation by being honest at this stage, one with supervision and one without. If, however, he deviates he induces the agent to exert maximal effort in a low potential

task. Under assumption (A), the payoff from being honest are strictly higher than that from deviating as it is evaluated relative to his private updated belief β_2^ℓ . The same line of reasoning can then be extended to any belief above β_3^{NI} as well so that the supervisor always prefers honestly discouraging the agent and getting her to experiment more often than making her implement a low potential idea.

What happens when $c < \frac{\kappa k - (\kappa k)^2}{2}$? The following corollary identifies the honest equilibrium.

Corollary 1 *When $c < \frac{\kappa k - (\kappa k)^2}{2}$, $\beta_j^{NI} \leq \beta^{truth} < \beta_{j+1}^{NI}$ exists such that for all $t > 1$ for $j \geq 1$*

1. *truth-telling is an equilibrium strategy for the supervisor for $\beta_t \geq \beta^{truth}$, and*
2. *babbling is an equilibrium strategy for the supervisor for $\beta_t < \beta^{truth}$.*

The agent's equilibrium strategy is given by

$$I_t^* = \begin{cases} 0 & \text{if } m_{t-1} = \ell \text{ and } \beta_t \geq \beta^{truth}, \text{ or } \beta_{j-1}^{NI} \leq \beta_t < \beta_j^{NI}, \\ 1 & \text{otherwise.} \end{cases}$$


The agent's optimal effort is given by

$$e^* = \begin{cases} 1 & \text{if } m_{t-1} = \hat{h}, \\ \beta_t(q + (1 - q)k) & \text{otherwise.} \end{cases}$$

In this case, $\beta^{truth} > \beta_1^{NI}$ and can lie between any β_j^{NI} and β_{j+1}^{NI} . We can then again construct an honest equilibrium above β^{truth} and a babbling one below. That all of these beliefs are above β_0^{NI} ensures that the agent experiments once more when a low potential idea is revealed to her in the presence of future babbling and makes such a strategy an equilibrium. The two cases discussed here are depicted in Figure 5.

It is worth emphasizing at this stage the key intuition driving the results in Propositions 2 and 3. What action the agent chooses depends on whether she thinks she is capable of drawing a better idea, and the expected strategy of the supervisor. If the agent has produced a low potential idea, the supervisor needs to incorporate the downwards effect that his negative message has on the belief about her ability. A lower belief discourages the agent at two levels. First is the discouragement to experiment, i.e., stopping experimentation too early. Second is the discouragement to implement, i.e., exerting low effort in implementing the idea. The low self-opinion belief phase arises when the first effect dominates where the concern of the supervisor is the agent abandoning experimentation.

1. Honest equilibria when $c \geq \frac{\kappa k - (\kappa k)^2}{2}$



2. Honest equilibria when $c < \frac{\kappa k - (\kappa k)^2}{2}$




Figure 5: Honest equilibria for different c ranges

The high self-opinion belief phase arises when the second effect dominates where the concern of the supervisor is the agent exerting a low effort after experimentation.

We conclude this section by presenting an important corollary and our second main result.

Corollary 2 *The expected performance of the agent is better under a higher self-opinion.*

To see this, first note that the supervisor induces a weakly higher number of rounds of experimentation under a prior $\beta' > \beta$. If $\beta_j^{NI} \leq \beta < \beta_{j+1}^{NI}$, then either $\beta_j^{NI} \leq \beta < \beta' < \beta_{j+1}^{NI}$ or $\beta' > \beta_{j+1}^{NI}$. In the former case, the agent experiments an equal number of rounds under the two beliefs. However, in the latter case, the agent experiments more often under belief β' than under β . The reason is that it is easier to support the mutual expectation of honesty and repeated experimentation under a higher belief so that the agent experiments weakly more often under β' .

However, this has consequences on the agent's overall performance. Honest feedback by the supervisor allows the agent to match her effort more closely to the actual potential maximizing the probability of success. If the agent abandons seeking supervision (and experiments one final round) in the k th round under belief β , then she should still be seeking honest supervision in the round k under belief β' . While the agent with belief β exerts an inefficiently low amount of effort in a high potential idea in round k , an agent with β' will exert the efficient level of effort of 1. An inefficient level of effort reduces the probability of success in a high potential idea.

Finally, if the idea in round k is a low potential one, then an agent with lower belief exerts an inefficiently high level of effort in its implementation while the agent with a higher belief experiments again. Therefore, there is a magnifying effect of a higher belief that results from the combined effect of better experimentation and better implementation. Conditional on being high-ability, an a priori

better agent who has a higher belief in her ability does better in expectation.

It is also worth noting that in this context an a priori better agent (who has a higher self-opinion) will face more criticism from the supervisor for the same reason. An agent with a higher belief in her ability receives discouraging messages more often conditional on producing the same number of low potential ideas. However, the agent’s incentive to experiment more often arises precisely out of the supervisor offering honest criticism. In equilibrium, an agent with a higher belief expects to receive honest feedback more often and is therefore willing to experiment more often. In return, the supervisor expecting more experimentation offers more honest feedback to the agent. When the agent’s belief is lower, he fears to discourage the agent with negative messages. In this sense, an agent with a higher belief is more receptive to criticism, and that increases her chances of being successful.

4.3 Welfare analysis

The previous result (Corollary 2) only talks about the benefit of a higher self-opinion. However, the agent also pays a higher cost under a higher self-opinion owing to the aforementioned magnifying effect. This particularly hurts a low-ability agent who only pays a higher costs of experimentation and/or implementation under a higher belief.

The first part of this section shows that the above is not a concern even when evaluating the agent’s welfare under a higher self-opinion. We show, through a series of lemmas below that the ex-ante expected utility of the agent is always higher under a higher belief.¹⁴ The reason is that under a higher belief the agent places a greater ex-ante weight on being high-ability and believes that she is less likely to find herself in the worst situation.

The second part of the section then analyzes if holding an incorrect higher belief could also be welfare improving. Surprisingly, we show that this is possible. The reason is the discontinuous change in the supervisor’s feedback strategy as it switches from babbling to honesty.¹⁵

4.3.1 Welfare effect of a correct increase in self-opinion

Lemma 4 *Any increase in the prior from β to β' within the region of beliefs $0 < \beta < \beta' < \beta_0^{NI}$, $\beta_0^{NI} \leq \beta < \beta' < \beta_1^{NI}$, and $\beta_j^{NI} \leq \beta < \beta' < \beta_{j+1}^{NI}$ for $j > 1$ is welfare improving for the agent.*

This lemma relates to increasing the beliefs of the agent in such a way that only the cost of exerting effort increases in the eventuality that the project is implemented with a low potential idea or after

¹⁴The supervisor is always better off with a higher self-opinion agent because in expectation such an agent performs better. At the same time, the supervisor doesn’t have to bear any costs.

¹⁵We prove all the statements here assuming that $c \geq \frac{\kappa k - (\kappa k)^2}{2}$ or that the truth-telling threshold $\beta^{\text{truth}} \leq \beta_1^{NI}$. However, this is not required as the proofs go through with a higher β^{truth} as well.

not seeking supervision. In such a situation, welfare may increase on account of better implementation (because of higher effort) but may reduce on account higher costs of implementing.

Lemma 5 *An increase in the prior from the region $\beta_0^{FI} \leq \beta < \beta_0^{NI}$ to the region $\beta_0^{NI} \leq \beta' < \beta_1^{NI}$ is welfare improving for the agent.*

When the belief increases in such a manner, the agent is expected to conduct a costly round of experimentation which she did not earlier. Moreover, she is not expected to receive any feedback in this round. At the same time, her optimal effort choice increases unambiguously which is both more costly and more beneficial in expectation. From Lemma 4, we know that increasing the effort is always welfare improving when the belief increases. In addition, the increase in belief also makes it worthwhile to conduct experimentation without supervision from Lemma 1. This leads to an overall increase in welfare.

Lemma 6 *Let $2c < q(1 - (q + (1 - q)k)^2)$. An increase in the prior from $\beta = \beta_{j+1}^{NI} - \epsilon$ to $\beta' = \beta_{j+1}^{NI}$ is welfare improving for the agent.*

Finally, this lemma establishes that just pushing up the belief from an arbitrary region $\beta_j^{NI} \leq \beta < \beta_{j+1}^{NI}$ to the next region $\beta_{j+1}^{NI} \leq \beta' < \beta_{j+2}^{NI}$ is welfare improving. In doing so, the agent is expected to pay not only an additional cost of experimentation c but also that of some minimal increase in effort cost in the event of implementing without supervision.

Proposition 4 *Let $2c < q(1 - (q + (1 - q)k)^2)$. An increase in the prior from β to β' is welfare improving for the agent.*

The above proposition combines the information from the three lemmas and concludes that any increase in prior is welfare improving. This highlights the importance of agent's self-opinion – the agent's confidence in her ability is critical for the overall success of the project.

4.3.2 Welfare effect of overconfidence

Still more interesting is to explain the effect of overconfidence in our environment. To introduce the notion of overconfidence, consider the following. Let the agent and the supervisor hold a common prior belief β about the agent's ability when the true belief is b .

Definition 1 *The agent and the supervisor are overconfident about the agent's ability if $\beta > b$.*

Under the above definition of overconfidence, we prove the following proposition:

Proposition 5 *Overconfidence is sometimes, but not always, welfare improving.*

To understand the intuition, consider the welfare of the agent when the correct belief is $b = \beta_1^{NI} - \epsilon$ but the common prior is β_1^{NI} . In such a situation, her overconfidence will drive her to experiment once with a round of honest feedback by the supervisor (and then potentially once more without any feedback). This would not have been possible under the true belief wherein she would have simply experimented without any feedback. However, the discontinuous benefit that arises from the change in supervisor's feedback strategy at a higher belief (i.e. receiving honest feedback) outweighs the additional cost that the agent pays for an additional round of experimentation.

In fact, she is able to reduce her inefficient cost of implementation when the supervisor honestly reveals that her idea was a low potential one under the overconfident belief. To see this note that under the true belief she would exert $(\beta_1^{NI} - \epsilon)(q + (1 - q)k)$. Whereas under the overconfident belief she would exert $\beta_0^{NI}(q + (1 - q)k)$. Thus, overconfidence (and holding an incorrect self-opinion) can be welfare improving.

However, the above argument relies on the discontinuous change in behavior of the supervisor at the threshold. It then follows that when the supervisor's behavior does not change, there might not be a benefit of being overconfident. To illustrate this, we show that overconfidence is welfare reducing when the common prior is β_0^{NI} but the true belief is any $b < \beta_0^{NI}$. In such a situation, holding the incorrect belief only adds to an added cost of experimentation and implementation without any corresponding benefit. Contrasting this with Lemma 6, it is immediate to see that overconfidence is different from a correct increase in belief.

5 Extensions

5.1 Benevolent supervisor and time-constrained players

We start out by discussing what happens when the supervisor also bears the cost of experimentation and implementation. In some situations, it is possible that a benevolent supervisor partially internalizes the costs borne by the agent. Such internalization may arise from the expert's (i.e. the supervisor's) prior experience from when he as an apprentice (agent), or simply because he works on the project with the agent.

For the two players $i \in \{A, S\}$, agent (A) and supervisor (S), let the cost of experimentation be c_i and the cost of implementation be $\frac{\phi_i \epsilon^2}{2}$. The difference between these costs for the two players captures any preference conflict between them. In so far as $c_S < c_A$ and $\phi_S < \phi_A = 1$, the preference conflict

persists. For a given $(c_S, \phi_S) > 0$, there will be a “full information” threshold for the supervisor as well. Call this threshold β_{S0}^{FI} . This reflects the preferences of the supervisor and determines what are the maximum number of rounds the supervisor desires the agent to experiment (or the belief threshold equivalently) with full information about the potential of the ideas.

In the limiting case of $c_S = \phi_S = 0$ studied in the main text, this threshold did not exist – the supervisor wanted the agent to continue experimenting with complete information until she ended up with a high potential idea. However, when $c_S < c_A$ and $\phi_S < \phi_A$, we have $\beta_{S0}^{FI} < \beta_{A0}^{FI}$ so that the supervisor would still like the agent to experiment more than she would like. In this case, all our results from the main text go through as the fear of discouragement and the agent abandoning experimentation still persists.

One possible interpretation of such a situation are time-constrained players. To keep things simple, let $\phi_S = \phi_A = 1$ so that the supervisor fully internalizes the time cost of implementing to the agent. Now let c_S denote the time cost that the supervisor pays for providing feedback to the agent. This could happen when the supervisor has some alternate tasks to perform or requires time to understand the true potential of the agent’s ideas. The following proposition follows from our discussion.

Proposition 6 *Let $\phi_S = \phi_A = 1$.*

1. *If $c_S < c_A$ then Propositions 1, 2 and 3 capture the optimal strategies of the agent and the supervisor.*
2. *If $c_S \geq c_A$ then the supervisor offers honest feedback until he reaches the belief β_{S0}^{FI} and the agent experiments with ideas till that point absent a high potential idea.*

The intuition is as follows. When the supervisor is time-constrained, he cares both about success, and about costly supervision from the agent experimenting in pursuit of success. In turn, this eliminates the fear of discouragement. Notably, now it is more costly for the supervisor to keep offering feedback beyond a point over letting the agent implement a low potential idea. We can then get honest equilibria for some additional ranges of beliefs. Thus, a more time-constrained supervisor can potentially offer more honest feedback. The next corollary identifies the condition that makes this possible.

Corollary 3 *Let $\phi_S = \phi_A = 1$. If $c_S \geq c_A$ such that $\beta_{S0}^{FI} < \beta_1^{NI}$ then the region of beliefs where honest equilibria exist is larger in the case of $c_S \geq c_A$ than $c_S < c_A$.*

Observe that in the case of $c_S < c_A$ honest equilibria exist in the region of beliefs from β_1^{NI} upwards (depending on c_A). But from the above proposition, honest equilibria in the case of $c_S \geq c_A$

exist starting from β_{S0}^{FI} . Thus, the latter case provides the possibility of more honesty if $\beta_{S0}^{FI} < \beta_1^{NI}$. However, since there is no closed form solution of β_{S0}^{FI} , it is not straightforward to translate this into a condition with only the costs.

Finally, note that if the supervisor does not internalize the cost of exerting effort, there is no benefit (in terms of more honest feedback) of even partially internalizing the costs of experimentation.

Proposition 7 *If $\phi_S = 0$, then the equilibrium strategies are given by Propositions 1, 2 and 3.*

To understand the intuition, let $c_S = c_A$ and consider whether honesty is an equilibrium strategy for $\beta_{A0}^{FI} \leq \beta < \beta_{A1}^{FI}$ (after all, if the supervisor internalizes the full cost of experimentation then the belief thresholds should match). At this belief, if the supervisor is expected to be honest, then following a negative message the agent abandons experimentation and exerts a low effort level on the idea. If instead, she receives a positive message, she exerts 1 on her idea. Now, for a supervisor who has seen a low potential idea and does not internalize the cost of implementation, there is a strictly positive deviation to giving a positive message. This breaks down the honest equilibrium (and the existence of β_{S0}^{FI}).¹⁶

The issue arises here because the supervisor wants the agent to exert the maximal effort independent of the potential of the idea produced. The supervisor fears discouragement leading to lower effort in implementation which precludes honesty.

5.2 Perfect recall of previous ideas

Here we describe what happens if the agent and the supervisor have perfect recall of all the previous ideas. If that is the case, then in each round of experimentation the supervisor can potentially make announcements about each of the previous ideas. Given our attention to pure strategies, there are two kinds of honest and informative strategies that a supervisor may employ: immediate honesty and delayed honesty.

In the immediately honest strategy, the supervisor reveals to the agent the outcome of her experimentation immediately after she experiments. This is implicitly what we assumed all throughout Section 4. In a strategy of delayed honesty, the supervisor provides uninformative messages for certain rounds and then reveals honestly some or all the previous outcomes. Observe that a variety of delayed honesty strategies are possible – the supervisor may babble for any arbitrary number of rounds and then provide information for any arbitrary number of those rounds, and this may change over time. If

¹⁶It is possible to derive a belief threshold above which the supervisor is expected to be honest in equilibrium for a generic ϕ_S and given c_S and c_A . This is necessarily different from β_{S0}^{FI} because that is contingent on the equilibrium best response of the agent to the supervisor's strategy.

the supervisor reveals to the agent the $\tau' \leq \tau$ outcomes of her experimentation after τ rounds starting from a prior β_t the agent's updated posterior in round $t + \tau$ is

$$\beta_{t+\tau}^\ell = \frac{(1-q)^{\tau'} \beta_t}{1 - q\beta_t \sum_{s=0}^{\tau'-1} (1-q)^s} \text{ if } m_t = \ell \text{ for all } \tau' \text{ ideas, and} \quad (3)$$

$$\beta_{t+\tau}^h = 1 \text{ otherwise.} \quad (4)$$

The case of $\tau = \tau' = 1$ corresponds to immediate honesty where the agent expects the supervisor to reveal the outcome of the experimentation immediately after each round of experimentation. All other cases fall under delayed honesty.

In case the supervisor is expected to babble, the agent's posterior belief is the same as her prior belief. We will assume that when the supervisor is expected to lie about an idea the agent does not consult the supervisor regarding that idea. This rules out the possibility of the supervisor privately learning and not revealing to the agent the outcome, and the arising deviations.¹⁷

Note first that the result of Proposition 1 remains unaltered. If the agent does not want to experiment with an immediately honest strategy, she does not want to experiment with a delayed honesty strategy. By experimenting when the supervisor is expected to reveal the outcomes after a delay, the agent only bears a higher cost of experimentation to receive feedback when she is almost convinced that she cannot produce a high potential idea. Thus, implementing the outside option is the best response of the agent, and all strategies of information revelation are an equilibrium.

Corollary 4 *Under perfect recall of ideas, for any belief $\beta_0^{FI} \leq \beta < \beta_1^{NI}$, babbling is the unique equilibrium strategy.*

The result of babbling being a unique equilibrium in the region of beliefs $\beta_0^{FI} \leq \beta < \beta_1^{NI}$ even under perfect recall follows almost directly from Proposition 2. To illustrate this point, start again with a prior belief $\beta_0^{FI} \leq \beta_1 < \beta_1^{FI}$. In the absence of commitment, a supervisor who observes only low potential ideas from all the experimentation rounds (after delaying) is tempted to deviate and call any arbitrary idea a high potential one. This is for the same reason as before – when such a message is believed, the agent exerts maximal effort on such an idea assuming it is a high potential one. The supervisor gains from such a deviation because he increases the effort of the agent on a low potential idea in the absence of more experimentation. As a result, babbling is the unique equilibrium and

¹⁷A formal definition of strategies in this case is complicated. But it is easy to describe what a strategy for the two players are in words. A strategy for the supervisor when the agent consults him in round t is a mapping from all the ideas she observes to the set of messages, one for each round of experimentation. A strategy for the agent in round t is a mapping from the observed messages to a decision to experiment again or implement. If she decides to implement, she must also decide which idea to implement given the message history.




Figure 6: Terminal belief possibilities in potential delayed equilibria

the agent best responds by implementing the low potential outside option idea. The same reasoning can then be extended to all the beliefs which when updated with a negative message lead to the agent abandoning experimentation (as the supervisor is going to babble in the following round). This happens all the way up to the belief β_1^{NI} as before.

For beliefs above β_1^{NI} , we have already identified the condition for *immediate honesty* to arise in Proposition 3. It is, however, possible to have other equilibria with some delayed honesty. We identify here a critical feature of such equilibria (if they exist) that allows us to compare it with the immediately honest equilibrium.

Observation 1 *In a delayed equilibrium, the supervisor can only induce as many rounds of experimentation as the ones for which he provides honest feedback eventually.*

The above observation merely states that if the supervisor never provides feedback on some rounds of experimentation that the agent performs, then the agent has no incentive to experiment. Since the agent never consults the supervisor for rounds in which he is expected to babble, there is no benefit to the agent from experimenting these extra rounds. This allows us to focus attention on those strategies in which the outcome of all the rounds of experimentation is eventually revealed.

Proposition 8 *The number of rounds of experimentation that an equilibrium strategy of delayed honesty induces can be no more than that induced by the equilibrium immediate honesty strategy.*

What matters when evaluating the supervisor's incentive to be honest at the time of final revelation is the belief from truthfully announcing that all the ideas produced are low potential. Say that the belief after such a revelation at round τ is β_τ^ℓ . This belief can be in one of the following three ranges: $\beta_\tau^\ell \geq \beta_1^{NI}$, $\beta_0^{NI} \leq \beta_\tau^\ell < \beta_1^{NI}$ or $\beta_\tau^\ell < \beta_0^{NI}$ (See Figure 6).

Observe that a terminal belief in the first and second range can also be attained by an immediately honest strategy, which is also an equilibrium. For any prior β_1 , for the agent to experiment more rounds than what she does under immediately honest strategy her terminal belief after all the revelations should fall in the third case, i.e. $\beta_\tau^\ell < \beta_0^{NI}$. However, we argue that such a strategy cannot be an

equilibrium. This is for the same reason as before – a supervisor who has only observed low potential ideas will prefer to deviate and claim any one of the ideas to be of high potential than inducing the agent to stop experimenting with a lower belief where the supervisor only babbles. Thus, equilibrium experimentation possibilities under perfect recall can be no more than those under limited recall.

5.3 Alternate interpretations

Our model more generally speaks to the following type of settings. An informed sender of information (supervisor) communicates with a less informed receiver of information (agent) who needs to take a costly action dynamically. Consider, for instance, an entrepreneur who works on a project experimenting with ideas, *privately observing their potential*, and implementing one of them. However, she relies on the finances of a venture capitalist (VC) who pays for such experimentation and implementation. While the entrepreneur would prefer to continue experimenting until she receives a high potential idea, the VC would like to cut funding for experimentation when he is sufficiently pessimistic.

In such a setting, the entrepreneur is the supervisor, while the VC is the agent.¹⁸ Costs c and $e^2/2$ are the money promised by the agent to the supervisor for experimenting with and implementing ideas. Let $\alpha \in \{0, q\}$ be the state of the project which is determined ex-ante and remains persistent but potentially unknown to both the parties. $\theta \in \{\ell, h\}$ denotes the potential of the idea produced by the entrepreneur. The VC decides in each period, whether to fund experimentation for one extra round or force the entrepreneur to implement the last idea.

We then provide answers to the following questions: When can the entrepreneur credibly release information? How many chances of experimentation can the entrepreneur extract from the VC with her revelation strategy? Notably, our inefficiency result shows that even though the VC would like to continue financing the entrepreneur’s experimentation and the entrepreneur would like to continue experimenting, she calls off the project too early. However, there are benefits to be had from the VC both correctly and incorrectly believing that the project is good.

6 Conclusion

In this paper, we showed how an employee responds to criticism influences whether she receives feedback or not. Supervisors may not provide honest feedback to employees who do not believe in their ability. In turn, this hurts their performance and potentially their future careers. Moreover, it

¹⁸Which player is the agent and which one is the supervisor is *not* determined by who is experimenting and implementing, but by who holds the information and who pays for the action.

also hurts organizations as the supervisors provide inefficiently low levels of honest feedback. In this sense, organizations should seek to hire employees that *believe* in their ability to succeed. In fact, our model shows that overconfidence can sometimes be welfare-improving.

Our results are based on a model of feedback provision in an agent-supervisor environment. The agent experiments with ideas to try to solve a problem at hand and a supervisor offers feedback on whether her ideas have the potential to be successful. We showed the results for when the supervisor has no commitment power and uses cheap talk messages to communicate with the agent. We identified the region of beliefs for which the supervisor could only uniquely babble in equilibrium leading to inefficiency in the relationship. Driven by the fear of discouraging the agent to the point of abandonment of experimentation, the supervisor is not able to offer any credible information to the agent. We then showed if there are possible equilibria in which the supervisor can honestly communicate his information to the agent. A necessary and sufficient condition for honesty above the babbling threshold was found to be the costs of experimentation being sufficiently high.

However, our analysis focused only on pure strategy equilibria. The problem involving mixed strategies is a complicated one that requires determining how the agent responds to the current message when, in the future, there can be more mixing. Our work shows the further scope of looking at mixed communication strategies in such dynamic environments in the absence of commitment. One may also think of introducing new complications in the model such as those involving different priors of the agent and the supervisor.

References

- Ali, O. (2017). Disclosure in multistage projects. *Working paper*.
- Aumann, R. J., & Hart, S. (2003, November). Long cheap talk. *Econometrica*, *71*(6), 1619–1660.
- Bandura, A. (1977). Self-efficacy: toward a unifying theory of behavioral change. *Psychological review*, *84*(2), 191.
- Bénabou, R., & Tirole, J. (2002). Self-confidence and personal motivation. *The Quarterly Journal of Economics*, *117*(3), 871–915.
- Bizzotto, J., Rüdiger, J., & Vigier, A. (2018). Dynamic persuasion with outside information. *Working paper*.
- Boleslavsky, R., & Lewis, T. R. (2016). Evolving influence: Mitigating extreme conflicts of interest in advisory relationships. *Games and Economic Behavior*, *98*, 110–134.
- Che, Y.-K., & Horner, J. (2015). Optimal design for social learning. *Working Paper*.

- Ely, J. C. (2017). Beeps. *The American Economic Review*, *107*(1), 31–53.
- Ely, J. C., & Szydlowski, M. (2017). Moving the goalposts. *Working paper*.
- Forges, F., & Koessler, F. (2008). Long persuasion games. *Journal of Economic Theory*, *143*(1), 1–35.
- Fuchs, W. (2007). Contracting with repeated moral hazard and private evaluations. *The American Economic Review*, *97*(4), 1432–1448.
- Golosov, M., Skreta, V., Tsyvinski, A., & Wilson, A. (2014). Dynamic strategic information transmission. *Journal of Economic Theory*, *151*, 304–341.
- Goltsman, M., Hörner, J., Pavlov, G., & Squintani, F. (2009). Mediation, arbitration and negotiation. *Journal of Economic Theory*, *144*(4), 1397–1420.
- Henry, E., & Ottaviani, M. (2019, March). Research and the approval process: The organization of persuasion. *American Economic Review*, *109*(3), 911–55. Retrieved from <http://www.aeaweb.org/articles?id=10.1257/aer.20171919> doi: 10.1257/aer.20171919
- Hirshleifer, D., Low, A., & Teoh, S. H. (2012). Are overconfident ceos better innovators? *The Journal of Finance*, *67*(4), 1457–1498.
- Honryo, T. (2018). Dynamic persuasion. *Journal of Economic Theory*, *178*, 36 - 58. Retrieved from <http://www.sciencedirect.com/science/article/pii/S002205311830485X> doi: <https://doi.org/10.1016/j.jet.2018.08.007>
- Hörner, J., & Skrzypacz, A. (2016). Learning, experimentation and information design. *Working Paper*.
- Kamenica, E., & Gentzkow, M. (2011). Bayesian persuasion. *The American Economic Review*, *101*(6), 2590–2615.
- Koellinger, P., Minniti, M., & Schade, C. (2007). i think i can, i think i can: Overconfidence and entrepreneurial behavior. *Journal of economic psychology*, *28*(4), 502–527.
- Kremer, I., Mansour, Y., & Perry, M. (2014). Implementing the “wisdom of the crowd”. *Journal of Political Economy*, *122*(5), 988–1012.
- Krishna, V., & Morgan, J. (2004). The art of conversation: eliciting information from experts through multi-stage communication. *Journal of Economic theory*, *117*(2), 147–179.
- Lizzeri, A., Meyer, M. A., & Persico, N. (2002). *The incentive effects of interim performance evaluations*. University of Pennsylvania, Center for Analytic Research in Economics and the Social Sciences.
- Morris, S. (2001). Political correctness. *Journal of Political Economy*, *109*(2), 231–265. Retrieved from <https://doi.org/10.1086/319554> doi: 10.1086/319554

- Orlov, D. (2013). Optimal design of internal disclosure. *Working paper*.
- Orlov, D., Skrzypacz, A., & Zryumov, P. (2018). Persuading the principal to wait. *Working paper*.
- Renault, J., Solan, E., & Vieille, N. (2013). Dynamic sender–receiver games. *Journal of Economic Theory*, *148*(2), 502–534.
- Renault, J., Solan, E., & Vieille, N. (2017). Optimal dynamic information provision. *Games and Economic Behavior*, *104*, 329–349.
- Smolin, A. (2017). Dynamic evaluation design. *Working Paper*.

Appendices

Appendix A: Proofs from main text

We present general proofs in mixed strategies, wherever we can. The first section provides some new mathematical notation for this purpose.

Mathematical notation for mixed strategies

We focus attention on limited recall of previous ideas so that when the agent experiments one more round, she does not recall the previous ideas she has worked on. As a result, the supervisor does not need to make back dated messages about all the previous ideas. A strategy for the agent ρ_t in round t is a mapping from the last observed message to a possible mixed decision to continue experimenting with ideas or implementing the last one . We let

$$\rho_t^{m_{t-1}} = \Pr(I_t = 1 \mid m_{t-1})$$

be the probability that the agent decides to implement the project following the last message.

Similarly, when the supervisor is called upon, a strategy for the supervisor σ_t in round t is a mapping from the last idea to a possible mixed message about its potential. We let

$$\sigma_t^{\theta_t} = \Pr(m_t = \theta_t \mid \theta_t)$$

be the probability of the supervisor being honest about the potential of the observed idea. Depending on the expected strategy of the supervisor, the agent conditions her action only on the last message received.

Let the sequence $\hat{\sigma} = \{\hat{\sigma}_t^h, \hat{\sigma}_t^\ell\}_{t=1}^T$ denote the conjectured strategy of the supervisor, and let $\hat{\rho} = \{\hat{\rho}_t^h, \hat{\rho}_t^\ell\}_{t=1}^T$ denote the conjectured strategy of the agent. Given the conjectured strategy of the supervisor, the agent updates beliefs about the two unknowns – her ability and the potential of her previous ideas. The belief about her ability is β_t . Let the belief about whether her idea was as announced by the supervisor be denoted by λ_t . Observe that:

1. the public history h_t^A at the beginning of round t can be summarized by the current public belief β_t about the ability of the agent and by the belief about the true potential of the last idea produced λ_t , while

2. the private history of the supervisor h_t^S at the beginning of round t can be summarized by the current private belief β_t about the ability of the agent.¹⁹

We can now informally describe the notion of equilibrium. We say that a pair of sequences of conjectured strategies σ and ρ constitute an equilibrium if (1) they are both the best responses to each other given the beliefs β_t and λ_t for each t , and (2) the beliefs β_t and λ_t are consistent with what the players are conjectured to do, i.e. σ and ρ . Strategies expressed in the text without a hat constitute an equilibrium.

When both the messages are expected in equilibrium, either one of the messages will lead to a higher and the other to a lower β_t , or β_t remains the same with both the messages. We will call the former *informative* strategy and the latter *babbling* (or lying) strategy. The supervisor is expected to babble in equilibrium in round $t - 1$ if $\hat{\sigma}_{t-1}^h = 1 - \hat{\sigma}_{t-1}^\ell$, i.e. when the probability with which the supervisor is expected to reveal a true high potential idea is the same as the probability with which the supervisor incorrectly calls a low potential idea a high one. Thus, the agent is equally likely to get a positive or a negative message, and in turn does not learn from the messages. When the supervisor is expected to be informative, we will assume without loss of generality that he does so by increasing the posterior after a positive message of $m_{t-1} = h$ (and the posterior beliefs fall after a negative message $m_{t-1} = \ell$). So, we assume that $\hat{\sigma}_{t-1}^h > 1 - \hat{\sigma}_{t-1}^\ell$ for informativeness.

We will restrict attention here to informative strategies in which $\sigma^h = 1$, i.e. the supervisor always truthfully announces that the project has a high potential to succeed when he sees so. The supervisor cannot credibly commit to lying when $\theta_t = h$. In any informative strategy, a positive message $m_t = h$ should increase the posterior belief β_{t+1} of the agent. When the supervisor sees $\theta_t = h$, he has no incentive to discourage the agent. If discouragement leads to another round of experimentation, then the supervisor faces the risk of abandoning the current high potential idea and never getting a new one. Alternately, if discouragement leads to implementation then she will do so with a lower effort. In neither case a supervisor who has observed a high potential idea is better off discouraging the agent. Going forward, we assume $\sigma_t^h = 1$, and with some replace σ_t^ℓ with σ_t . Then the posterior beliefs about

¹⁹Note that we are currently not making any notational distinction between the private and the public beliefs about ability. This is to keep things simple. The two will coincide as long as the supervisor is honest. When the supervisor is not honest, the beliefs diverge only when the agent best responds to a dishonest message by experimenting again. This plays a role only in checking for deviations when constructing other informative equilibria.

ability is

$$\beta_t^\ell = \frac{(1-q)\beta_{t-1}}{1-q\beta_{t-1}} \quad (5)$$

$$\beta_t^{\hat{h}} = \frac{(1-\hat{\sigma}_{t-1}(1-q))\beta_{t-1}}{1-\hat{\sigma}_{t-1}(1-q\beta_{t-1})} \quad (6)$$

where $\beta_t^{m_{t-1}} = \Pr(\alpha = q|m_{t-1})$ is the posterior belief of the agent about her ability after receiving message m_{t-1} given the conjecture $\hat{\sigma}_{t-1}$. And

$$\lambda_t^\ell = 1 \quad (7)$$

$$\lambda_t^{\hat{h}} = \frac{q}{q + (1 - \hat{\sigma}_{t-1})(1 - q)} \quad (8)$$

where $\lambda_t^{m_{t-1}} = \Pr(\theta_{t-1} = m_{t-1}|m_{t-1})$ is the belief about whether the supervisor's message m_{t-1} matches the true potential of the idea given the conjectured $\hat{\sigma}_{t-1}$.

Thus, the value of a negative message under any informative strategy is the same as in a truth-telling strategy. When an agent receives $m_t = \ell$ then she can be sure that $\theta_t = \ell$ and she revises her belief about her ability downwards to the maximum extent. Under this condition, the agent must decide what to do following a message of $m_t = \hat{h}$ since a positive message cannot be trusted.

Proof of Lemma 2

Proof.

Part 1: Existence of β_0^{FI}

For a given set of parameters, there is no straightforward closed form solution to the equation in condition C2. We therefore need to establish the existence of belief threshold(s). First, it can be verified that both the LHS and RHS of condition C2 are monotonically increasing and convex in β . We have

$$\begin{aligned} \frac{\partial \text{LHS}}{\partial \beta} &= \frac{q}{2} + \frac{(k\beta')^2}{2} \left(\frac{2}{\beta} - q \right) > 0 \\ \frac{\partial^2 \text{LHS}}{\partial \beta^2} &= \frac{k^2(1-q)^2}{(1-\beta q)^3} > 0 \end{aligned}$$

and

$$\begin{aligned}\frac{\partial \text{RHS}}{\partial \beta} &= k^2 \beta \geq 0 \\ \frac{\partial^2 \text{RHS}}{\partial \beta^2} &= k^2 > 0.\end{aligned}$$

Second, we show that if $2c < q(1 - k^2)$ then the threshold belief β_0^{FI} is unique. Consider the range of beliefs $0 \leq \beta \leq 1$. Since $c > 0$ and LHS at $\beta = 0$ is zero, RHS cuts the LHS from above at least once. Now, under the assumption $2c < q(1 - k^2)$, it can be verified that RHS at $\beta = 1$ is lower than LHS at $\beta = 1$. Since both LHS and RHS are monotonically increasing, they must intersect at exactly one point. Call that belief β_0^{FI} . Thus, β_0^{FI} exists and is unique.

Third, we need to show that if there exists a unique threshold belief β_0^{FI} , then $2c < q(1 - k^2)$. If there is a unique belief threshold then it must be the case that there is a unique point of intersection of LHS and RHS in condition C2. Again, RHS cuts the LHS from above because at $\beta = 0$ $c > 0$. Therefore, given the monotonicity of the two functions, a sufficient condition for uniqueness is $\text{LHS}|_{\beta=1} > \text{RHS}|_{\beta=1}$. This gives $\frac{q}{2} + (1 - q)\frac{k^2}{2} > \frac{k^2}{2} + c$, which can be rearranged to $2c < q(1 - k^2)$.

Lastly, we need to show that the agent does not experiment when $2c \geq q(1 - k^2)$. This is so because then the RHS is always above the LHS, so that even experimentation once is not beneficial. When $2c \geq q(1 - k^2)$ we have that $\text{LHS}|_{\beta=1} \leq \text{RHS}|_{\beta=1}$. Given that both LHS and RHS of condition (C2) are increasing convex functions, a concern is that there might be two points of intersection. However, it is easy to verify that the slope of the RHS is lower than the slope of the LHS at both $\beta = 0$ and $\beta = 1$. This precludes such a possibility. Therefore, the agent does not want to experiment when $2c \geq q(1 - k^2)$ as the RHS is always above the LHS.

Part 2: Optimal decision rule I_t^{FI}

Condition C2 is the condition for experimenting in the worst case scenario, that is when the agent knows she is going to stop after another ℓ idea. Therefore, it follows that $I_t^{FI} = 0$ in $\beta \geq \beta_0^{FI}$ if $\theta_{t-1} = \ell$, i.e the agent continues experimenting.

Next, note that the agent cannot continue experimenting forever after ℓ ideas because at the limit the value of experimentation goes to $-c$. This is so because at the limit the belief about ability goes to zero while the cost of experimentation is a positive constant. Thus, what we need to show is that the agent does not want to experiment even once when condition C2 does not hold, i.e. $I_t^{FI} = 1$ for beliefs $\beta_t < \beta_0^{FI}$ if $\theta_{t-1} = \ell$ is the optimal decision rule.

Suppose not. Say that for some belief $\tilde{\beta} < \beta_0^{FI}$, it does not pay to experiment just once but it pays to experiment at least \tilde{T} times and then stop (Note from above, she does not want to experiment forever). Now at round $\tilde{T} - 1$ when belief is $\tilde{\beta}_{\tilde{T}-1}$ it must be that condition C2 holds i.e.

$$\frac{\tilde{\beta}_{\tilde{T}-1}q}{2} + (1 - \tilde{\beta}_{\tilde{T}-1}q)\frac{(\tilde{\beta}_{\tilde{T}}k)^2}{2} \geq \frac{(\tilde{\beta}_{\tilde{T}-1}k)^2}{2} + c$$

But now since $\tilde{\beta}_{\tilde{T}-1} \leq \tilde{\beta} < \beta_0^{FI}$ and we know that for any belief $\beta < \beta_0^{FI}$ condition C2 does not hold, this is a contradiction.

Finally, we have already shown the proof of the choice of e^{FI} in the main text. ■

Proof of Lemma 3

Proof. Fix the parameters such that $2c < (q + (1 - q)k)^2 - k^2$. Since, $q(1 - k^2) > (q + (1 - q)k)^2 - k^2$, both β_0^{NI} and β_0^{FI} exist and are unique. To compare β_0^{NI} and β_0^{FI} , we only need to compare the LHS of the equation that defines condition (C1) with the LHS of the equation that defines condition (C2). We can then compare them with a common RHS.

Observe that the LHS of both the conditions are increasing and convex in β . Further, as $\beta \rightarrow 0$ the LHS in both the conditions also tend to zero. Thus, to establish a relationship between them it is sufficient to look at the behaviour of the LHS as $\beta \rightarrow 1$. This is equal to $\frac{(q+(1-q)k)^2}{2}$ for condition C1 and $\frac{q+(1-q)k^2}{2}$ for condition C2. Again, it can be shown that $\frac{(q+(1-q)k)^2}{2} < \frac{q+(1-q)k^2}{2}$ which is equivalent to $q(1 - k^2) > (q + (1 - q)k)^2 - k^2$. This implies that the LHS of condition C1 lies below the LHS of condition C2 for all $\beta > 0$. Thus, $\beta_0^{NI} > \beta_0^{FI}$. ■

Proof of Proposition 2

Proof. We prove this statement in steps by considering different regions of starting prior β_1 . There exists a $j \geq 0 \in \{0, 1, 2, \dots\}$ where belief β_j^{FI} is such that $\beta_j^{FI} < \beta_0^{NI} \leq \beta_{j+1}^{FI}$. The value that j takes depends on the parameters.

Step 1: Proving babbling is a unique equilibrium for $\beta_0^{FI} \leq \beta_1 < \beta_1^{FI}$

Consider any informative strategy $\hat{\sigma}_1 \in (0, 1]$ including the truth-telling strategy. In any such strategy a message $m_1 = \ell$ is only used when $\theta_1 = \ell$. So the agent believes such a message ($\lambda_2^\ell = 1$) with the posterior about ability $\beta_2^\ell < \beta_0^{FI}$ which makes the agent experiment only once at $t = 1$ and then exert $e = \beta_2^\ell k$ (see Proposition 1). A message $m_1 = \ell$ instead leads to a higher belief $\beta_2^\ell \in (\beta_1, 1]$,

which can either push the agent to implement her idea with a higher effort or to experiment again (depending on $\hat{\sigma}_1$ and $\hat{\sigma}_2$).

If the agent best responds to $m_1 = h$ implementing her idea, she exerts effort $e = \beta_2^h(\lambda_2^h + (1 - \lambda_2^h)k) > \beta_2^\ell k$. In this case, the supervisor type $\theta_1 = \ell$ is better off deviating and sending a message $m_1 = h$ and getting a higher expected probability of success of $\beta_2^h \beta_2^\ell (\lambda_2^h + (1 - \lambda_2^h)k)k$ instead of $(\beta_2^\ell k)^2$. If the agent best responds to $m_1 = h$ by experimenting again, then also the supervisor type $\theta_1 = \ell$ is better off always sending the message $m_1 = h$. This is because the supervisor always prefers experimentation when the current idea is low potential. Thus, the supervisor has an incentive to deviate in either case.

Thus, only the babbling strategy remains which is always an equilibrium. The agent's equilibrium strategy is to implement her outside information idea, i.e. $I_1 = 1$ with $e = \beta_1 k$ since $\beta_1 < \beta_0^{NI}$ (see Lemma 1).

Step 2: Proving babbling is a unique equilibrium for $\beta_1^{FI} \leq \beta_1 < \beta_0^{NI}$

If $j = 0$, then either $\beta_0^{FI} \leq \beta_1 < \beta_0^{NI} < \beta_1^{FI}$ or $\beta_0^{FI} < \beta_0^{NI} \leq \beta_1 < \beta_1^{FI}$. In either case, the scenario highlighted in Step 2 does not exist. Step 1 is sufficient in this case.

If $j = 1$ then it is enough to show that babbling is the unique equilibrium in the range $\beta_1^{FI} \leq \beta_1 < \beta_0^{NI}$ with the knowledge that if the posterior $\beta_2 < \beta_1^{FI}$ then the supervisor babbles (from Step 1 above). Note that any informative messaging strategy conjecture for $t = 1$ with $\hat{\sigma}_1 \in (0, 1]$ must lead to a posterior $\beta_2^\ell < \beta_1 < \beta_2^h$. Now, as before the value of message $m_1 = \ell$ is the same as in truth-telling so that $\beta_2^\ell \in [\beta_0^{FI}, \beta_1^{FI})$. From Step 1 above, the supervisor is then expected to babble in $t = 2$ and the agent best responds by choosing to implement her low potential idea ($I_2 = 1$) from $t = 1$ with effort $e = \beta_2^\ell k$. A message $m_1 = h$ again leads to a higher belief $\beta_2^h \in (\beta_1, 1]$, which can either push the agent to implement her idea with a higher effort or to experiment again (depending on $\hat{\sigma}_1$ and $\hat{\sigma}_2$). As before now, the supervisor type $\theta_1 = \ell$ is better off deviating and sending a message $m_1 = h$. Thus, babbling is the unique equilibrium strategy of the supervisor.

If $j \in \{2, 3, \dots\}$, then it needs to be shown that babbling is a unique equilibrium strategy in the ranges $\beta_1^{FI} \leq \beta_1 < \beta_2^{FI}$, \dots , $\beta_{j-1}^{FI} \leq \beta_1 < \beta_j^{FI}$ and $\beta_j^{FI} \leq \beta_1 < \beta_0^{NI}$. Consider first the range $\beta_1^{FI} \leq \beta_1 < \beta_2^{FI}$. Any posterior β_2^ℓ for priors $\beta_1^{FI} \leq \beta_1 < \beta_2^{FI}$ must map in to the range of beliefs highlighted in Step 1. This implies that supervisor type $\theta_1 = \ell$ cannot credibly commit to sending a message $m_1 = \ell$. Such a message leads to the agent implementing with effort $e = \beta_2^\ell k$. This makes babbling a unique equilibrium strategy for $\beta_1^{FI} \leq \beta_1 < \beta_2^{FI}$. The same logic applies to all the ranges of prior belief up to β_j^{FI} . Then, in the range $\beta_j^{FI} \leq \beta_1 < \beta_0^{NI}$ the proof is identical to the above

described $j = 1$ case.

Therefore, babbling is the unique equilibrium strategy of the supervisor and the agent does not experiment, i.e. $I_1 = 1$ and $a = \beta_1 k$.

Step 3: Proving babbling is a unique equilibrium for $\beta_0^{NI} \leq \beta_1 < \beta_1^{NI}$

For $j = 0$, we have already shown that babbling is a unique equilibrium strategy for $\beta_0^{FI} \leq \beta_1 < \beta_0^{NI} < \beta_1^{FI}$ or $\beta_0^{FI} < \beta_0^{NI} \leq \beta_1 < \beta_1^{FI}$. Note that since $\beta_0^{FI} < \beta_0^{NI}$, it must be the case that $\beta_1^{FI} < \beta_1^{NI} < \beta_2^{FI}$. So, it remains to show that babbling is unique for $\beta_1^{FI} \leq \beta_1 < \beta_1^{NI}$. This argument is the same as the one presented below.

Any informative mixing for $j \geq 1$ leads to $\beta_2^\ell < \beta_0^{NI}$. The supervisor babbles in the range of posteriors $\beta_0^{FI} \leq \beta_2^\ell < \beta_0^{NI}$ from Step 1 and 2 above (and for $j = 0$ case the supervisor babbles in the range $\beta_0^{FI} \leq \beta_2^\ell < \beta_1^{FI}$), and the agent chooses to implement thereafter (from Lemma 1). A message $m_1 = \hbar$, on the other hand, is believed and the agent best responds by either implementing with a higher belief or experimenting again. Therefore, the supervisor can do better by lying instead when he observes $\theta_1 = \ell$ when he is expected to be informative. ■

Proof of Proposition 3

Proof. We prove the proposition in two parts.

Part 1: To show that if $\sigma_1 = 1$ is an equilibrium for $\beta_1^{NI} \leq \beta_1 < \beta_2^{NI}$, then it must be an equilibrium for all $\beta_2^{NI} \leq \beta_1 < 1$.

Consider the region of priors $\beta_2^{NI} \leq \beta_1 < \beta_3^{NI}$. We check whether $\hat{\sigma}_1 = 1$ is an equilibrium. Here, the supervisor has an incentive to reveal the truth about $\theta_1 = \ell$ if the expected probability of success by sending $m_1 = \ell$ is higher than that from sending the message $m_1 = \hbar$. If he sends a message $m_1 = \ell$, the agent at most experiments two more times - consulting the supervisor after one (which is at β_2^ℓ where the supervisor is again honest given the premise) and not doing so after the other. Therefore, the expected probability of success by sending $m_1 = \ell$ is

$$\beta_2^\ell q + (1 - \beta_2^\ell q)(\beta_3^\ell)^2(q + (1 - q)k)^2$$

By lying the supervisor convinces the agent that her idea has a high potential to succeed ($\hat{\lambda}_2^\hbar = 1$) and that she is of ability q ($\hat{\beta}_2^\hbar = 1$). She then exerts $e = 1$ to implement her idea. However, the supervisor has an updated belief of β_2^ℓ knowing that $\theta_1 = \ell$. Thus, expected probability of success by

sending $m_1 = h$ is $\beta_2^\ell k$.

The supervisor has an incentive to reveal the truth at this stage if

$$\beta_2^\ell k \leq \beta_2^\ell q + (1 - \beta_2^\ell q)(\beta_3^\ell)^2(q + (1 - q)k)^2.$$

It is easy to check that the above condition is always holds with a strict inequality sign under Assumption (A). So, $\sigma_1 = 1$ is an equilibrium for $\beta_2^{NI} \leq \beta_1 < \beta_3^{NI}$.

Now, for any $\beta_1 \geq \beta_3^{NI}$ the supervisor can make the agent experiment (and if any of the following ideas has a high potential to succeed make them exert $e = 1$ on it) at least three more times by honestly revealing $\theta = \ell$. Given Assumption (A), this should always be an equilibrium.

Part 2: To show that $\sigma_1 = 1$ is an equilibrium for $\beta_1^{NI} \leq \beta_1 < \beta_2^{NI}$ if and only if $c \geq \frac{\kappa k - (\kappa k)^2}{2}$ where $\kappa \equiv \frac{k}{(q + (1 - q)k)^2}$ and $k < (q + (1 - q)k)^2$.

Suppose $c \geq \frac{\kappa k - (\kappa k)^2}{2}$ where $\kappa \equiv \frac{k}{(q + (1 - q)k)^2}$ and $k < (q + (1 - q)k)^2$. Now consider the conjectured strategy $\hat{\sigma}_1 = 1$ for $\beta_1^{NI} \leq \beta_1 < \beta_2^{NI}$. When the supervisor observes $\theta_1 = \ell$, his expected probability of success by sending message $m_1 = \ell$ is

$$(\beta_2^\ell)^2(q + (1 - q)k)^2.$$

Following $m_1 = \ell$, the agent experiments once more but does not consult the supervisor thereafter. Thus, she implements her idea of unknown potential by exerting effort $e = \beta_2^\ell(q + (1 - q)k)$. On the other hand by sending a message $m_1 = \hat{h}$ when the agent expects supervisor to be honest leads her to exert $e = 1$ in implementing a $\theta_1 = 1$ idea. This is so because she believes in the supervisor's message, $\hat{\lambda}_2^{\hat{h}} = 1$ and $\hat{\beta}_2^{\hat{h}} = 1$. The expected probability of success is then $\beta_2^\ell k$.

Truth-telling is an equilibrium if

$$\begin{aligned} (\beta_2^\ell)^2(q + (1 - q)k)^2 &\geq \beta_2^\ell k \\ \implies \beta_1 &\geq \frac{k}{qk + (1 - q)(q + (1 - q)k)^2} := \beta^{\text{truth}}. \end{aligned}$$

$\hat{\sigma}_1 = 1$ is an equilibrium if for all $\beta_1 \in [\beta_1^{NI}, \beta_2^{NI})$, it is also the case that $\beta_1 \geq \beta^{\text{truth}}$. This can happen iff $\beta^{\text{truth}} \leq \beta_1^{NI}$. This condition then be rearranged given β^{truth} from above and $\beta_0^{NI} =$

$(\frac{2c}{(q+(1-q)k)^2-k^2})^{\frac{1}{2}}$ (from Lemma 1), and using the fact that $\beta_1^{NI} = \frac{\beta_0^{NI}}{1-q(1-\beta_0^{NI})}$. This gives us

$$c \geq \frac{\kappa k - (\kappa k)^2}{2}$$

where $\kappa \equiv \frac{k}{(q+(1-q)k)^2}$ and we need that $k < (q + (1 - q)k)^2$. But this is also our premise. Thus, $\sigma_1 = 1$ is an equilibrium.

Alternately, suppose $\sigma_1 = 1$ is an equilibrium for $\beta_1^{NI} \leq \beta_1 < \beta_2^{NI}$. Then it must be the case that $\beta_1 \geq \beta^{\text{truth}}$ for all $\beta_1 \in [\beta_1^{NI}, \beta_2^{NI})$. Specifically, it must be that $\beta_1^{NI} \geq \beta^{\text{truth}}$. This condition can then be rearranged to yield

$$c \geq \frac{\kappa k - (\kappa k)^2}{2}$$

where $\kappa \equiv \frac{k}{(q+(1-q)k)^2}$ and with an added constraint $k < (q + (1 - q)k)^2$. ■

Proof of Lemma 4

Proof. It is immediate to see that an increase in belief from β to β' such that

1. $\beta_0^{FI} \leq \beta < \beta' < \beta_0^{NI}$ is welfare improving. This is because $\frac{(\beta k)^2}{2} > \frac{(\beta' k)^2}{2}$ which we get by replacing the optimal effort $e = \beta k$ in the expected utility function.
2. $\beta_0^{NI} \leq \beta < \beta' < \beta_1^{NI}$ is welfare improving. This is because $\frac{(\beta(q+(1-q)k))^2}{2} > \frac{(\beta'(q+(1-q)k))^2}{2}$ which we get by replacing the optimal effort $e = \beta(q + (1 - q)k)$ in the expected utility function.

Now consider an increase in belief from β to β' such that $\beta_j^{NI} \leq \beta < \beta' < \beta_{j+1}^{NI}$ such that $j > 1$. Denote the ex-ante expected utility or welfare of the agent at prior β by $W(\beta)$. We have that

$$W(\beta) = \beta \frac{q}{2} [1 + (1-q) + \dots + (1-q)^{j-1}] - \beta c [1 + (1-q) + \dots + (1-q)^j] + \beta (1-q)^j [K e - \frac{e^2}{2}] - (1-\beta) [(j+1)c + \frac{e^2}{2}]$$

where $K = q + (1 - q)k$. Similarly, we can write $W(\beta')$ keeping in mind that the maximum number of attempts is still $j + 1$.

Now, comparing term-by-term, it is obvious that everything other than the comparison of $\beta'(1 - q)^j K e' - ((1 - \beta') + \beta'(1 - q)^j) \frac{e'^2}{2}$ with $\beta(1 - q)^j K e - ((1 - \beta) + \beta(1 - q)^j) \frac{e^2}{2}$ in $W(\beta')$ is greater than that in $W(\beta)$. Thus it is sufficient to show that

$$\beta'(1 - q)^j K e' - ((1 - \beta') + \beta'(1 - q)^j) \frac{e'^2}{2} > \beta(1 - q)^j K e - ((1 - \beta) + \beta(1 - q)^j) \frac{e^2}{2}$$

which can be rearranged to

$$\beta'(1-q)^j K e' - (1 - \beta'(1 - (1-q)^j)) \frac{e'^2}{2} > \beta(1-q)^j K e - (1 - \beta(1 - (1-q)^j)) \frac{e^2}{2}$$

where $e = K\beta_{j+1}^\ell$ and $e' = K\beta'_{j+1}^\ell$.

Now it is easy to check that $Ke - \frac{e^2}{2}$ is increasing in beliefs. So that

$$\begin{aligned} K e' - \frac{e'^2}{2} &> K e - \frac{e^2}{2} \\ \implies K e' - (1 - \beta'(1 - (1-q)^j)) \frac{e'^2}{2} &> K e - (1 - \beta(1 - (1-q)^j)) \frac{e^2}{2} \\ \implies \beta'(1-q)^j K e' - (1 - \beta'(1 - (1-q)^j)) \frac{e'^2}{2} &> \beta(1-q)^j K e - (1 - \beta(1 - (1-q)^j)) \frac{e^2}{2} \end{aligned}$$

where in the second step the inequality is preserved because a greater number is added to the LHS than the RHS. And in the third step the inequality is again preserved because Ke' (which is greater than Ke) on the LHS is multiplied with a greater number than Ke in the RHS. Hence, the welfare has increased. ■

Proof of Lemma 5

Proof. Using the language introduced in Lemma 4, we can write $W(\beta)$ and $W(\beta')$ where $\beta < \beta_0^{NI}$ and $\beta_0^{NI} \leq \beta < \beta_1^{NI}$ as

$$W(\beta) = \frac{(\beta k)^2}{2} \text{ and } W(\beta') = \frac{(\beta' K)^2}{2} - c$$

Now, if the agent finds herself in $[\beta_0^{NI}, \beta_1^{NI})$ then Condition (C1) must be slack. This means

$$\frac{(\beta' K)^2}{2} - c > \frac{(\beta' k)^2}{2} > \frac{(\beta k)^2}{2}$$

where the second inequality follows from the fact that $\beta' > \beta$. Hence, the welfare has increased. ■

Proof of Lemma 6

Proof. We show here the proof of how an increase in belief from $\beta = \beta_1^{NI} - \epsilon$ to $\beta' = \beta_1^{NI}$ is welfare improving. The general proof of an increase in the prior from $\beta_{j+1}^{NI} - \epsilon$ to β_{j+1}^{NI} follows the same argument.

We can write the ex-ante expected welfare in the two cases as follows:

$$W(\beta_1^{NI} - \epsilon) = (\beta_1^{NI} - \epsilon)Ke - \frac{e^2}{2} - c$$

$$W(\beta_1^{NI}) = \beta_1^{NI}\frac{q}{2} + \beta_1^{NI}(1-q)Ke' - (c + \frac{e'^2}{2})(1 - \beta_1^{NI}q) - c$$

where $e = (\beta_1^{NI} - \epsilon)K$ and $e' = \beta_0^{NI}K$.

Now, if $W(\beta_1^{NI}) > W(\beta_1^{NI} - \epsilon)$, then substituting for e and e' , letting $\epsilon \rightarrow 0$, and simplifying the inequality by using $\beta_0^{NI} = \frac{(1-q)\beta_1^{NI}}{1-q\beta_1^{NI}}$ gives

$$\beta_1^{NI}\frac{q}{2} - c(1 - \beta_1^{NI}q) - \frac{(\beta_1^{NI}K)^2}{2} > -\frac{K^2}{2}(1 - q)\beta_1^{NI}\beta_0^{NI}.$$

If the above inequality holds, then we are done.

Let $2c < q(1 - K^2)$. Under this assumption, Condition (C2) must hold in a way that k is replaced with K as

$$\beta_1^{NI}\frac{q}{2} - c > \frac{(\beta_1^{NI}K)^2}{2} - (1 - \beta_1^{NI}q)\frac{(\beta_0^{NI}K)^2}{2}.$$

Now, the inequality is preserved if the c on the LHS is reduced. Then rearranging gives

$$\beta_1^{NI}\frac{q}{2} - (1 - \beta_1^{NI}q)c - \frac{(\beta_1^{NI}K)^2}{2} > -(1 - \beta_1^{NI}q)\frac{(\beta_0^{NI}K)^2}{2}.$$

It is now straightforward to verify that $(1 - \beta_1^{NI}q)\frac{(\beta_0^{NI}K)^2}{2} = \frac{K^2}{2}(1 - q)\beta_1^{NI}\beta_0^{NI}$, so that our original hypothesis on welfare comparison holds. ■

Proof of Proposition 4

Proof. From Lemma 4 and 5, it is immediate that an increase in belief of up to, but not including the level β_1^{NI} is welfare improving. Now, from Lemma 6, an epsilon increase in belief that pushes the agent in to experimentation with supervision is also welfare improving. Finally, from Lemma 4, any increase in belief of up to but not including the level β_2^{NI} is welfare improving. This reasoning can then be extended for any increase in belief. ■

Proof of Proposition 5

Proof. To prove the statement, we consider two particular situations, and show how in each the welfare at the correct and overconfident beliefs differ. Let $W(\beta; b)$ be the ex-ante expected utility of

the agent when the common prior is β but the correct belief is b .

Part 1: Showing that overconfidence can be welfare improving

Let $\beta = \beta_1^{NI}$ but $b = \beta_1^{NI} - \epsilon$. The two expected utility functions can be written as

$$W(b; b) = \frac{(bK)^2}{2} - c$$

$$W(\beta_1^{NI}; b) = \frac{bq}{2} + b(1-q)\beta_0^{NI}K^2 - (1-bq)\left(c + \frac{(\beta_0^{NI}K)^2}{2}\right) - c$$

We need to show if $W(\beta_1^{NI}; b) > W(b; b)$. In order to do so, first observe that $\frac{bq}{2} > \frac{(bK)^2}{2}$. This follows immediately from Assumption (A). So if we are able to show that

$$b(1-q)\beta_0^{NI}K^2 - (1-bq)\left(c + \frac{(\beta_0^{NI}K)^2}{2}\right) \geq 0$$

then we are done. Rearranging the above and recognizing that $\frac{b(1-q)}{1-bq} = \beta_0^{NI} - \epsilon'$ where $\epsilon' \neq \epsilon$, we need that

$$\frac{(\beta_0^{NI}K)^2}{2} \geq \epsilon'\beta_0^{NI}K^2 + c$$

But we know from Condition (C1) that

$$\frac{(\beta_0^{NI}K)^2}{2} = \frac{(\beta_0^{NI}k)^2}{2} + c.$$

Therefore, it is possible to find an ϵ' (and consequently ϵ) such that welfare improves under overconfidence. This requires $\epsilon' \leq \beta_0^{NI} \frac{k^2}{2K^2}$.

Part 2: Showing that overconfidence can be welfare reducing

Let $\beta = \beta_0^{NI}$ but $b < \beta_0^{NI}$. The two expected utility functions can be written as

$$W(b; b) = \frac{(bk)^2}{2}$$

$$W(\beta_0^{NI}; b) = b\beta_0^{NI}K^2 - \frac{(\beta_0^{NI}K)^2}{2} - c$$

This time we need to show that $W(\beta_0^{NI}; b) < W(b; b)$. Again using Condition (C1) to substitute for $-\frac{(\beta_0^{NI}K)^2}{2} - c$ in $W(\beta_0^{NI}; b)$, we can reduce the above to

$$b < \beta_0^{NI} \left(\frac{2K^2}{k^2} - 1 \right),$$

which must always be true because $\frac{2K^2}{k^2} - 1 > 1$. ■

Proof of Proposition 6

Proof. Let $\phi_S = \phi_A = 1$. Consider the supervisor who has seen a $\theta_{t-1} = \ell$ and reveals it honestly to the agent. His value function is given by

$$\mathcal{V}_S^\ell(\beta_t) = \max \left\{ \frac{(\beta_t k)^2}{2}, \frac{\beta_t q}{2} - c_S + (1 - \beta_t q) \mathcal{V}^\ell(\beta_{t+1}) \right\}.$$

where the first term is the value that the supervisor would get if he gets the low idea implemented and the second term is what he would get if he gets experimentation again. Given his costs, he would then like the agent to continue experimenting for as long as

$$\frac{\beta q}{2} + (1 - \beta q) \frac{(\beta' k)^2}{2} \geq \frac{(\beta k)^2}{2} + c_S,$$

which gives a belief threshold β_{S0}^{FI} . However, under an honest strategy, the agent would like to continue experimenting for as long as

$$\frac{\beta q}{2} + (1 - \beta q) \frac{(\beta' k)^2}{2} \geq \frac{(\beta k)^2}{2} + c_A,$$

which gives a belief threshold β_{A0}^{FI} .

Now, if $c_S < c_A$ then $\beta_{S0}^{FI} < \beta_{A0}^{FI}$ so that the supervisor would like the agent to experiment beyond β_{A0}^{FI} . The supervisor then fears discouraging the agent through honest revelation for any prior belief that leads the agent to a belief lower than β_{A0}^{FI} . Therefore, the results of Propositions 1, 2 and 3 hold.

On the other hand if $c_S \geq c_A$, then $\beta_{S0}^{FI} > \beta_{A0}^{FI}$. The agent would like to experiment more than β_{S0}^{FI} . Consider a belief $\beta_{S0}^{FI} \leq \beta_1 < \beta_{S1}^{FI}$ and consider the expected strategy of honesty. When the supervisor has seen a low potential idea, then by announcing it truthfully he gets an effort of $e = \beta_2^\ell k$ which is also optimal from the point of view of the supervisor because $\phi_S = 1$. This is so because it is an equilibrium strategy for the supervisor to babble tomorrow. So, even though the agent at this stage would like to experiment again but in the absence of honesty tomorrow, and $\beta_2^\ell < \beta_{A0}^{NI}$ she prefers to implement. By deviating and calling it a high potential idea, he induces an effort of 1 on a low-potential idea. However, this is suboptimal from his perspective, since he would also the full cost of implementation. Thus, there is no incentive to lie and honesty is an equilibrium. ■

Proof of Proposition 8

Proof. Consider a prior $\beta_j^{NI} \leq \beta_1 < \beta_{j+1}^{NI}$. In an immediately honest equilibrium strategy, the agent experiments for j rounds with subsequent messages $m = \ell$ before reaching the babbling region so that $\beta_0^{NI} \leq \beta_j^\ell < \beta_1^{NI}$. In addition, the agent experiments one extra round without supervision. Now, consider any strategy that reveals some $j' \leq j$ outcomes together. Let the round of eventual revelation be denoted by τ . Now, the agent is induced to experiment a higher number of rounds in this strategy iff $\beta_\tau^\ell < \beta_0^{NI} \leq \beta_j^\ell$. Say that this is the case. We determine whether such a strategy is an equilibrium.

Observe that at $\beta_\tau^\ell < \beta_0^{NI}$ the agent best responds by abandoning experimentation and implementing any one of her low potential ideas with an effort $\beta_\tau^\ell k$. If the supervisor is honest, his expected payoff is $(\beta_\tau^\ell k)^2$. By deviating, and calling any one of the low potential ideas a high one, the supervisor is able to induce an effort of 1 by the agent on that idea. This gives the supervisor an expected payoff of $\beta_\tau^\ell k$. Since the latter is greater than the former, such an eventually honest strategy cannot be an equilibrium. ■

Appendix B: Additional proofs not in the main text

Comparative statics of β_0^{FI}

Lemma 7 β_0^{FI} is increasing in e , increasing in k , and decreasing in q .

Proof. Consider, first, an exogenous increase in e . It is easy to verify that an increase in e raises the value of the RHS (i.e. of implementing the idea) in condition C2 for every belief level β . This raises the β_0^{FI} .

Second, consider the effect of an exogenous increase in k .

$$\begin{aligned}\frac{\partial \text{LHS}}{\partial k^2} &= (1 - \beta q) \frac{(\beta')^2}{2} \\ \frac{\partial \text{RHS}}{\partial k^2} &= \frac{\beta^2}{2}.\end{aligned}$$

Now, since $\beta > \beta'$ and $1 > \beta q$, $\frac{\partial \text{LHS}}{\partial k^2} < \frac{\partial^2 \text{RHS}}{\partial k^2}$. Thus, the value from implementing increases by more than the value from experimenting, which leads to a higher β_0^{FI} .

Finally, consider an exogenous increase in q . The RHS remains unchanged with an increase in q . For the LHS,

$$\frac{\partial \text{LHS}}{\partial q} = \frac{\beta}{2} - k^2 \beta \beta' \left(1 - \frac{\beta'}{2}\right).$$

This is positive if $\frac{1}{2} > k^2 \beta' \left(1 - \frac{\beta'}{2}\right)$, which is true since $\frac{\partial k^2 \beta' \left(1 - \frac{\beta'}{2}\right)}{\partial \beta'} = k^2(1 - \beta') > 0$ and at the limits the inequality holds. As $\beta' \rightarrow 0$, we have that $k^2 \beta' \left(1 - \frac{\beta'}{2}\right) \rightarrow 0$ and as $\beta' \rightarrow 1$, $k^2 \beta' \left(1 - \frac{\beta'}{2}\right) \rightarrow \frac{k^2}{2}$.

■

An exogenous increase in k makes executing a low potential idea more attractive and therefore, leads to a higher β_0^{FI} and reduces the incentives to experiment for long. The agent desires to finish the project with a sufficiently high belief so that he can exert a higher effort in implementing a low potential idea (if need be), thereby maximizing the probability of success even with a poor idea. Finally, an increase in q lowers the belief threshold. This is so because conditional on being of high-ability, a higher q increases the chances of coming up with a high potential idea. Therefore, in a world in which ability is unknown it makes experimentation more attractive and pushes the agent to experiment for longer.

Comparative statics of β_0^{NI}

It is straightforward to derive how β_0^{NI} behaves with a change in parameters. A decrease in the probability of coming up with a high potential idea q or an increase in the cost of experimentation c has the effect of increasing the threshold. Finally, an increase in k can have a non-monotonic effect on β_0^{NI} depending on the initial value. For $k < \frac{1-q}{2-q}$, an increase in k decreases β_0^{NI} . For $k > \frac{1-q}{2-q}$, an increase in k increases β_0^{NI} . The intuition behind a non-monotonic relation between k and β_0^{NI} is as follows. k measures the success rate (for any given effort level) from a bad idea when the agent is of high-ability. When the agent does not observe the value of θ from experimentation, then she experiments only as a gamble (and this gamble is not worth taking more than once). When k increases from a sufficiently low k to begin with, it makes this gamble more attractive – the agent reasons that even if the gamble fails (i.e. $\theta = \ell$ is the outcome of the gamble), she is more likely to succeed because of a higher k . On the other hand, when k increases further from an already high level, then the gamble becomes less attractive. This is so because the agent already has an outside option $\bar{\theta} = \ell$ available which then becomes relatively more attractive to finish.

Appendix C: Committed supervisor

A note on the enforcement of commitment

Here we present the case of the supervisor committing to an information policy before the agent starts experimenting with ideas. Before we do so, we should understand how such a commitment may be enforced. An information disclosure policy is a sequence of revelation strategies about the observed potential of ideas produced by the agent to which the supervisor is committed. One may imagine the policy as a sequence of public tests - the supervisor may or may not observe the true potential of the idea but he designs tests that will reveal to the agent (and to the supervisor) the true potential of the idea. Thus, commitment to information disclosure policy is akin to commitment to test designs. This interpretation is in the spirit of Kamenica and Gentzkow (2011) and Smolin (2017).

Another way in which such a commitment may be enforced is through “presentation” of ideas to multiple supervisors. Many co-supervisors rather than one main supervisor may work to discipline each other. This requires that if the optimal disclosure policy involves mixing by the supervisors then they all should agree on such a mixing and then enforce it (say by punishing deviations with full disclosure). Alternately, one supervisor’s recommendation may be cross-examined by another supervisor who has also observed the agent’s idea. However, these interpretations are not immediate and might not be realistic in many settings. An apprentice working on a project might only be assigned one expert due to cost concerns. It is also not obvious how a supervisor might commit to a test design that reveals his private information to the agent. Because of this limitation, we present the commitment case as an extension of the model in Section 4. We consider here only the flavour of an optimal policy by discussing the incentives of the supervisor and the agent, and showing how the supervisor can achieve better outcomes (relative to the equilibrium outcome) for both himself and the agent by committing to information disclosure policies.

Immediate honesty

Consider first the policy in which the supervisor is committed to revealing the true potential of the idea after each round of experimentation. We call this a policy of *immediate honesty*. As illustrated in Lemma 2 such a policy induces the agent to experiment with continued low potential ideas all the way down to the belief β_0^{FI} . It is immediate that the agent prefers to experiment more under this policy relative to the equilibrium outlined in Proposition 3. Immediate honesty guarantees maximum possible learning to the agent and in the least cost, which allows the agent to match effort to the true

potential of the idea. This helps retain the attractiveness of experimentation insofar as condition (C2) holds. The prior β_1 determines how many more rounds the agent ends up experimenting under this policy relative to the equilibrium.

That the supervisor prefers such a policy is not immediate in the region of beliefs in which the supervisor is honest in equilibrium as well. While on the one hand such a policy induces more experimentation (and therefore, a higher probability of the agent producing a high potential idea), it also depresses the effort of the agent when she does not ever produce a high potential idea. The agent exerts a higher effort in equilibrium on an idea of unknown potential (see Proposition 3) because of a higher belief. Let $\beta_1 > \beta_1^{NI}$ such that under both the equilibrium and the immediately honest policy the agent experiments for t rounds until β_1^{NI} , then in equilibrium the agent experiments for one additional round (without supervision) while under the immediately honest policy she does so for t' additional rounds with supervision. Note that t and t' are functions of β_1 . The supervisor prefers the immediately honest policy over the equilibrium policy iff

$$\begin{aligned} (\beta_{t+1}^\ell)^2(q + (1 - q)k)^2 &< \beta_{t+1}^\ell q + (1 - \beta_{t+1}^\ell)\beta_{t+2}^\ell q + \\ &+ (1 - \beta_{t+1}^\ell)(1 - \beta_{t+2}^\ell)\beta_{t+3}^\ell q + \\ &+ \dots + (1 - \beta_{t+1}^\ell)(1 - \beta_{t+2}^\ell)\dots(1 - \beta_{t+t'}^\ell)(\beta_{t+t'+1}^\ell k)^2. \end{aligned}$$

Until round t both policies yield the same payoff to the supervisor. The LHS captures the additional payoff from one more round of experimentation in $t + 1$. The RHS captures increase in the payoff from t' additional rounds of experimentation with the agent implementing a low potential idea in round $t + t' + 1$. A sufficient condition for the above to be satisfied is $q > (q + (1 - q)k)^2$, which we know is satisfied from Assumption (A). Lemma 8 follows from the above discussion.

Lemma 8 *The immediately honest policy is pareto superior to the equilibrium policy.*

Thus, both the supervisor and the agent stand to gain if the supervisor commits to honesty. However, as we show below, the supervisor can do better than immediate honesty.

Delayed honesty

The supervisor's preferred policy is driven by the desire to make the agent experiment more when she has low potential ideas but implement immediately if she gets a high potential idea. Thus, while on the one hand he wants to be honest with the agent, he also wants the agent to experiment as often

as possible. We show how the supervisor can fulfil these two objectives through a delayed disclosure policy which we call *delayed honesty* and quantify the gain attainable over immediate honesty.²⁰

A policy is a combination of a disclosure time and what to recommend at that disclosure time. A disclosure timing rule is a mapping from the current belief β_t to a choice of round τ at which the supervisor requires the agent to show her ideas to him (or equivalently the number of rounds the agent is required to experiment). He then makes a comment about each of the τ ideas according to a recommendation policy which is a mapping of $\{\ell, \mathbf{h}\}^\tau$ onto itself. A recommendation policy is honest if the supervisor honestly reveals the type of all the ideas that the agent has produced. We restrict attention to honest recommendation policies for the time being and analyse what is the optimal disclosure time τ^* . At the disclosure time τ , the agent and the supervisor update their belief about the ability sequentially according to Bayes' rule. Thus, if the supervisor reveals that any of the ideas are high potential they both update their belief to 1 and otherwise revise their belief downwards by τ times

$$\beta_\tau^\ell = \frac{(1-q)^\tau \beta_1}{1 - q\beta_1 \sum_{t=0}^{\tau-1} (1-q)^t}.$$

Fix a prior $\beta_1 \geq \beta_0^{FI}$ and consider a disclosure policy that requires the agent to experiment at least τ times to receive feedback from the supervisor. We are interested in finding out the *maximum* number of rounds of delay. Let the disclosure policy be such that after the agent discovers all her ideas were of low potential she quits experimentation and implements any one her ideas, i.e. $\beta_{\tau+1}^\ell < \beta_0^{FI}$.²¹ We say that such a policy is *implementable* if the agent prefers to experiment τ times and receiving feedback to not experimenting and implementing her outside option idea.²² This yields the following implementability constraint (IC)

$$\underbrace{\frac{1}{2}\beta_1[1 - (1-q)^\tau(1 - (\beta_{\tau+1}^\ell k)^2)]}_{\text{expected benefit of experimentation}} \geq \underbrace{\frac{(\beta_1 k)^2}{2}}_{\text{opportunity cost}} + \underbrace{\tau c}_{\text{actual cost}}. \quad (\text{IC})$$

Observe that since the agent is expected to carry out multiple rounds of experimentation without knowing their outcome, she evaluates the possibility of attaining a high potential idea relative to β_1 . Conditional on being high-ability, with probability $(1-q)^\tau$ she expects to attain only low potential

²⁰Since we are not focussing on delayed partial disclosure, we will omit any mathematical complexity that comes with it such as that of defining mixed strategies. We will focus on the supervisor using pure strategies.

²¹If there is any implementable delayed policy that leads to a posterior above β_0^{FI} , then the same can be achieved by an immediately honest policy by inducing the same number of rounds of experimentation. We will refer to delayed honesty policy as the one which leads to posteriors below β_0^{FI} so that more number of rounds are induced than in immediately honest policy.

²²There is no expected benefit to the agent by experimenting less than τ times since given the policy the supervisor does not reveal any information to the agent when this is the case.

ideas to implement, and with the remaining probability she expects at least one high potential idea. Therefore, with probability $\beta_1(1 - (1 - q)^\tau)$ she receives $1/2$ and with probability $\beta_1(1 - q)^\tau$ she will revise her belief down to $\beta_{\tau+1}^\ell$ after the supervisor honestly reveals all her τ ideas are low potential. At this point, she will implement any one of her low potential ideas to obtain an expected benefit of $\frac{(\beta_{\tau+1}^\ell k)^2}{2}$. Finally, there is no benefit of experimentation if the agent is of low-ability type. This is captured in the LHS of IC condition as the expected benefit of experimentation.

If the agent instead opts for implementing her low potential outside option idea, she expects to receive a payoff of $\frac{(\beta_1 k)^2}{2}$. As illustrated in the RHS, she must forego this expected benefit when she decides to experiment, in addition to paying the cost of experimentation c for τ rounds. The IC condition thus puts a limit on the maximum number of rounds the agent is willing to experiment when she is at a belief β_1 and the supervisor is committed to revealing all the information after those rounds.

We next analyse the supervisor's incentives under such a policy. The supervisor's ex-ante expected payoff from a τ -implementable policy is

$$\beta_1[1 - (1 - q)^\tau(1 - (\beta_{\tau+1}^\ell k)^2)].$$

The supervisor, like the agent, only sees the potential of the ideas once they are presented to him – he evaluates the probability of at least one high potential idea among the τ attempts according to β_1 . Does the supervisor benefit from a higher or a lower τ ? While on the one hand a higher τ reduces the probability of the agent only producing low potential ideas, but on the other hand it also depresses the effort of the agent in case of such event. The following lemma shows that the first order effect of reduced probability dominates the second order effect of reduced effort so that the supervisor is always better off inducing a higher τ .

Lemma 9 *Under assumption (A), the supervisor's payoffs are increasing in the number of rounds the agent experiments τ .*

Proof. Consider the expected probability of success from a τ -implementable policy:

$$\beta_1[1 - (1 - q)^\tau(1 - (\beta_{\tau+1}^\ell k)^2)] \tag{9}$$

Now consider the expected probability of success from a $\tau + 1$ -implementable policy:

$$\beta_1[1 - (1 - q)^{\tau+1}(1 - (\beta_{\tau+2}^\ell k)^2)] \tag{10}$$

Subtracting equation (9) from (10) and looking at the condition for it being positive, we get

$$q + (1 - q)(\beta_{\tau+2}^\ell k)^2 - (\beta_{\tau+1}^\ell k)^2 > 0$$

This always the case since $q > k$ from Assumption (A), which implies $q > (\beta_{\tau+1}^\ell k)^2$. Therefore, the payoff of the supervisor is increasing in the number of rounds of experimentation. ■

Supervisor's maximization problem therefore reduces to getting the agent to experiment as many rounds as possible. This is solely determined by the IC condition. It is immediate that the expected benefit of experimentation to the agent under such a policy, although increasing in β_1 , is bounded above by 1/2. Consequently, for a higher β_1 the agent should want to experiment more number of rounds but up to a limit. This limit is imposed by the bounded benefits on the one hand, and the increasing cost of experimentation on the other. Our objective is to determine the maximum (β_1, τ) combination that is implementable with such a policy.

For this purpose, fix τ . Now, if there exists a prior belief that makes the IC condition bind, then it must be the *minimum* prior that does so. Define this minimum prior belief by $\bar{\beta}^\tau$. So for any belief $\beta_1 \geq \bar{\beta}^\tau$ the agent finds it optimal to at least experiment τ times. Observe that $\bar{\beta}^\tau$ must be increasing in τ since the agent must have a higher belief to induce him to experiment more often by paying a higher cost. Let $\bar{\beta}^{\bar{\tau}}$ be the maximum of this increasing sequence so that $\bar{\tau}$ gives the maximum number of rounds that are implementable and $\bar{\beta}^{\bar{\tau}}$ is the minimum prior that can induce those many rounds. Proposition 9 follows from the above discussion.

Proposition 9 *The maximum number of rounds τ^* the supervisor can delay honestly revealing the outcomes and therefore induce experimentation at prior β_1 is given by*

$$\bar{\beta}^{\tau^*} \leq \beta_1 < \bar{\beta}^{\tau^*+1} \text{ if } \beta_1 \leq \bar{\beta}^{\bar{\tau}},$$

and is equal to $\bar{\tau}$ if $\beta_1 > \bar{\beta}^{\bar{\tau}}$.

We end this section with the following observation.

Observation 2 *The supervisor weakly prefers a policy of delayed honesty to immediate honesty when delayed honesty is implementable, i.e. when $\beta_1 \leq \bar{\beta}^{\bar{\tau}}$.*

Ali (2017) derives the same result when determining the optimal dynamic disclosure policy in a slightly different environment. In his setting, the agent needs two consecutive successes in order to be successful in the project. The experiments yield success with a positive probability only if the project

is of a good type. Ali shows that the more informed party always has an incentive to delay information revelation while the less informed party would prefer early revelation. While we do not solve for the optimal policy here, we showed here delaying may be preferred to immediately revealing the outcome by the supervisor.

For priors above $\bar{\beta}^r$, a combination of immediate honesty and delayed honesty may be preferred by the supervisor. The prospect of finding out the outcome of experimentation immediately after experimenting makes the agent assess future costs probabilistically. Since it might be determined immediately that the last idea had a high potential to succeed, the agent then does not have to bear future costs of experimenting. This reduces the expected cost of experimentation to the agent and makes her willing to experiment. So for higher beliefs, where the agent is not willing to pay a lump sum cost for experimenting with delayed honesty, the supervisor can induce experimentation with immediate honesty. The supervisor can then commit to delayed honesty when the agent reaches a lower belief. However, immediate honesty might provide too much incentive to the agent and the supervisor might do better by committing to a mixed revelation for high beliefs.²³

²³We do not consider these policies in this paper as our primary objective is to highlight the tensions when the supervisor does not have commitment power. We merely want to show how the supervisor can do better when there is commitment in the relationship, and what incentives shape a “preferred” policy.