Effective supervision*

Zeinab Aboutalebi
Department of Economics,
European University Institute †

Ayush Pant
Department of Economics,
Ashoka University ‡

October, 2022

Abstract

When can an employer effectively motivate its employees to experiment through supervision? Only when both the employees *and* the supervisors are confident in the employees' ability. We obtain this result in the context of an agent-supervisor model with experimentation and feedback. An agent works on a project, comprising an idea generation phase followed by an idea implementation phase. A supervisor provides unverifiable feedback on her ideas. We investigate when the supervisor can effectively persuade the agent to continue experimenting with ideas. Giving feedback involves the following tradeoff for the supervisor: while honest feedback encourages the agent to discard bad ideas, it can also be demotivating. Our framework provides novel insight into the gender performance gap, as less confident women may underperform due to ineffective supervision. We also discuss relevant policy implications of our results.

Keywords: Feedback; Self-confidence; Dynamic cheap talk; Experimentation

JEL Classification Numbers: C73, D83, M53

---

*This work was supported by the Economic and Social Research Council [grant number ES/J500203/1]. An older version of this paper was circulated under the title "Feedback on ideas." 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, Rahul Deb, Daniel Habermacher, James Fenske, 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.

†Email: zeinab.aboutalebi@eui.eu

‡Email: ayush.pant@ashoka.edu.in

1 Introduction

Strategic disclosure of the information is a potent tool in designing incentives and enhancing performance of employees (Halac, Kartik, and Liu (2017), Ely (2017)). Notably, how a supervisor provides information to the employee can influence what tasks she works on and her effort levels. However, when the supervisor's information is unverifiable, and commitment to a disclosure policy cannot be enforced, he may not be able to influence employees' decisions. For instance, universities assign supervisors to graduate students who, because of their experience, can provide feedback to the students on their ideas. However, since the supervisor cannot commit to an information policy, he may not always provide informative feedback and be able to influence students' choices.[1] Thus, knowing when supervision can improve and alter employees' incentives is valuable.

To illustrate the problem, consider an employer with two employees, Anne (she) and Bob (he), who experiment with ideas for separate projects and implement one each. The employer would like both to produce successful projects. However, without information on their ideas, both under-experiment. To fix the issue, the employer considers allocating them a supervisor with the expectation that supervision would fuel experimentation. However, the supervisor's information is unverifiable, and so is the feedback he provides. Moreover, he may be apprehensive about giving honest feedback; while encouraging employees to discard bad ideas, critical feedback can demoralize and discourage idea generation and implementation effort. So, a supervisor's feedback may be contingent on confidence levels.

In this paper, we formalize the role of confidence in giving and receiving critical feedback and its impact on the employee's performance. In the illustrated situation,

[1] Similarly, organizations often assign more experienced supervisors to early-career creatives. For instance, a partner in a law firm supervises an associate developing a litigation strategy; a project manager in a tech firm supervises an engineer solving a bug in app development; a senior designer in an architecture firm supervises a junior designer looking for a design solution. A supervisor cannot commit to an information disclosure policy in all these examples.

suppose Anne and Bob are identical in all respects other than either (1) Bob is more self-confident than Anne, or (2) the supervisor is more confident in Bob's ability than Anne's. We show that in either case, the supervisor favors Bob over Anne by providing him with more critical feedback while greenlighting Anne's bad idea earlier. As a result, supervision only helps employees with high self-confidence or those in which the supervisors have high confidence.

We develop a supervisor-agent model where a firm tasks an agent (she) with a project comprising two phases – experimentation and implementation, and a supervisor (he) for providing her feedback. In the experimentation phase, the agent sequentially generates ideas at a cost, receives feedback from the supervisor about the merits of her ideas, and then 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 to all, and the agent and supervisor may not have the same prior confidence in ability. The supervisor's (and the firm's) objective is to achieve success in the project, but he does not internalize the agent's cost of effort. Success is contingent on the idea and the implementation effort. The misalignment of preferences means that the supervisor may try to mislead the agent.

The nature of feedback varies with the agent's self-confidence and the supervisor's confidence in the agent's ability. We begin by showing that the supervisor does not provide honest feedback at low self-confidence levels. At such levels, the agent's pessimism induces her to quit experimentation after bad news. As a result, in this interval, the supervisor is more concerned with the effort choice in the implementation phase and prefers to mislead the agent if she has a poor idea. A higher confidence of the supervisor does not help as the agent's effort is not contingent upon the supervisor's confidence.

Next, we show that while agent's higher self-confidence is necessary, it is not sufficient for the supervisor to give critical feedback. At these levels, when bad news yields at least one extra round of experimentation before implementation, the supervisor cares about the result of experimentation and her effort when implementing. Supervisor's honesty is possible only when both agent's self-confidence, and the supervisor's confidence are high enough. The former matters for implementation effort, while the latter matters for the agent's idea generation ability. Indeed, higher confidence of the supervisor can now partially mitigate the agent's lower self-confidence.

Fixing a high enough supervisor's confidence, the reversal of the supervisor's incentives to be honest at high levels of self-confidence creates a performance difference between the high and low self-confident agents. A more self-confident agent receives honest critical feedback more often and gets more opportunities to get the best idea and match her effort to the idea quality. Therefore, the presence of a supervisor for such self-confident agents can have a magnifying effect on their performance. On the other hand, a less self-confident agent cannot benefit from supervision, and having a supervisor does not affect her performance.

This key result of our paper provides new intuition about how a confidence gap between men (Bob) and women (Anne) in research-driven industries translates into a performance gap. A vast amount of literature documents the gender confidence gap. For example Goodman, Cunningham, and Lachapelle (2002) emphasize that the main cause of female dropouts from STEM fields is their lack of confidence in their abilities to pursue degrees in these fields.[2] We emphasize that one pathway to the underperformance of women with lower self-confidence is through the reluctance of the supervisors to be critical. As evidence of our result, Jampol and Zayas (2021) find that women are more likely to receive upwards distorted "white lies" feedback that potentially hinders their performance.

Such performance differences between less and more confident agents have welfare

[2] Niederle and Vesterlund (2007), Gneezy and Rustichini (2004), Gnther, Arslan Ekinci, Schwierenc, and Strobeld (2010) and Shurchkov (2012) attribute women's underperformance and unwillingness to participate in competitive tasks to the same confidence gap.

and policy implications. Notably, confidence-building exercises that *incorrectly* raise Anne's self-confidence can also improve her welfare. The discontinuous change in the supervisor's feedback strategy as she incorrectly goes from low to high self-confidence gives rise to this possibility. Therefore, as a first step toward providing incentives to experiment, organizations should build their employees' self-confidence. Only then will the provision of supervision start having a positive impact.

We further show that for more complex tasks, namely those in which it is harder to obtain a good idea, effective supervision may be sustained for lower self-confidence levels. As a first-order effect of increasing task difficulty, success becomes less likely for the same effort, which reduces the agent's incentives to experiment. However, as a second-order effect, it also makes failure less predictive of the ability. Therefore, the agent is less discouraged both in experimentation and implementation decisions. As a result, the supervisor finds it easier to give honest feedback for lower levels of self-confidence.

Finally, we extend our model to include agents' learning-on-the job with honest supervision. Indeed, firms may create supervisor-agent relationships to facilitate the learning of early-career employees who gain experience through supervision (Fudenberg and Rayo (2019)). We show that the magnifying effect of higher self-confidence, through feedback, on performance is even more prominent with the possibility of learning on the job. More confident agents may experiment to their maximum limit as feedback fuels learning, increasing the supervisor's incentive to provide further feedback. Those with lower confidence, on the other hand, might not even experiment once.

Related Literature. Our paper primarily contributes to the literature on dynamic persuasion.[3] Persuasion is usually modeled either as disclosure of verifiable information (Grossman (1981), Milgrom (1981)) or as information design (Kamenica and Gentzkow (2011), Ely (2017)). Our paper is closely related to Ely and Szydlowski (2020), Orlov,

[3] The problem of persuasion deals with convincing the party taking action to take it in favor of the one with the informational advantage.

Skrzypacz, and Zryumov (2020), Smolin (2021), and Ali (2017) that fall in the latter category.

Ely and Szydlowski (2020) show that by committing to a delayed information disclosure policy, the principal can persuade the agent to exert effort on more challenging projects that require more effort. Smolin (2021) shows that a minimally informative feedback policy, which only reveals whether the agent should quit given the past performance, is optimal to induce him to work longer. Our main point of departure is to focus on a situation where the informed party cannot commit ex-ante to a disclosure policy, and the messages are essentially cheap talk. In our setting, the *equilibrium* feedback strategy preserves full information revelation for highly self-confident agents, *and* persuades them not to quit, i.e., implement a bad idea, early on, despite no commitment.[4]

Similar to our model, Orlov et al. (2020) look at a principal-agent setting where production is contingent on two random variables. The agent wants to persuade the principal to wait for a high realization of one by controlling information on the other. However, unlike their model, we explore a setting where information about one variable (idea quality) affects the beliefs about the other (ability). This difference helps us obtain a different characterization of the efficacy of supervision. Notably, we can pin down the role of confidence in critical feedback and its effect on performance.[5]

We also contribute to the literature on dynamic performance evaluation and feedback. This literature largely looks at strategic feedback provision in tournaments (Lizzeri, Meyer, and Persico (2002), Aoyagi (2010), Perry and Gershkov (2009)). Ederer (2010) builds a similar model to ours under tournaments. He sets up a two-period model

[4]Persuasion is usually not the concern in cheap talk literature (Crawford and Sobel (1982), Golosov, Skreta, Tsyvinski, and Wilson (2014)) as the problem is viewed from the receiver's perspective, and informative communication need not imply persuasion. Chakraborty and Harbaugh (2010) is a notable exception, but they consider a static model with different tradeoffs for the sender. In dynamic models of cheap talk, it is usually impossible to have any influential communication without commitment (Boleslavsky and Lewis (2016)).

[5]Further, Orlov et al. (2020) retain the assumption of information design even in the no dynamic commitment version. Thus, the agent in their model has the tools to manage the principal's beliefs, unlike ours.

where the principal chooses between giving no and full information about interim performance to two agents with hidden ability who make effort choices for two periods. We differ from this work and contribute to the literature by offering a unique setting where a single agent receives dynamic feedback on his idea. We show that even without commitment, it is possible to sustain honest feedback.

Finally, we are related to the literature on feedback and self-confidence starting from Bénabou and Tirole (2002, 2003). Our paper relates to Bénabou and Tirole (2003) in the tradeoff studied. While they look at the role of reward and punishment in managing an agent's self-confidence and performance, we study information transmission to analyze the effect of self-confidence on performance. Koszegi (2006) introduces the utility of ego to capture the role of self-confidence on effort and performance in a decision-theoretic framework. There is vast experimental and theoretical behavioral economics and psychology research based on Bandura (1977) that captures such effects absent strategic motivations. Our essential contribution is to explore strategic aspects of information provision on self-confidence and performance.[6]

The paper is organized as follows: In Section 2, we present our base model. Starting with analyzing the agent's decision problem (under some natural restrictions) in Section 3, we proceed to present the main analysis in Section 4. We introduce the possibility of the agent's learning to understand her ideas with feedback in Section 5. Finally, we present alternate interpretations of our model, discuss policy implications, and conclude in Section 6.

2 A model of feedback on ideas

An agent (she) works on a project with a supervisor (he). The project involves two distinct stages that occur sequentially. The first stage involves experimenting with

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

ideas, and the second stage requires implementing a chosen idea. The agent exerts effort toward success in the project in both stages, while the supervisor provides feedback to the agent in the first stage. The supervisor has no commitment power and provides feedback based on what he observes. Success generates positive benefits to both the agent and the supervisor. Only the agent pays the cost of effort.

***Stage 1*: Experimentation with ideas.** The process of idea generation involves multiple rounds $r = 1, 2, \ldots$. In each round r, the agent decides whether she wants to draw a new idea. An idea is defined by its quality, q, which is either good, g, or bad, b. The agent's ability, a, determines the quality of the idea drawn. Ability is either high or low. We let $\Pr(q_r = g|a = \text{high}) = \theta \in (0, 1)$ and $\Pr(q_r = g|a = \text{low}) = 0$, i.e., only a high-ability agent can produce a good idea.

Ability, in contrast to the idea, is persistent throughout the game and is initially unobserved to both the agent and the supervisor. π_r and Π_r respectively denote the agent's and the supervisor's belief that the agent is high-ability at the beginning of round r. Both are non-degenerate and there is common knowledge about the other's belief.

We assume that the agent possesses a bad outside option idea at the beginning that we denote by $q_0 = b$.[7]

Stage 1 actions and timing. At the beginning of each round of experimentation, the agent decides whether to experiment again or to implement her most recent idea. Each round of experimentation costs the agent c and a new idea is realized.[8] The supervisor privately observes the quality q_r of the idea drawn.[9] He then sends a costless

[7]This assumption is not necessary for the analysis, but helps make relevant comparisons for the agent at any stage. In particular, the first stage decision is the same as all potential future decisions – implement a bad quality idea or experiment again.

[8]Such time and effort costs could arise from seeking inspiration, making online searches, looking up for data, reading material and exploring the literature.

[9]In Section 5 we relax the assumption that quality gets revealed only to the supervisor. We extend the analysis to the case were the agent gradually gains the ability to understand the quality of her own ideas without needing a supervisor. In general, however, it is reasonable to assume that the supervisor is better equipped to understand the quality of ideas that early-career agents may generate. This may be the case due to more experience.

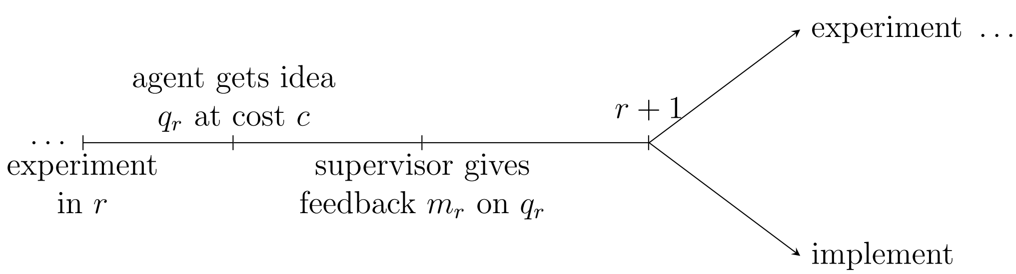

Figure 1: Summary of timing when the agent chooses to experiment in round r

message, $m_r \in \{g, b\}$ about the quality.[10]

Alternately, the agent may choose to implement and move to the second stage in round r. The move is permanent and the agent cannot return to experimenting with ideas again. Figure 1 summarizes timing and actions in Stage 1.

***Stage 2*: Implementation of idea.** Once the agent moves to implementation stage in round $r+1$, she can only implement her last idea q_r. Thus, we assume that the agent cannot implement an idea abandoned in previous rounds.

Stage 2 actions and timing. When implementing in round $r+1$, the agent exerts effort $e \in \mathbb{R}^+$ at per unit cost k for some $k > 0$ to complete the project. Success of a completed project depend on the idea and the effort. A good idea succeeds with probability $g(e) \in [0,1]$ and a bad idea with $b(e) \in [0,1]$. We make the following assumptions on the success functions: (1) $g(e) > b(e)$ for all $e > 0$. (2) $g(0) = b(0) = 0$. (3) $g' > 0$, $g'' < 0$ and $b' > 0$, $b'' < 0$ so that probability of success is increasing and concave in effort.

Assumption (1) describes the sense in which having a good idea is beneficial; for any positive effort level, the probability of success is greater with a good idea. Assumption (2) is a simplification and is made without loss of generality. It states that neither idea will succeed without implementing it. Assumption (3) is a technical assumption that ensures that the agent exerts positive effort on both ideas. Figure 2 summarizes the

[10] This restriction is without loss of generality as the type space is binary and what matters are the equilibrium mappings from the supervisor type to the message space, i.e., the meaning of the messages. Here, the messages g and b have their natural meaning and are understood as the idea quality.

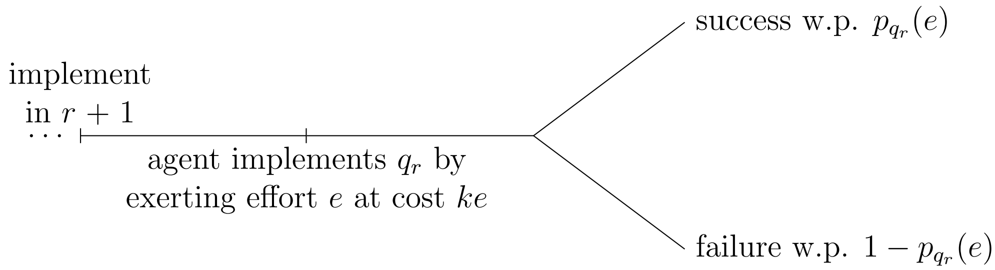

Figure 2: Summary of timing when the agent chooses to implement in round $r+1$

timing and actions in Stage 2.

Payoffs. A successful project yields a benefit of 1, and 0 if it fails. The agent gets the payoff from completion net of the costs of experimentation and implementation. The supervisor only cares about getting a successful project. That the supervisor does not bear the costs of the project captures the preference misalignment between the players. Once payoffs are realized, the game ends.

3 The agent's problem

We start by discussing the agent's problem of deciding when to stop experimenting with ideas and how much effort to exert in implementation. When making these decisions, she faces two unknowns – the quality of her idea and her ability. She can update beliefs about both through the information available to her. Let the probability that the current idea is good be λ; the probability that she is high-ability is her self-confidence, π. Her decision is, therefore, a function of these beliefs.

Consider first the agent's decision to implement. At this stage, ability does not play a role. When she believes her idea is good with probability λ, she chooses e to maximize her expected payoff $\lambda g(e)+(1-\lambda)b(e)-ke$. The FOC implicitly defines the optimal effort $e^*(\lambda)$. Let the resulting value function be $V(\lambda)=\lambda g(e^*(\lambda))+(1-\lambda)b(e^*(\lambda))-ke^*(\lambda)$. It is straightforward to derive the following comparative static results.

Lemma 1 *The optimal implementation effort e^* is increasing in λ if $g' > b'$, and* *decreasing otherwise. The expected payoff from implementing V is increasing and convex in λ.*

All proofs of results that appear in the main text are in Appendix A.

The lemma highlights the two important cases we consider. The first is when the increase in marginal probability of success from effort is higher for the better idea. As a result, effort is increasing in the belief about idea. Here, the outcomes of the two stages, idea quality and implementation effort, are *complements*. The second is when the reverse is true, and effort is decreasing in the belief about idea. In this case, idea quality and implementation effort are *substitutes*.

We now consider the agent's experimentation decision. Such decision is based on her self-confidence and the belief about her idea. As a starting point, we first discuss here the "best" case of when the agent can fully observe the quality, the full information case. We then turn to the worst case and discuss agent's decisions when she has no information on the ideas, the no information case.[11]

Under full information, there is a unique threshold on self-confidence, F_0, above which the agent experiments repeatedly after bad ideas but implements one below it. Potential learning about ideas and updating on self-confidence permits repeated experimentation. However, when the self-confidence is too low, the agent does not find it worthwhile to pay its cost any further.

In the worst case of no information, there is a unique threshold on self-confidence, N_0, above which the agent experiments once as a gamble and below which she never experiments and implements her bad idea. She does not experiment more than once as there is no learning about either the ideas or self-confidence.

It can be shown that as long as the cost of experimentation is not too large, i.e., $c < V(\theta) - V(0)$, both N_0 and F_0 exist with $F_0 < N_0$.[12] Further, it can be shown that for any starting self-confidence level $\pi_1 \geq F_0$, the agent strictly prefers to be fully

[11] All the results presented here get a formal treatment in the online Appendix B.

[12] Unless otherwise stated, we assume this cost condition holds for the rest of the analysis.

informed about her ideas than having no information. Indeed, one way to persuade the agent to experiment again after a bad idea (up to some limit) is to be honest with her. Our objective, as a result, is to determine if and when a supervisor, who privately observes quality and provides unverifiable feedback, can engage in such beneficial truthful communication without commitment. Specifically, we want to determine equilibria in which, beginning from any confidence level, the supervisor provides honest feedback at least once.

We conclude this section with an important definition. For any arbitrary self-confidence level π_r and $j = 1, 2, \ldots$, let $\phi_j(\pi_r) = \pi_{r+j}$ be the self-confidence after j rounds of application of Bayes' rule on π_r when the realized ideas were bad in all j rounds. Inverting $\phi_j^{-1}(\pi_{r+j}) = \pi_r$ gives the starting π_r that results in π_{r+j} after j rounds.

Definition 1 *Define $\phi_j^{-1}(N_0) := N_j$ and $\phi_j^{-1}(F_0) := F_j$ for $j = 1, 2, \ldots$ for the no information and full information thresholds, N_0 and F_0, respectively.*

Therefore, N_j (F_j) is the starting confidence level, which, when correctly updated about bad ideas j times, leads to the terminal level N_0 (F_0).

4 Strategic supervisor

4.1 Preliminaries

The game in the experimentation phase is one of dynamic cheap talk where the supervisor provides feedback without incurring any cost. The feedback is unverifiable and independent of the true quality. Our solution concept is (perfect) Bayesian equilibrium.

Round r begins for the agent after having observed the last message sent by the supervisor m_{r-1}. Therefore, a realized history for the agent includes the set of all previous messages sent by the supervisor up to and including m_{r-1}, and the sequence of past decisions made. Round r begins for the supervisor after observing the last idea of the agent q_r. Accordingly, a realized history for the supervisor includes, in addition to the history viewed by the agent, the sequence of all the realized ideas from past experimentation attempts.[13] A strategy for the supervisor in round r is a mapping from the realized history to possibly mixed message space $\{b, g\}$.[14]

If the agent expects the supervisor to be honest at a self-confidence level of π_r in round r, the agent's updated self-confidence in round $r+1$ is $\frac{(1-\theta)\pi_r}{1-\theta\pi_r}$ if $m_r = b$ and 1 otherwise. If the supervisor uses the same message independent of the realized history, the supervisor is said to lie or babble. The agent does not learn anything. We will assume that the agent does not consult the supervisor when she expects the supervisor to lie. So, the supervisor cannot privately learn and not reveal to the agent the outcome.

4.2 Analysis

To begin with, babbling is always an equilibrium for any self-confidence level π. The agent does not learn about her idea as the supervisor is expected to give uninformative feedback. Expectations of supervisor babbling are self-fulfilling and neither party can profitably deviate. In terms of the agent best response, this situation is equivalent to the agent making decisions without any information discussed in Section 3.

[13] Let $d^r := (d_1, \ldots, d_r)$ and $m^r := (m_1, \ldots, m_r)$ be the sequence of decisions made by the agent and the public messages given by the supervisor until round r. Denote the set of histories for the agent and the supervisor at the beginning of round r by H_r^A and H_r^S respectively. The history for the agent at the beginning of round r is

$$h_r^A = (d^{r-1}, m^{r-1}) \in H_r^A \subset (\{\text{experiment}\}^{r-1} \times \{b, g\}^{r-1}).$$

This is also the public history of the play of the game up to round r. In addition to the public history, the supervisor observes $q^r := (q_1, \ldots, q_r)$ and an extra decision of the agent to experiment, i.e., $d_r = \text{experiment}$. The history for the supervisor at the beginning of round r is

$$h_r^S = (\theta^r, d_r, h_r^A) \in H_r^S \subset (\{b, g\}^r \times \{\text{experiment}\}^r \times \{b, g\}^{r-1}).$$

[14] Note that we want to identify equilibria that induce at least one round of honesty. Thus, we look at mixed strategies only in so far as as honesty cannot be sustained at a given confidence level.

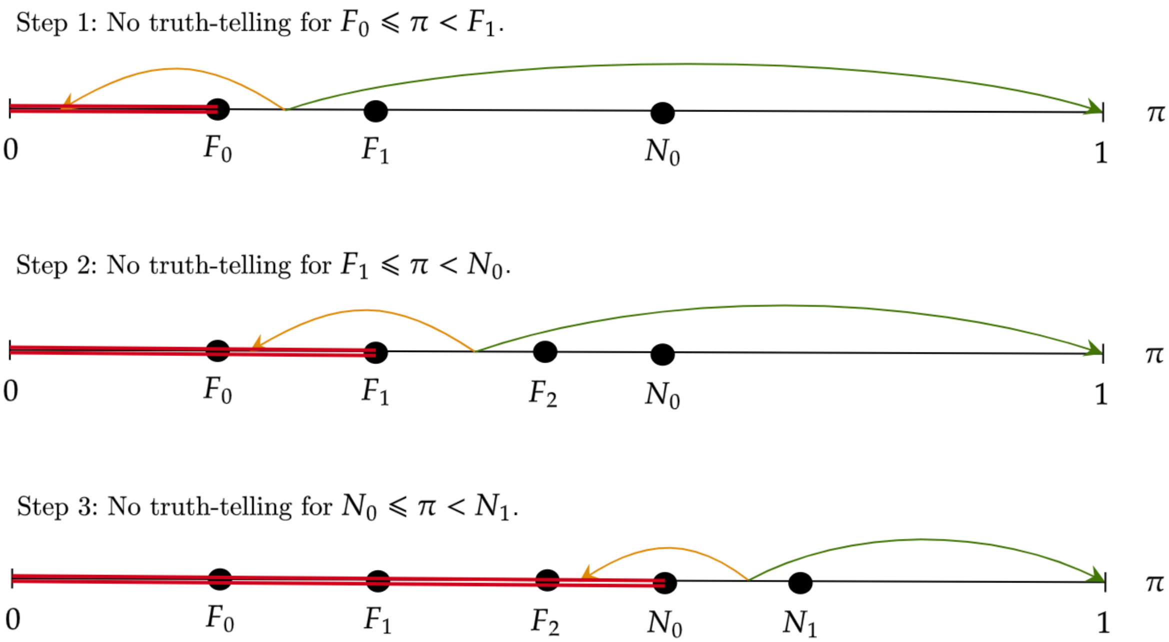

Figure 3: No truth-telling for self-confidence levels between F_0 and N_1

Observation 1 *For any self-confidence lower than the full-information threshold F_0, any communication strategy is an equilibrium and none induces the agent to experiment.*

The self-confidence region $\pi < F_0$ reflects pessimism. In this range, the prospect of obtaining a good idea is so low that the agent does not find it beneficial to even try. So, the feedback stage of the supervisor is not reached and any strategy is equilibrium.

In what follows we determine if equilibria with at least one round of honest feedback exist (in addition to babbling) for different ranges of self-confidence starting with lower ones. Our first main result defines the range for which truth-telling can never be an equilibrium. We call this the region of *low* self-confidence.

Proposition 1 *The supervisor does not give honest feedback even once in the low self-confidence region, $F_0 \leq \pi < N_1$, independent of whether idea quality and implementation effort are complements or substitutes and the supervisor's confidence in the agent's ability.*

The intuition for this proposition is illustrated in steps using Figure 3. Suppose first that idea and effort are complements. In this case, effort is increasing in the belief about idea, and in particular $e^*(0) < e^*(1)$. As $g(e^*(1))$ is the maximal probability of success, the supervisor reveals the truth when he observes good quality. The concern, however, is whether he reveals a bad idea honestly.

In the range $F_0 \leq \pi < F_1$ (see Step 1 of Figure 3), if the agent expects negative feedback in equilibrium, her self-confidence falls below F_0. Here, she abandons experimentation and chooses the *lowest* implementation effort of $e^*(0)$. By deviating and providing positive feedback, however, the supervisor induces a maximal implementation effort of $e^*(1)$. As the probability of success is increasing in effort, the supervisor always finds it beneficial to encourage the agent at this stage; truth-telling cannot be an equilibrium.

When honesty cannot be sustained, we revert to the no-information babbling equilibrium. In the absence of any information, the agent does not experiment because her self-confidence is lower than N_0. Now, the previous argument applies to those self-confidence levels that can get discouraged to levels below F_1 after negative feedback (see Step 2). In fact, the same logic can now be extended to all self-confidence levels below N_1 (illustrated in Step 3).

The absence of honest equilibria here results from the agent not experimenting further after either feedback. The argument for when idea and effort are substitutes is also the same. In this case, the agent's optimal effort declines if she more likely believes her idea to be good. So, now the supervisor is skeptical to provide honest *positive* feedback. As an illustration, consider the range $F_0 \leq \pi < F_1$. The supervisor, who does not internalize the cost of effort, now wants to call a good idea bad. In doing so, the agent exerts $e^*(0) > e^*(1)$, which increases the likelihood of success. The same unraveling argument of the complementarity case then follows.

It is worth emphasizing that since the supervisor is aware of the true quality and the agent does not experiment, his payoffs are independent of his confidence in the agent's ability.

A possibility of truthful feedback opens up for higher levels of self-confidence due to the agent's different best response to uninformative feedback. Notably, the agent experiments once in the region between N_0 and N_1. The previous threat point for the supervisor now potentially disappears. Our next result determines whether this one extra round of experimentation (without the supervisor's assistance) is sufficient for the supervisor to be honest. We will say the agent in this region has a *high* self-confidence.

Idea and effort are complements. Note that here the concern is giving negative feedback. Assuming that negative feedback leads to one *last* experimentation, the supervisor does not deviate if

$$\begin{aligned} \Pi_2\theta\, g(e^*(\pi_2\theta)) \,+\, (1-\Pi_2\theta)\, b(e^*(\pi_2\theta)) \,\geq\, b(e^*(1)) \\ \iff \Pi_2 \,\geq\, \Pi_2^{\text{truth}}(\pi_2) \,:=\, \frac{b(e^*(1)) - b(e^*(\pi_2\theta))}{\theta[g(e^*(\pi_2\theta)) - b(e^*(\pi_2\theta))]}, \end{aligned} \tag{C1}$$

where Π_2 and π_2 are respectively the supervisor's and the agent's confidence in the ability after seeing a bad idea, and $\Pi_1^{\text{truth}} = \phi_1^{-1}(\Pi_2^{\text{truth}})$.

It is easy to see that higher confidence of both players will more likely generate honesty. However, to further unpack their role, begin by considering the case of common confidence levels. Determine the belief about good idea that equates the expected probability of success under experimentation (without supervision) with the maximum probability of success that the (current) bad idea can generate. Denote this belief by $\bar{\lambda}$ (see Figure 4), which is implicitly given by the condition,

$$\bar{\lambda}\, g(e^*(\bar{\lambda})) \,+\, (1-\bar{\lambda})\, b(e^*(\bar{\lambda})) \,=\, b(e^*(1)). \tag{C1*}$$

For any belief greater than $\bar{\lambda}$, agent's effort and the weight on idea being good is higher. The expected probability of success with experimentation is larger than the maximum probability of success with a bad idea. So, for $\pi_2\theta \geq \bar{\lambda}$, the supervisor has

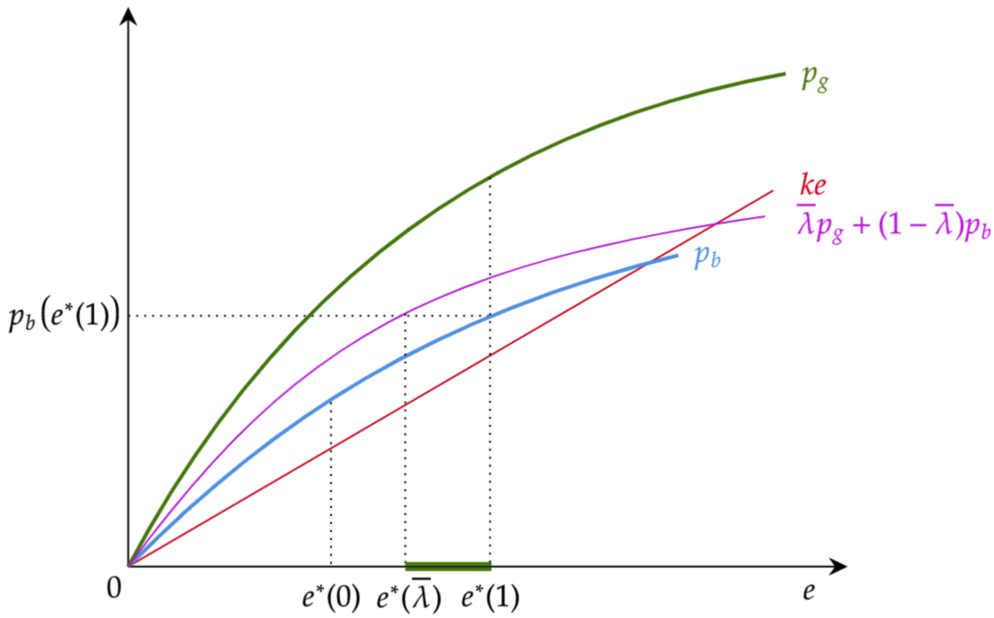

Figure 4: Honesty when idea and effort are complements

an incentive to be honest. Notably, if $N_0\theta \geq \bar{\lambda}$, the supervisor provides honest feedback for all levels of (common) confidence above N_1. Further, when confidence levels differ for the two players, then whether $\pi_2\theta$ is above (below) $\bar{\lambda}$ determines whether Π_2 can be smaller (larger) than π_2 to get honesty. Lastly, when $N_0\theta < \bar{\lambda}$, there must still be a confidence level sufficiently high so that honesty is still possible. We summarize these observations below.

Proposition 2 *Consider a self-confidence level $\pi_1 \geq N_1$, let $\theta > \bar{\lambda}$, and suppose idea and effort are complements. The supervisor gives honest feedback if and only if his confidence $\Pi_1 \geq \Pi_1^{truth}(\pi_1)$ where Π_1^{truth} is decreasing in π_1. In addition, the supervisor gives honest feedback for all self-confidence levels*

1. *above N_1 if $N_1 \geq \frac{\bar{\lambda}}{\theta(1+\bar{\lambda}-\theta)}$, and above $\frac{\bar{\lambda}}{\theta(1+\bar{\lambda}-\theta)}$ if $N_1 < \frac{\bar{\lambda}}{\theta(1+\bar{\lambda}-\theta)}$, where Π_1^{truth} is at most equal to the agent's self-confidence level;*

2. *between N_1 and $\frac{\bar{\lambda}}{\theta(1+\bar{\lambda}-\theta)}$ if $N_1 < \frac{\bar{\lambda}}{\theta(1+\bar{\lambda}-\theta)}$, where Π_1^{truth} is strictly higher than the agent's self-confidence.*

Proposition 2 shows the importance of the supervisor's confidence in addition to the agent's self-confidence. From Proposition 1 we know that at low self-confidence levels, supervisor's honesty cannot be sustained for any confidence level. Yet, to receive honest feedback both confidence levels should be high. In this sense, agent's self-confidence is necessary but not sufficient. Indeed, two agents with the same self-confidence level may get different feedback if the supervisor does not believe in their abilities equally. Alternately, even if the supervisor believes in the abilities of two agents equally, a difference in their self-confidence levels can cause the supervisor to give different levels of feedback.

When the agent and the supervisor hold the same ex-ante confidence level π_1, the corollary below captures the supervisor's truth-telling incentives.

Corollary 1 *Consider a (common) confidence level $\pi_1 \geq N_1$ and suppose idea and effort are complements.*

1. *Let $\theta > \bar{\lambda}$. The supervisor gives honest feedback for all confidence levels above N_1 if $N_1 \geq \frac{\bar{\lambda}}{\theta(1+\bar{\lambda}-\theta)}$, and only above $\frac{\bar{\lambda}}{\theta(1+\bar{\lambda}-\theta)}$ if $N_1 < \frac{\bar{\lambda}}{\theta(1+\bar{\lambda}-\theta)}$.*
2. *If $\theta \leq \bar{\lambda}$, the supervisor does not give honest feedback for any confidence level.*

Idea and effort are substitutes. In this situation, the supervisor may want to lie about good ideas. To see that no such incentives exist for honestly revealing bad ideas, note that

$$\Pi_2\theta\, g(e^*(\pi_2\theta)) \,+\, (1-\Pi_2\theta)\, b(e^*(\pi_2\theta)) > b(e^*(1))$$

because $g(e^*(\pi_2\theta)) > b(e^*(\pi_2\theta)) > b(e^*(1))$. However, the supervisor will reveal a good idea only if

$$g(e^*(1)) \;\geq\; \theta\, g(e^*(\pi_2\theta)) \,+\, (1-\theta)\, b(e^*(\pi_2\theta)), \tag{C2}$$

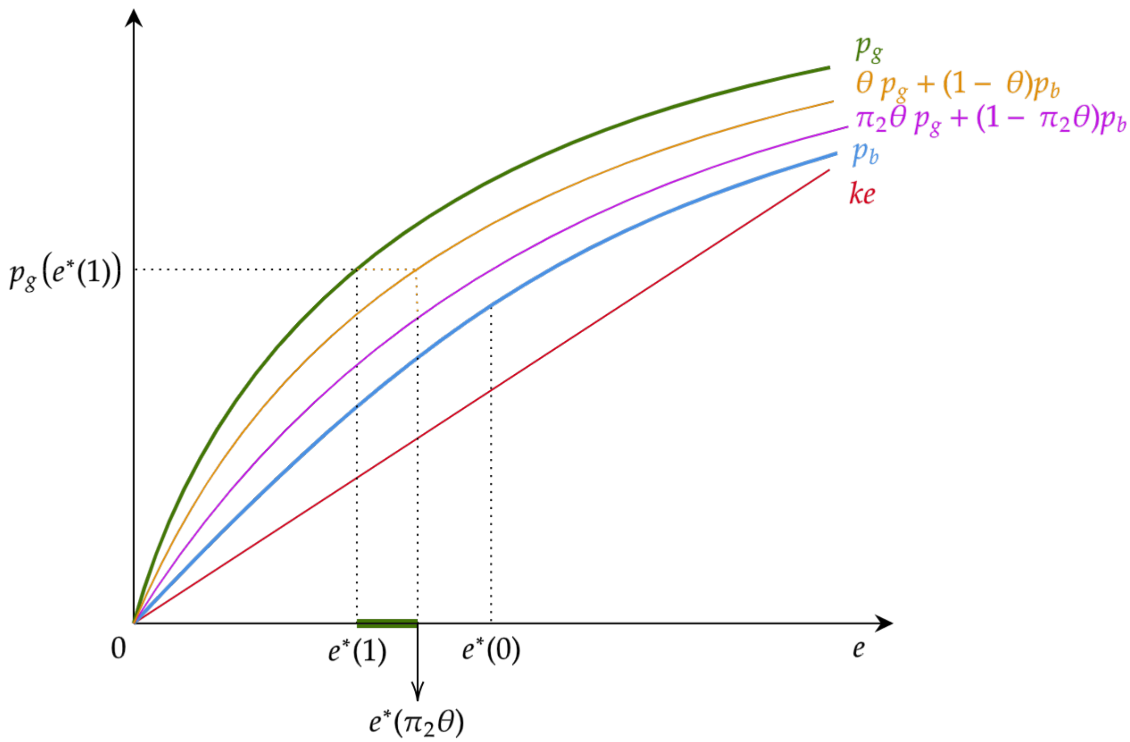

Figure 5: Honesty when idea and effort are substitutes and $g(e^*(1)) \geq b(e^*(0))$

where the condition holds with equality at the self-confidence level π_2^{truth} with $\pi_1^{\text{truth}} = \phi_1^{-1}(\pi_2^{\text{truth}})$. The RHS of condition (C2) captures the supervisor's benefit from calling a good idea a bad one. Doing so gets the agent to experiment again without feedback. The benefit of such discouragement is that the supervisor has learnt that the agent is high-ability, so he expects the agent to generate a good idea with a higher probability. In addition, the agent exerts a higher effort. However, the cost is that the agent may produce a bad idea which reduces the probability of success.

To understand if condition (C2) holds, two subcases must be considered. Honest feedback is possible only when despite $e^*(1) < e^*(0)$, $g(e^*(1)) \geq b(e^*(0))$. The reason is because the inequality implies $g(e^*(1)) > b(e^*(\pi_2\theta))$ and a threshold π_1^{truth} exists. However, when $g(e^*(1)) < b(e^*(0))$, $\theta\, g(e^*(\pi_2\theta)) + (1-\theta)\, b(e^*(\pi_2\theta)) > \theta\, g(e^*(\theta)) + (1-\theta)\, b(e^*(\theta)) > g(e^*(1))$ and there cannot be any honest feedback.

Proposition 3 *Consider a self-confidence level $\pi_1 \geq N_1$ and let idea and effort be* *substitutes. When $g(e^*(1)) \geq b(e^*(0))$, the supervisor gives honest feedback for all levels of self-confidence*

1. *above N_1 if either $\theta \leq \frac{g(e^*(1))-b(e^*(0))}{g(e^*(0))-b(e^*(0))}$ or if $\theta > \frac{g(e^*(1))-b(e^*(0))}{g(e^*(0))-b(e^*(0))}$ and $\pi_1^{truth} \leq N_1$;*

2. *above π_1^{truth} if $\theta > \frac{g(e^*(1))-b(e^*(0))}{g(e^*(0))-b(e^*(0))}$ and $\pi_1^{truth} > N_1$.*

When $g(e^*(1)) < b(e^*(0))$, the supervisor does not provide honest feedback for any level of self-confidence.

4.3 Changing project difficulty

We interpret θ, the probability of producing a good idea, as the project difficulty. A lower θ means that the project is more difficult. We determine here the effect of making the project more difficult on honest feedback and experimentation. At a first pass, doing so should reduce experimentation as both N_0 and F_0 increase. However, there is a countervailing force that influences the equilibrium honesty and agent's experimentation: self-confidence.

For simplicity, suppose the agent and the supervisor hold the same confidence π_1. Whether the supervisor provides honest feedback is contingent on $e^*(\pi_2\theta)$, the agent's effort upon discouragement. However, negative feedback is less informative about ability when the project is more difficult; the decline in π_2 is smaller when θ is lower. As a result, $\theta\,\pi_2(\theta;\pi_1)$ is non-monotonic in θ.

Consider first the case of effort increasing in belief about idea (complements). The supervisor's incentives to be honest are increasing in $\pi_2\theta$, which may be decreasing in θ. Thus, there is a possibility of more rounds of honest feedback and more experimentation with a smaller θ. The reason is that with a smaller θ, agent's effort after experimentation without supervision does not decline much. In the case of effort decreasing in belief (subsititutes), the incentives to be honest are again increasing in $\pi_2\theta$. So, under a more difficult project, if $\pi_2\theta$ is higher, $e^*(\pi_2\theta)$ is lower and that pushes the supervisor to honestly reveal a good idea.

Proposition 4 *Suppose the project is made more difficult by reducing θ to θ'. Let N_1 and N_1' be the respective no-information thresholds.*

1. *When idea and effort are complements, let $\overline{\lambda} < \frac{1+\overline{\lambda}}{2} \leq \theta' < \theta$ such that $N_1 < N_1' < \frac{\overline{\lambda}}{\theta(1+\overline{\lambda}-\theta)}$. Then the supervisor provides honest feedback for lower self-confidence levels when the project is more difficult.*

2. *When idea and effort are substitutes, let $\max\left\{\frac{g(e^*(1))-b(e^*(0))}{g(e^*(0))-b(e^*(0))}, \frac{1}{1+\sqrt{1-\pi_1^{truth}}}\right\} < \theta' < \theta$ such that $N_1 < N_1' < \pi_1^{truth}$. Then the supervisor provides honest feedback for lower self-confidence levels when the project is more difficult.*

4.4 Welfare effect of "overconfidence"

We are interested here in determining how the ex-ante expected utility of the agent behaves when her ex-ante self-confidence level differs from the interim self-confidence. To fix ideas, suppose the agent has an *ex-ante* self-confidence of π_1 before joining the organization. However, before the project starts, her self-confidence level is $\widetilde{\pi}_1 \neq \pi_1$.

Definition 2 *An agent is overconfident (relative to her ex-ante self-confidence) if $\widetilde{\pi}_1 > \pi_1$, i.e., if her interim self-confidence exceeds her ex-ante levels.*

We want to determine whether and when the agent would ex-ante prefer to be overconfident in the interim. To answer this question, we evaluate the ex-ante expected utility, $W(\widetilde{\pi}_1; \pi_1)$, of the agent when she makes decisions according to her interim self-confidence, $\widetilde{\pi}_1$. With some abuse of notation, denote $W(\pi_1; \pi_1)$ by $W(\pi_1)$.

In similar contexts where effort is contingent on the prior, there would be no reason for the welfare to increase under an incorrect belief, i.e., $W(\widetilde{\pi}_1; \pi_1) > W(\pi_1)$. Since this problem is akin to choosing a prior, one would not find it beneficial to choose anything different from their actual (or correct) belief.[15] Yet, we show that this is not always the case in our setting.

Proposition 5 *Suppose the cost of experimentation is at least* $\theta^2(1-N_0)(g(e^*(N_0\theta)) - b(e^*(N_0\theta)))$. *Then there exists a threshold* π^{over} *on the ex-ante self-confidence such that for all levels* $\pi \geq \pi^{over}$, *the agent would be strictly better off under a higher interim level that allows her at least one extra round of experimentation with supervisor feedback. Moreover,* $F_0 < \pi^{over} < N_0$.

The intuition follows from the nature of the equilibrium. There is a discontinuous change in the supervisor's equilibrium strategy at N_1 – he starts being honest. Thus, there exists an agent who at her current ex-ante self-confidence level slightly below N_1 would not receive any feedback but will do so after incorrectly having it bumped up to N_1. The supervisor's honest feedback makes it worthwhile for the agent to be overconfident. A simple corollary follows resulting from the magnifying effect of higher self-confidence.

Corollary 2 *An agent with a higher ex-ante self-confidence weakly prefers larger increases in the interim , i.e., she prefers to be "more" overconfident.*

5 Learning-by-doing with more feedback

One reason to have such agent-supervisor relationships is to permit agent's learning when she is still early in her career. Our model naturally extends to an environment in which an agent learns to understand the quality of her ideas from supervisor feedback. Suppose the agent receives a private signal s_r of the quality of her idea after every experimentation. This signal could either correctly reflect her idea, i.e., $s_r = q_r$, or it may be empty, $s_r = \varnothing$. Let $\psi_r = \Pr(s_r = q_r)$. To capture learning, assume that at

[15] Consider the problem of $\max_e \pi e - \frac{e^2}{2}$ for a belief π. The optimal e is given by $e^* = \pi$, where π is the correct prior belief. Making belief π' a choice variable the maximum would occur at $\pi' = \pi$.

the beginning $\psi_1 \in [0, 1)$, and *only* conditional on receiving honest feedback in $r-1$, $\psi_r > \psi_{r-1}$, and otherwise $\psi_r = \psi_{r-1}$.[16] Naturally, there exists a $r = R$ where $\psi_R = 1$ with repeated honest feedback. Until round R, it is always beneficial to consult with the supervisor if the agent's own signal is empty and the supervisor provides honest feedback.

As before, start with the agent's decision-making problem for a fixed learning level $\psi \geq 0$ without supervision. The optimal decision can again be summarized by a threshold on self-confidence. Call this threshold P_0^{ψ}, where we use the letter P generically for *partial* information threshold. The threshold specifies the minimum self-confidence level below which the agent does not experiment further with a bad idea for a fixed ψ. It can further be shown that $P_0^{\psi} < P_0^{\psi'}$ for $\psi > \psi'$, i.e., the threshold P_0^{ψ} is decreasing in ψ.[17]

Once we have P_0^{ψ}, define $\phi_j^{-1}(P_0^{\pi}) := P_j^{\pi}$ for $j = 1, 2, \ldots$ and for each π as in Definition 1. P_j^{ψ} is the starting belief which when (correctly) updated about bad ideas j times leads to the terminal belief P_0^{ψ}.

We can now present how Proposition 1 is altered when there is learning-by-doing.

Proposition 6 *The supervisor does not provide honest feedback in the region of low self-confidence,* $F_0 \leq \pi_1 < P_1^{\psi_2}$.

While the main intuition is the same as in Proposition 1, there are two points of departure. First, this proposition only deals with the ex-ante self-confidence level, not any generic level. The supervisor's equilibrium behavior differs under the situation when the agent starts with a self-confidence level in $[F_0, P_1^{\psi_2})$ from when it ends up in this region through some past experimentation. With more experimentation and

[16] We assume there is no learning on the job, in terms of an increased π for the next round, if the agent received her own accurate signal. The agent may understand that her idea is bad, but to learn she needs the supervisor to tell her why her idea is bad. To keep the model tractable, we further assume that the agent's signals are public, and the supervisor only provides feedback when she has not received her own signal.

[17] The proofs have been omitted for brevity and can be found in Appendix B.

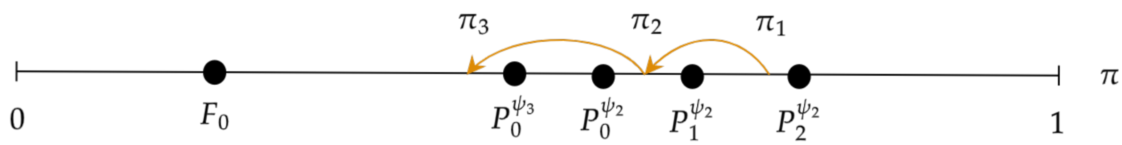

Figure 6: Truth-telling with learning-by-doing when idea and effort are complements for self-confidence levels $P_1^{\psi_2} \leq \pi_1 < P_2^{\psi_2}$

feedback the agent's self-confidence threshold for further experimentation declines. As a result, supervisor's incentives to provide honest feedback to the agent can also change with more experimentation.

For the same reason, the upper bound of the low self-confidence region is given by $P_1^{\psi_2}$. The supervisor incorporates the effect of one round of feedback and the learning it induces as ψ goes from ψ_1 to ψ_2. If self-confidence level π_2 falls below $P_0^{\psi_2}$, the agent does not experiment further in the second round. In this case, the supervisor has the same incentives as described in Proposition 1.

Next, we identify a sufficient conditions to get supervisor honesty in the high self-confidence region, i.e., above $P_1^{\psi_2}$.

Idea and effort are complements. There are two effects of negative feedback. While the agent is less willing to experiment on account of a lower self-confidence, she is also now *more* willing to experiment due to learning-by-doing. More honest feedback increases ψ and reduces the threshold P_0^{ψ}. The supervisor incorporates the effect of reduced P_0^{ψ} when providing feedback. In the relevant worst case incentive constraint, the supervisor is honest one last time in round one. This situation arises when $P_1^{\psi_2} \leq \pi_1 < P_2^{\psi_2}$ and $\pi_3 < P_0^{\psi_3}$ so that in the first round the aforementioned positive effect dominates the negative effect and not thereafter. Negative feedback pushes the agent to experiment once again in round two after which she stops. See Figure 6. Thus, the supervisor is honest if

$$\begin{aligned}&\psi_2\left[\pi_2\theta\, g(e^*(1)) + (1-\pi_2\theta)\, b(e^*(0))\right]\\&\quad + (1-\psi_2)\left[\pi_2\theta\, g(e^*(\pi_2\theta)) + (1-\pi_2\theta)\, b(e^*(\pi_2\theta))\right] \geq b(e^*(1)).\end{aligned} \tag{C3}$$

We identify below the sufficient condition for (C3) to hold true for any self-confidence level above $P_1^{\psi_2}$, and therefore, guarantees honest equilibria. For this purpose, let

$$\bar{\bar{\lambda}}\, g(e^*(1)) + (1-\bar{\bar{\lambda}})\, b(e^*(0)) = b(e^*(1)). \tag{C3*}$$

We also already have $\bar{\lambda}$ from condition (C1*).

Observation 2 *Fix a self-confidence level $\pi_1 \geq P_1^{\psi_2}$ and suppose idea and effort are complements.*

1. *If $P_0^{\psi_2}\theta \geq \max\{\bar{\lambda}, \bar{\bar{\lambda}}\}$, the supervisor provides honest feedback for all self-confidence levels above $P_1^{\psi_2}$.*

2. *If $\bar{\lambda} > P_0^{\psi_2}\theta \geq \bar{\bar{\lambda}}$, there exists a threshold on the level of learning, $\underline{\psi}_2$, such that if $\psi_2 \geq \underline{\psi}_2$, the supervisor provides honest feedback for all self-confidence levels above $P_1^{\psi_2}$.*

3. *If $\bar{\bar{\lambda}} > P_0^{\psi_2}\theta \geq \bar{\lambda}$, there exists a threshold on the level of learning, $\bar{\psi}_2$, such that if $\psi_2 \leq \bar{\psi}_2$, the supervisor provides honest feedback for all self-confidence levels above $P_1^{\psi_2}$.*

4. *If $\min\{\bar{\lambda}, \bar{\bar{\lambda}}\} > P_0^{\psi_2}\theta$, there is no level of learning for which the supervisor gives honest feedback for any self-confidence level.*

Idea and effort are substitutes. The supervisor now has an incentive to call a good idea bad. But incorrect feedback does not help the agent learn for the next period. Suppose the agent expects the supervisor to be honest and experiments as a best response. The worst case is again when she experiments one last time in round two. The supervisor truthfully reveals a good idea if

$$\begin{aligned}&\psi\left[\theta\, g(e^*(1)) \,+\, (1-\theta)\, b(e^*(0))\right]\\ &\quad +\,(1-\psi)\left[\theta\, g(e^*(\pi_2\theta)) \,+\, (1-\theta)\, b(e^*(\pi_2\theta))\right] \,\leq\, g(e^*(1)).\end{aligned} \tag{C4}$$

Unlike in the previous case, $g(e^*(1)) > \theta\, g(e^*(1)) + (1-\theta)\, b(e^*(0))$. Thus, if condition (C2) holds at the self-confidence level $P_0^{\psi_2}$, then the supervisor gives honest feedback.

Observation 3 *Consider a self-confidence level $\pi_1 \geq P_1^{\psi_2}$ and suppose idea and effort are substitutes. The supervisor provides honest feedback for self-confidence levels above $P_1^{\psi_2}$ if either $\theta \leq \frac{g(e^*(1))-b(e^*(0))}{g(e^*(0))-b(e^*(0))}$ or $\theta > \frac{g(e^*(1))-b(e^*(0))}{g(e^*(0))-b(e^*(0))}$ and $\pi^{truth} < P_1^{\psi_2}$.*

Observe how honesty is possible for $F_0 \leq \pi < P_1^{\psi_2}$ when the prior self-confidence level is above $P_1^{\psi_2}$ but not when it is between F_0 and $P_1^{\psi_2}$. In fact, the agent with self-confidence just above $P_1^{\psi_2}$ may get honest feedback and experiment all the way to F_0, while an agent just below does not.[18] Thus, an agent with a self-confidence slightly above $P_1^{\psi_2}$ gets a boost in her performance due to repeated feedback, while not the one below the threshold.

6 Implications and conclusion

6.1 Alternate interpretations

It is possible to reinterpret our model to describe settings where a less informed receiver (agent) seeks feedback from an informed sender (supervisor) after a costly effort to decide her future course of action. Consider, for instance, an entrepreneur who works on a project experimenting with ideas, *privately* observing their quality and implementing one of them. However, he relies on a venture capitalist (VC) to finance such experimentation and implementation of the project. The VC, in turn, makes

[18]For instance, this may happen when $\pi_j \geq P_0^{\psi_j}$ from $j = 3$ onward.

decisions based on the entrepreneur's recommendations. Preferences are such that the entrepreneur would prefer to continue experimenting until he receives a good idea; the VC would like to cut funding for experimentation when she is sufficiently pessimistic. In such a setting, the entrepreneur is the supervisor, while the VC is the agent.

Implication 1 *A VC-entrepreneur relationship can be "inefficient" because even though the VC may like to continue financing the entrepreneur's experimentation, she calls off the project too early.*

Still another way to interpret our model is in the context of developmental feedback. Employees in organizations often seek input from their supervisors on improving their current performance and building their skills before applying for promotion. Suppose an employee works every period to develop her skills getting feedback from her supervisor. Building adequate skill is akin to getting a good idea. She may choose to apply for a promotion at any time, at which time she conducts a particular task. Success in this task depends on her skill level and effort.

Implication 2 *Less confident women will get fewer opportunities to grow in the organizational hierarchy due to less developmental feedback.*

6.2 Affirmative action and other policies

A common policy prescription to reduce gender differences in STEM fields is to have quotas for early-career women scientists. Whether the different stakeholders believe such quotas are in place or not can influence a potential scientist's self-confidence when starting her career. Dr. Julia Omotade, a cellular neuroscientist and a woman of color, recalled her experience from Ph.D.[19],

[19]This quote is taken from an opinion piece that appeared in ASBMB Today magazine. See Omotade, King, and Kahn (2017).

> "Upon entering graduate school, I received a merit-based fellowship that the institution used to attract top applicants. Importantly, this fellowship is not associated with any diversity initiative. When several colleagues became aware of this fellowship, they asked how much "*diversity money*" I was receiving. Thus, instead of an accomplishment, this fellowship instantly was transformed in my mind into an automatically generated handout based on statistics or an attempt to meet a diversity quota."

Social and organizational psychology literature highlights that opponents of affirmative action policies (AAP) may view them as giving preferential treatment to certain groups. As a result, it delegitimizes the achievements of the targeted groups. Evidence suggests that nonbeneficiaries may react adversely to the mere knowledge of such policies being practiced in their organizations (Heilman, McCullough, and Gilbert (1996), Kidder, Lankau, Chrobot-Mason, Mollica, and Friedman (2004)). Within academia, forcing admissions, grants and tenure committees to make decisions based on ethnicity, gender and race may lead to adverse perceptions among potential supervisors of early-career scientists. Moreover, if the benefiting groups internalize such criticism, it can lower their self-confidence (Nacoste (1989), Heilman (1994), Hattrup (1998)). For instance, when women believe their organization practices AAP, it may lower their desire to take up leadership positions (Islam and Zilenovsky (2011)).[20]

The knowledge of being admitted under a quota can lead to differences between ex-ante and interim levels of self-confidence. For example, if the agent believes she has been admitted under AAP, her interim self-confidence may decline. On the flip side, an agent may experience a boost in her interim self-confidence if she is admitted to a high-ranking school that is not known to practice AAP.

[20] These studies do not find outright negative relations between AAP and the relevant variables. For instance, Hattrup (1998) finds that less confident women were more likely to attribute their performance and hiring for the job to preferential selection than to merit-based selection. Similarly, Kidder et al. (2004) finds that framing AAP as diversity initiatives that do not induce a sense of loss for the nonbeneficiary groups reduces backlash.

Implication 3 *Early-career diversity students in academia and STEM fields will perform better if these hires are made without publicizing affirmative action policies among the students and supervisors.*

Another organizational policy to increase gender diversity is to have more gender-balanced decision-making committees. One reason to believe this exercise will benefit women is that they can better understand the signals generated by other women (Lang (1986), Bjerk (2008), Flabbi, Macis, Moro, and Schivardi (2019)). For instance, in the presence of gender segregation among scientific subfields, it may be useful to have women evaluators for women candidates. Women evaluators, in turn, will be better able to identify more confident and more able women candidates.

Implication 4 *Early-career diversity scientists' performance will improve in universities and departments with more diverse supervisors and decision-makers.*

There is already some evidence supporting the above prediction. Gaule and Piacentini (2018) find that (chemistry) students with advisors of the same gender tend to be more productive during their Ph.D. and are also more likely to become professors themselves. Hale and Regev (2014) and Boustan and Langan (2019) similarly find positive effects of a greater share of female faculty on both admissions and completion rates of female Ph.D. candidates among U.S. economics departments. We posit that one reason for observing such results is that the female faculty members make better decisions regarding women candidates and support them with more critical feedback.

6.3 Long-run statistical discrimination

The model showed that the agent with a lower self-confidence or supervisor confidence is more likely to fail. With this result, we can identify one route from taste-based discrimination to statistical discrimination. Consider the situation where two agents, a man and a woman, with same self-confidence levels face a supervisor who has a lower confidence in women because of his biases. The difference in supervisor confidence will induce a performance difference between the two agents.So, even without intent or explicitly targeting the woman, the supervisor can cause women to fail more often. This greater likelihood of failures solidifies beliefs about the lower ability and fuels statistical discrimination. In turn, managers in organizations are less likely to allocate high-stake or career-promoting tasks to minorities, hurting their promotion and career trajectories.

Implication 5 *Taste-based discrimination or other supervisor biases against minorities can turn into long-run statistical discrimination in organizations.*

6.4 Conclusion

In this paper, we have shown 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 or to those employees in whose ability they do not believe. In turn, this hurts their performance and potentially their future careers. Moreover, it also hurts organizations as supervisors provide too little critical feedback. Our model highlights the importance of confidence-building exercises for young creative professionals and academics. Our findings suggest that to improve the long-term professional outcomes of women and other minorities organizations should begin by investing in their confidence-building.

References

Ali, O. (2017). Disclosure in multistage projects. *Working paper*.

Aoyagi, M. (2010). Information feedback in a dynamic tournament. *Games and Economic Behavior*, *70*(2), 242-260.

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.

Bénabou, R., & Tirole, J. (2003). Intrinsic and extrinsic motivation. *The Review of Economic Studies*, *70*(3), 489–520.

Bjerk, D. (2008). Glass ceilings or sticky floors? statistical discrimination in a dynamic model of hiring and promotion. *The Economic Journal*, *118*(530), 961–982.

Boleslavsky, R., & Lewis, T. R. (2016). Evolving influence: Mitigating extreme conflicts of interest in advisory relationships. *Games and Economic Behavior*, *98*, 110–134.

Boustan, L., & Langan, A. (2019). Variation in women's success across phd programs in economics. *Journal of Economic Perspectives*, *33*(1), 23–42.

Chakraborty, A., & Harbaugh, R. (2010). Persuasion by cheap talk. *The American Economic Review*, *100*(5), 2361–2382.

Chow, Y., Robbins, H., & Siegmund, D. (1971). Great expectations: The theory of optimal stopping. In (pp. 52–55). Houghton Mifflin.

Crawford, V. P., & Sobel, J. (1982, November). Strategic information transmission. *Econometrica: Journal of the Econometric Society*, 1431–1451.

Ederer, F. (2010). Feedback and motivation in dynamic tournaments. *Journal of Economics & Management Strategy*, *19*(3), 733-769.

Ely, J. C. (2017). Beeps. *The American Economic Review*, *107*(1), 31–53.

Ely, J. C., & Szydlowski, M. (2020). Moving the goalposts. *Journal of Political Economy*, *128*(2), 468–506.

Ferguson, T. S. (2006). *Optimal stopping and applications.*

Flabbi, L., Macis, M., Moro, A., & Schivardi, F. (2019). Do female executives make a difference? the impact of female leadership on gender gaps and firm performance. *The Economic Journal*, *129*(622), 2390–2423.

Fudenberg, D., & Rayo, L. (2019, November). Training and effort dynamics in apprenticeship. *American Economic Review*, *109*(11), 3780-3812.

Gaule, P., & Piacentini, M. (2018). An advisor like me? advisor gender and post-graduate careers in science. *Research Policy*, *47*(4), 805–813.

Gneezy, U., & Rustichini, A. (2004). Gender and competition at a young age. *American Economic Review*, *94*, 377–581.

Golosov, M., Skreta, V., Tsyvinski, A., & Wilson, A. (2014). Dynamic strategic information transmission. *Journal of Economic Theory*, *151*, 304–341.

Goodman, I., Cunningham, C. M., & Lachapelle, C. (2002). Womens experiences in college engineering. *A Comprehensive Evaluation of Women in Engineering Programs*.

Grossman, S. J. (1981). The informational role of warranties and private disclosure about product quality. *The Journal of Law and Economics*, *24*(3), 461-483.

Gnther, C., Arslan Ekinci, N., Schwierenc, C., & Strobeld, M. (2010). Women cant jump? an experiment on competitive attitudes and stereotype threat. *Journal of Economic Behavior and Organization*, *75*(3), 3959–401.

Halac, M., Kartik, N., & Liu, Q. (2017). Contests for experimentation. *Journal of Political Economy*, *125*(5), 1523–1569.

Hale, G., & Regev, T. (2014). Gender ratios at top phd programs in economics. *Economics of Education Review*, *41*, 55–70.

Hattrup, K. (1998). The role of self-perceptions in reactions to preferential and merit-based hiring 1. *Journal of Applied Social Psychology*, *28*(3), 225–234.

Heilman, M. (1994). Affirmative action: Some unintended consequences for working women. In *Research in organizational behavior* (pp. 125–169). JAI Press.

Heilman, M., McCullough, W., & Gilbert, D. (1996). The other side of affirmative action: Reactions of nonbeneficiaries to sex-based preferential selection. *Journal of Applied Psychology*, *81*(4), 346.

Hirshleifer, D., Low, A., & Teoh, S. H. (2012). Are overconfident ceos better innovators?

The Journal of Finance, *67*(4), 1457–1498.

Islam, G., & Zilenovsky, S. E. (2011). Affirmative action and leadership attitudes in brazilian women managers: The moderating influence of justice perceptions. *Journal of Personnel Psychology*, *10*(3), 139.

Jampol, L., & Zayas, V. (2021). Gendered white lies: Women are given inflated performance feedback compared with men. *Personality and Social Psychology Bulletin*, *47*(1), 57–69.

Kamenica, E., & Gentzkow, M. (2011). Bayesian persuasion. *The American Economic Review*, *101*(6), 2590–2615.

Kidder, D. L., Lankau, M. J., Chrobot-Mason, D., Mollica, K. A., & Friedman, R. A. (2004). Backlash toward diversity initiatives: Examining the impact of diversity program justification, personal and group outcomes. *International Journal of Conflict Management*.

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.

Koszegi, B. (2006). Ego utility, overconfidence, and task choice. *Journal of the European Economic Association*, *4*(4), 673-707.

Lang, K. (1986). A language theory of discrimination. *The quarterly journal of economics*, *101*(2), 363–382.

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.

Milgrom, P. R. (1981). Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics*, *12*(2), 380–391.

Nacoste, R. W. (1989). Affirmative action and self-evaluation. In *Affirmative action in perspective* (pp. 103–109). Springer.

Niederle, M., & Vesterlund, L. (2007). Do women shy away from competition? do men compete too much? *Quarterly Journal of Economics*, *122*, 1067–10101.

Omotade, J., King, J., & Kahn, R. A. (2017). Imposter syndrome and diversity graduate students. *ASBMB Today*, *16*(2). Retrieved from https://biomed.emory.edu/PROGRAM_SITES/BCDB/_includes/documents/sections/new-events/bcdb_imposter_syndrome.pdf

Orlov, D., Skrzypacz, A., & Zryumov, P. (2020). Persuading the principal to wait. *Journal of Political Economy*, *128*(7), 2542–2578.

Perry, M., & Gershkov, A. (2009). Tournaments with midterm reviews. *Games and Economic Behavior*, *66*(1), 162-190.

Shurchkov, O. (2012). Under pressure: gender differences in output quality and quantity under competition and time constraints. *Journal of the European Economic Association*, *10*(5), 1189–1213.

Smolin, A. (2021). Dynamic evaluation design. *American Economic Journal: Microeconomics*, *13*(4), 300–331.

Appendices

Mathematical notation for mixed strategies

A strategy for the agent ρ_r in round r is a mapping from the last observed message to a possible mixed decision to continue experimenting with ideas or implementing the last one. Let $\rho_r^{m_{r-1}} = \Pr(d_r = \text{implement} \mid m_{r-1})$ be the probability that the agent decides to implement the project following the last message.

Similarly, a strategy for the supervisor, σ_r in round r, is a mapping from the last idea to a possible mixed message about its quality. Let $\sigma_r^{q_r} = \Pr(m_r = g \mid q_r)$ be the probability of the supervisor that he calls the last idea good. Depending on the expected strategy of the supervisor, the agent conditions her action only on the last message received.

Let the sequence $\hat{\boldsymbol{\sigma}} = \{(\hat{\sigma}_r^g, \hat{\sigma}_r^b)\}_{r=1}^R$ denote the conjectured strategy of the supervisor, and let $\hat{\boldsymbol{\rho}} = \{(\hat{\rho}_r^g, \hat{\rho}_r^b\}_{r=1}^R$ 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 quality of the last idea produced. The belief about being high ability is $\pi_r^{m_{r-1}}$. Let the belief about her idea when the supervisor's message is m_r be denoted by $\lambda_r^{m_r}$. The public history, the one observed by the agent, h_r^A at the beginning of round r can be summarized by the current public beliefs $\pi_r^{m_{r-1}}$ and $\lambda_r^{m_r}$. The private history of the supervisor h_r^S at the beginning of round r can be summarized by the current private belief Π_r and q_r.

We can now informally describe the notion of equilibrium. We say that a pair of sequences of conjectured strategies $\boldsymbol{\sigma}$ and $\boldsymbol{\rho}$ constitute an equilibrium if (1) they are both the best responses to each other given the beliefs π_r, λ_r and Π_r for each r, and (2) the beliefs π_r, λ_r and Π_r are consistent with what the players are conjectured to do, i.e., $\boldsymbol{\sigma}$ and $\boldsymbol{\rho}$. Strategies expressed in the text without a hat constitute an equilibrium.

The supervisor is expected to babble in equilibrium in round r if $\hat{\sigma}_r^g = \hat{\sigma}_r^b$. Whenever the supervisor babbles, it might be useful to think of babbling in mixed strategies rather than in pure strategies. When the supervisor is expected to be informative, assume WLOG that $\hat{\sigma}_r^g > \hat{\sigma}_r^b$. Therefore, the agent updates belief about idea being good as

$$\lambda_r^b = \frac{(1-\hat{\sigma}_r^g)\pi_r\theta}{(1-\hat{\sigma}_r^g)\pi_r\theta + (1-\hat{\sigma}_r^b)(1-\pi_r\theta)} < \pi_r\theta \ , \ \lambda_r^g = \frac{\hat{\sigma}_r^g\pi_r\theta}{\hat{\sigma}_r^g\pi_r\theta + \hat{\sigma}_r^b(1-\pi_r\theta)} > \pi_r\theta, \quad (1)$$

and her self-confidence gets updated as

$$\pi_{r+1}^b = 1.\lambda_r^b + \frac{(1-\theta)\pi_r}{1-\pi_r\theta}.(1-\lambda_r^b) = \frac{(1-\hat{\sigma}_r^g)\pi_r\theta + (1-\hat{\sigma}_r^b)(1-\theta)\pi_r}{(1-\hat{\sigma}_r^g)\pi_r\theta + (1-\hat{\sigma}_r^b)(1-\pi_r\theta)} < \pi_r, \quad (2)$$

$$\pi_{r+1}^g = 1.\lambda_r^g + \frac{(1-\theta)\pi_r}{1-\pi_r\theta}.(1-\lambda_r^g) = \frac{\hat{\sigma}_r^g\pi_r\theta + \hat{\sigma}_r^b(1-\theta)\pi_r}{\hat{\sigma}_r^g\pi_r\theta + \hat{\sigma}_r^b(1-\pi_r\theta)} > \pi_r. \quad (3)$$

A Proofs from main text

Proof of Lemma 1

Proof. The proof of the first part follows from applying the Implicit Function Theorem on the first-order condition presented in the main text. Doing so establishes $\frac{de^*}{d\lambda} = \frac{-(g'-b')}{\lambda g'' + (1-\lambda)b''}$. As $\lambda g'' + (1-\lambda)b'' < 0$, whether $g' \lesseqgtr b'$ determines the sign of $\frac{de^*}{d\lambda}$. The proof of the second part follows from applying the Envelope Theorem on the first-order condition. The shape of the V function requires taking the second-order derivative and utilizing the expression for $\frac{de^*}{d\lambda}$. ∎

Proof of Proposition 1

Proof. We prove this statement in steps by considering different regions of starting prior π_1. There exists a $j \in \{0, 1, 2, \dots\}$ where belief F_j is such that $F_j < N_0 \leq F_{j+1}$. The value that j takes depends on the parameters.

Step 1: No honest equilibria starting from $F_0 \leq \pi_1 < F_1$

First, we check if there can be honest equilibrium for any self-confidence $F_0 \leq$

$\pi_1 < F_1$. If the agent expects honest feedback, the agent best responds to $m_1 = b$ by experimenting once in round one and then implementing with effort $e^*(0)$ (since $\pi_2 < F_0$, $\lambda_2^b = 1$, Observation 1). A message $m_1 = g$ instead leads the agent to implement with effort $e^*(1)$.

If $q_1 = b$ and idea quality and implementation effort are complements, then $e^*(1) > e^*(0)$ and the supervisor deviates to announcing $m_1 = g$. If instead $q_1 = g$ and idea quality and implementation effort are substitutes, then $e^*(1) < e^*(0)$ and the supervisor deviates to announcing $m_1 = b$. The agent's equilibrium strategy is to implement her outside information idea, i.e. $d_1 =$ implement with $e^*(\pi_1\theta)$ since $\pi_1 < N_0$ (Lemma 2).

Second, we also need to check if starting from any self-confidence $F_0 \leq \pi_1 < F_1$ there can be honest equilibria in the future rounds with mixing at least in the current round. Let the supervisor's mixed strategies for different π levels be fixed in this region, i.e., the supervisor's strategy remains the same every time the agent lands in the region $[F_0, F_1)$.

There cannot be any honesty following $m_1 = b$ in the future because the current mixing is informative and it must reduce the agent's self-confidence. Furthermore, as $m_1 = b$ lowers self-confidence and is followed by more mixing (less than full information), the agent's incentive to experiment decrease. Assuming the supervisor's strategy induces monotonically decreasing self-confidence due to $m = b$, there must be a threshold self-confidence level below which the agent does not experiment (after a bad idea). Call this threshold π^{final} and observe that $\pi^{\text{final}} > F_0$. Suppose the agent arrives at $\pi \leq \pi^{\text{final}}$ in round $r' \geq 1$ and the agent implements her idea. To show that there does not exist any honest equilibria it will be sufficient to show that there is at least one deviation from round r' onward.

Suppose idea and effort are complements.

- If the agent's best response following $m_{r'} = g$ is to implement, the supervisor deviates to always sending $m_{r'} = g$ because $\lambda_{r'}^g > \lambda_{r'}^b$ and $e^*(\lambda_{r'}^g) > e^*(\lambda_{r'}^b)$.

- If the agent's best response following $m_{r'} = g$ is to experiment for certain rounds and then implement after experimenting without feedback, the supervisor deviates to always sending $m_{r'} = g$. The reason is that this event can happen only if $\pi \geq F_1$, where π is the self-confidence when she implements after experimentation without feedback. So, $g(e^*(\pi\theta)) > b(e^*(\pi\theta)) > b(e^*(\lambda^b_{r'}))$ from (1).

- If the agent's best response following $m_{r'} = g$ is to experiment for certain rounds until she ends up again in the region $F_0 \leq \pi < F_1$, then if $\pi \geq \pi^{\text{final}}$, the supervisor deviates to always sending $m_{r'} = g$. The reason is because he can always guarantee himself a payoff of $b(e^*(\lambda^b_{r'}))$.

- If the agent's best response following $m_{r'} = g$ is to experiment for certain rounds until she ends up again in the region $F_0 \leq \pi < F_1$, then if $\pi < \pi^{\text{final}}$ and the supervisor prefers at least one round of experimentation with honest feedback, the supervisor deviates to always sending $m_{r'} = g$. Thus, we require,

$$\Pi^b_{r'}\theta\, g(e^*(1)) + (1 - \Pi^b_{r'}\theta)\, b(e^*(0)) > b(e^*(\lambda(\pi^{\text{final}}))). \tag{4}$$

  If the conjectured supervisor's strategy below π^{final} is $\hat{\sigma}^g = 1$ and any $\hat{\sigma}^b \in [0, 1)$, (4) is satisfied for any $\Pi^b_{r'} > 0$.

Suppose now idea and effort are substitutes. The argument largely proceeds as above with the exception that the supervisor deviates to sending $m_{r'} = b$ for $q_{r'} = g$. When the agent experiments further after $m_{r'} = g$, the supervisor has an incentive to report $m_{r'} = b$ because (1) the supervisor can get the maximum possible effort of $e^*(\lambda^b_{r'})$ on a good idea, and (2) the supervisor prevents a future bad idea.

For the remaining steps, the proof of no honest equilibria in the future proceeds as above has been omitted for brevity.

Step 2: No honest equilibria starting from $F_1 \leq \pi_1 < N_0$

If $j = 0$, then we are already done. If $j = 1$, then it is enough to show that honest equilibria are not possible in the range $F_1 \leq \pi_1 < N_0$. Here, if the posterior $\pi_2 < F_1$, then the supervisor is not honest (from Step 1 above) and that the agent best responds by implementing with effort $e^*(0)$. As before now, the supervisor is better off deviating to induce the agent to implement with a higher effort. Thus, honest equilibria will not survive.

If $j \in \{2, 3, \dots\}$, then it needs to be shown that honesty is not possible for $F_1 \leq \pi_1 < F_2$, ..., $F_{j-1} \leq \pi_1 < F_j$ and $F_j \leq \pi_1 < N_0$. Doing so is immediate using the above logic sequentially starting from $F_1 \leq \pi_1 < F_2$. Again, the agent does not experiment, i.e. $d_1 =$ implement and $e^*(\pi_1 \theta)$.

Step 3: No honest equilibria starting from $N_0 \leq \pi_1 < N_1$

For $j \geq 1$, the reasoning is exactly as in Step 2 for $N_0 \leq \pi_1 < N_1$. For $j = 0$, we have already shown that honest equilibria are not possible for $F_0 \leq \pi_1 < N_0 < F_1$ or $F_0 < N_0 \leq \pi_1 < F_1$. Note that since $F_0 < N_0$, it must be the case that $F_1 < N_1 < F_2$. Using the same argument as above, we can show our result for $F_1 \leq \pi_1 < N_1$. ■

Proof of Proposition 2

Proof. Fix $\pi_1 \geq N_1$ and let idea quality and implementation effort be complements. Then, as outlined in the text, the supervisor provides honest feedback for self-confidence above π_1 if $\Pi \geq \Pi_1^{\text{truth}}$ from condition (C1). From condition (C1*), we have the belief about good idea such that $\bar{\lambda}\, g(e^*(\bar{\lambda})) + (1 - \bar{\lambda})\, b(e^*(\bar{\lambda})) = b(e^*(1))$. Note that such a $\bar{\lambda}$ exists because convex combination of two concave functions is concave and e^* is increasing in λ (see Figure 4).

Part 1: Let $N_0 \theta \geq \bar{\lambda}$, or equivalently, $N_1 \geq \frac{\bar{\lambda}}{\theta(1+\bar{\lambda}-\theta)}$. This implies that for any common prior confidence level $\pi_1 \geq N_1$,

$$\pi_2 \theta\, g(e^*(\pi_2 \theta)) + (1 - \pi_2 \theta)\, b(e^*(\pi_2 \theta)) \geq b(e^*(1)). \tag{5}$$

The reason is that the LHS of equation (5) is increasing in π_2, and so, $N_0\theta \geq \bar{\lambda}$ reflects a sufficient condition for (5) to hold. Rearranging and using condition (C1) for supervisor honesty, we get that $\pi_2 \geq \Pi_2^{\text{truth}} \iff \pi_1 \geq \Pi_1^{\text{truth}}$.

Part 2: Let $N_0\theta < \bar{\lambda}$, or equivalently, $N_1 < \frac{\bar{\lambda}}{\theta(1+\bar{\lambda}-\theta)}$. Now, for $N_1 \leq \pi_1 < \frac{\bar{\lambda}}{\theta(1+\bar{\lambda}-\theta)}$,

$$\pi_2\theta\, g(e^*(\pi_2\theta)) \,+\, (1-\pi_2\theta)\, b(e^*(\pi_2\theta)) \,<\, b(e^*(1)). \tag{6}$$

Now, as above, $\pi_1 < \Pi_1^{\text{truth}}$. And for $\pi_1 \geq \frac{\bar{\lambda}}{\theta(1+\bar{\lambda}-\theta)}$, it is identical to Part 1.

Note that in all the cases, condition (C1) must hold and $\Pi_1 \geq \Pi_1^{\text{truth}}$ to support supervisor's honesty. ■

Proof of Proposition 3

Proof. Fix $\pi_1 \geq N_1$ and let idea quality and implementation effort be substitutes. Then the supervisor provides honest feedback for self-confidence above π_1 if condition (C2) holds. As described in the text, honest equilibria are not possible when $g(e^*(1)) < b(e^*(0))$. So, we focus here on the case of $g(e^*(1)) > b(e^*(0))$.

Now, the RHS of condition (C2) is decreasing in π_2, and thus, it takes the largest value for $\pi_2 = 0$. So, when $g(e^*(1)) \geq \theta g(e^*(0)) + (1-\theta)b(e^*(0)) \iff \theta \leq \frac{g(e^*(1))-b(e^*(0))}{g(e^*(0))-b(e^*(0))}$, the supervisor is honest for all levels of self-confidence above N_1. However, when $\theta > \frac{g(e^*(1))-b(e^*(0))}{g(e^*(0))-b(e^*(0))}$, then the comparison of $\bar{\pi}_2$ ($\bar{\pi}_1$) with N_0 (N_1) will determine whether the supervisor provides honest feedback. ■

Proof of Proposition 4

Proof.

Part 1: Idea quality and implementation effort are complements.

Begin by noting that $\bar{\lambda}$ does not change with a change in θ. $\bar{\lambda}$ is given by condition (C1*). Neither the LHS nor the RHS are determined by θ. Only the functions g and

b, and $k > 0$ determines $\bar{\lambda}$.

Second, we note that a decrease in θ increases N_1 to N_1'. However, $\frac{\bar{\lambda}}{\theta(1+\bar{\lambda}-\theta)}$ may either increase or decrease. Specifically, it is straightforward to verify that $\frac{d}{d\theta}\frac{\bar{\lambda}}{\theta(1+\bar{\lambda}-\theta)} \geq 0$ when $\theta \geq \frac{1+\bar{\lambda}}{2}$.

Third, note that when $N_1 \geq \frac{\bar{\lambda}}{\theta(1+\bar{\lambda}-\theta)}$, and θ decreases, the self-confidence region of honest feedback unambiguously falls because $N_1' > N_1$. However, when $N_1 < \frac{\bar{\lambda}}{\theta(1+\bar{\lambda}-\theta)}$, then the supervisor does not provide honest feedback for $N_1 \leq \pi_1 < \frac{\bar{\lambda}}{\theta(1+\bar{\lambda}-\theta)}$. Now, if $N_1 < N_1' < \frac{\bar{\lambda}}{\theta(1+\bar{\lambda}-\theta)}$, then the self-confidence region of honest feedback may increase. However, for this to happen either $N_1' \geq \frac{\bar{\lambda}}{\theta'(1+\bar{\lambda}-\theta')}$ or $\frac{\bar{\lambda}}{\theta'(1+\bar{\lambda}-\theta')} < \frac{\bar{\lambda}}{\theta(1+\bar{\lambda}-\theta)}$. But from the previous paragraph, we already know that $\frac{\bar{\lambda}}{\theta(1+\bar{\lambda}-\theta)}$ decreases when we reduce θ to θ'.

Part 2: Idea quality and implementation effort are substitutes.

Again, a decrease in θ increases N_1 to N_1'. Starting from condition (C2), there are two effects of reducing θ. First, the RHS (payoff from deviating) reduces on account of $g > b$. Second, however, a decrease in θ has a non-monotonic effect on $\pi_2\theta$. We are interested in when a reduced θ increases $\pi_2\theta$. Then, since effort is decreasing in λ, it would reduce the agent effort in case of deviation. It is easy to check that $\frac{d}{d\theta}\pi_2\theta < 0$ when $\theta > \frac{1}{1+\sqrt{1-\pi_1}}$. Note that $\frac{1}{1+\sqrt{1-\pi_1}}$ is larger for larger values of π_1. A sufficient condition for the RHS to decrease would be $\theta > \theta' > \frac{1}{1+\sqrt{1-\bar{\pi}_1}}$.

Note that as a result $\bar{\pi}_1' < \bar{\pi}_1$. The reason is that for $\pi_1 < \bar{\pi}_1$, $g(e^*(1)) < \theta\, g(e^*(\pi_2\theta)) + (1-\theta)\, b(e^*(\pi_2\theta))$. But with a lower θ', the RHS is smaller for any such π_1 under $\theta > \theta' > \frac{1}{1+\sqrt{1-\bar{\pi}_1}}$. ■

Proof of Proposition 5

Proof. To prove the statement, first we consider different ranges of ex-ante beliefs and determine the sufficient condition that makes exactly one additional round of experimentation with feedback welfare-improving.

Case 1: $F_0 \leq \pi_1 < N_0$.

First, note that an increase to $F_0 \leq \tilde{\pi}_1 < N_1$ cannot be welfare-improving. Here, $W(\pi_1) = V(0)$. When $F_0 \leq \tilde{\pi}_1 < N_0$, then $W(\tilde{\pi}_1; \pi_1) = W(\pi_1) = V(0)$. When $N_0 \leq \tilde{\pi}_1 < N_1$, then

$$W(\tilde{\pi}_1; \pi_1) \;=\; -c \,+\, \pi_1\theta\, g(e^*(\tilde{\pi}_1\theta)) \,+\, (1-\pi_1\theta)\, b(e^*(\tilde{\pi}_1\theta)) \,-\, k\, e^*(\tilde{\pi}_1\theta). \tag{7}$$

However, at $\lambda = \pi_1\theta$, the maximized value upon experimentation without information is $V(\pi_1\theta) - c$. So, $W(\tilde{\pi}_1; \pi_1)$ from (7) is lesser than $V(\pi_1\theta) - c$. Further, since $F_0 \leq \pi_1 < N_0$, it must be $V(\pi_1\theta) - c < V(0)$ from Lemma 2.

Second, consider an increase to $N_1 \leq \tilde{\pi}_1 < N_2$. In this case,

$$W(\tilde{\pi}_1; \pi_1) \;=\; -c \,+\, \pi_1\theta\, V(1) \,+\, (1-\pi_1\theta)\, W(\tilde{\pi}_2; \pi_2), \tag{8}$$

where $W(\tilde{\pi}_2; \pi_2) = -c + \pi_2\theta\, g(e^*(\tilde{\pi}_2\theta)) + (1-\pi_2\theta)\, b(e^*(\tilde{\pi}_2\theta)) - k\, e^*(\tilde{\pi}_2\theta)$ and $\tilde{\pi}_2$ follows from Bayesian updating $\tilde{\pi}_1$.

Note that at $\pi_1 > F_0$, $-c + \pi_1\theta\, V(1) + (1-\pi_1\theta)\, V(0) > V(0)$ from Lemma 3. But since $\pi_1 < N_0$, $W(\tilde{\pi}_2; \pi_2) < W(\pi_2) = V(\pi_2\theta) - c$ from Lemma 2. Thus, there may exist a threshold confidence level, π^{over}, such that for $\pi_1 > \pi^{\text{over}}$, $W(\tilde{\pi}_1; \pi_1) > V(0) = W(\pi_1)$.

To check for the existence of π^{over}, begin by observing that $W(\tilde{\pi}_1; \pi_1)$ is decreasing in $\tilde{\pi}_1$. Therefore, assign the minimal increase in self-confidence that will lead to experimentation with feedback, i.e., let $\tilde{\pi}_1 = N_1$. If the agent does not find it beneficial to be overconfident up to this level, then so would she not for any higher level. Thus, we want to check if

$$W(N_1; \pi_1) \;=\; -c \,+\, \pi_1\theta\, V(1) \,+\, (1-\pi_1\theta)\, W(N_0; \pi_2) > V(0), \tag{9}$$

where $W(N_0; \pi_2) = -c \,+\, \pi_2\theta\, g(e^*(N_0\theta)) \,+\, (1-\pi_2\theta)\, b(e^*(N_0\theta)) \,-\, k\, e^*(N_0\theta)$.

Observe that $W(N_0; \pi_2) < V(0) = -c + N_0\theta\, g(e^*(N_0\theta)) + (1 - N_0\theta)\, b(e^*(N_0\theta)) - k\, e^*(N_0\theta)$. So, we can rewrite $W(N_0; \pi_2) = V(0) - h(\pi_1\theta)$ for

$$
\begin{aligned}
h(\pi_1\theta) &= V(0) - [-c + \pi_2\theta\, g(e^*(N_0\theta)) + (1 - \pi_2\theta)\, b(e^*(N_0\theta)) - k\, e^*(N_0\theta)] \qquad (10)\\
&= \theta\, (N_0 - \pi_2)\, [g(e^*(N_0\theta)) - b(e^*(N_0\theta))]\\
&= \frac{\theta}{1 - \pi_1\theta}\, (N_0 - \pi_1(1 - \theta + N_0\theta))\, [g(e^*(N_0\theta)) - b(e^*(N_0\theta))].
\end{aligned}
$$

Thus, we want to show $-c + \pi_1\theta\, V(1) + (1 - \pi_1\theta)\, (V(0) - h(\pi_1\theta)) > V(0)$ where $h(\pi_1\theta)$ is given by (10). Substituting and simplifying gives

$$
\begin{aligned}
&- N_0\theta\, [g(e^*(N_0\theta)) - b(e^*(N_0\theta))] +\\
&\qquad + \pi_1\theta\, [V(1) - V(0) + \{g(e^*(N_0\theta)) - b(e^*(N_0\theta))\}(1 - \theta(1 - N_0))] > c \qquad (11)
\end{aligned}
$$

The LHS of (11) is linear in π_1. It takes values $-N_0\theta\, [g(e^*(N_0\theta)) - b(e^*(N_0\theta))]$ and $\theta\, [V(1) - V(0) + \{g(e^*(N_0\theta)) - b(e^*(N_0\theta))\}(1 - \theta)(1 - N_0)]$ for $\pi_1 = 0$ and $\pi_1 = 1$ respectively. It is straightforward to compare the LHS of (11) with that of condition (C-FI) to conclude that π^{over} exists and is above F_0 (see Figure). However, we need to confirm if it is below N_0.

To do so, we compare the slopes of $V(\pi_1\theta) - V(0)$ (from condition (C-NI) with the slope of LHS of equation (11) at $\pi_1 = N_0$. If the former is larger than the latter, then $\pi^{\text{over}} < N_0$. Note that $\frac{d}{d\pi} = \frac{dV(\lambda)}{d\lambda} . \frac{d\lambda}{d\pi} = [g(e^*(N_0\theta)) - b(e^*(N_0\theta))]\, \theta$, and the slope of the LHS of (11) $= \theta\, [V(1) - V(0) + \{g(e^*(N_0\theta)) - b(e^*(N_0\theta))\}(1 - \theta(1 - N_0))]$. So, we need

$$
\begin{aligned}
&\theta\, [V(1) - V(0) + \{g(e^*(N_0\theta)) - b(e^*(N_0\theta))\}(1 - \theta(1 - N_0))]\\
&\qquad\qquad > [g(e^*(N_0\theta)) - b(e^*(N_0\theta))]\, \theta,
\end{aligned}
$$

which simplifies to

$$V(1) - V(0) > \theta\,(1 - N_0)\,[g(e^*(N_0\theta)) - b(e^*(N_0\theta))]. \tag{12}$$

Note that we already know $V(1) - V(0) > c/\theta$. Thus, if $c/\theta \geq \theta\,(1 - N_0)\,[g(e^*(N_0\theta)) - b(e^*(N_0\theta))]$, then $\pi^{\text{over}} < N_0$. Thus, if $c \geq \theta^2(1 - N_0)[g(e^*(N_0\theta)) - b(e^*(N_0\theta))]$ there exists a threshold $F_0 < \pi^{\text{over}} < N_0$ such that the agent ex-ante prefers to hold an interim self-confidence of N_1 for all $\pi \geq \pi^{\text{over}}$.

Case 2: $N_0 \leq \pi_1 < N_1$.

In this case, $W(\pi_1) = V(\pi_1\theta) - c > V(0)$. Again, an increase to $N_0 \leq \tilde{\pi}_1 < N_1$ cannot be welfare-improving. For the case of increase in self-confidence to $N_1 \leq \tilde{\pi}_1 < N_2$, $W(\tilde{\pi}_1;\pi_1)$ is as in (8). As before, set $\tilde{\pi}_1 = N_1$ and $W(\tilde{\pi}_2;\pi_2) = V(0) - h(\pi_1\theta)$. After making the relevant substitutions from the previous case, we identify the condition that makes overconfidence preferable as

$$\begin{aligned} &- N_0\theta\,[g(e^*(N_0\theta)) - b(e^*(N_0\theta))] + \\ &\qquad + \pi_1\theta\,[V(1) - V(0) + \{g(e^*(N_0\theta)) - b(e^*(N_0\theta))\}(1 - \theta(1 - N_0))] \\ &\qquad\qquad > V(\pi_1\theta) - V(0). \end{aligned} \tag{13}$$

Observe that the LHS (13) is identical to that of (11). However, now we compare it with $V(\pi_1\theta) - V(0)$. Clearly, the relevant belief threshold now is below the previously determined threshold π^{over}.

Case 3: $N_j \leq \pi_1 < N_{j+1}$ for $j \geq 1$.

Again, we compare $W(\tilde{\pi}_1;\pi_1)$ with $W(\pi_1)$ for $\tilde{\pi}_1 = N_{j+1}$ so that $\tilde{\pi}_{j+2} = N_0$. Observe that for the first j tries, the overconfidence and the non-overconfidence situations are identical. However, in the $j+1$th attempt, if the agent is overconfident, she gets a lower payoff than if she were not (since $\pi_{j+2} < N_0$ while $\tilde{\pi}_{j+2} = N_0$). So, in this last round if

$$
\begin{aligned}
& -c + \pi_{j+1}\theta\, V(1) + \\
& \quad + (1-\pi_{j+1}\theta)\,(-c + \pi_{j+2}\theta\, g(e^*(N_0\theta)) + (1-\pi_{j+2}\theta)\, b(e^*(N_0\theta)) - k\, e^*(N_0\theta)) \\
& \qquad > V(\pi_{j+1}\theta) - c, \quad (14)
\end{aligned}
$$

then we are done. But as $N_0 \leq \pi_{j+1} < N_1$, condition (14) is identical to the one proved in the previous case.

To complete the proof, note that $W(\tilde{\pi}_1; \pi_1)$ is increasing in π_1. Thus, an agent with a higher ex-ante belief gains more from an increase in the interim belief that allows her more rounds of experimentation with feedback. Thus, a higher ex-ante belief makes it more beneficial to have a higher interim belief as well. ■

Proof of Observation 2

Proof. We are looking for the sufficient condition that makes it incentive compatible to provide honest feedback for all self-confidence levels above $P_1^{\psi_2}$. So, if condition (C3) is true for $\pi_2 = P_0^{\psi_2}$, then its true for all self-confidence levels above $P_1^{\psi_2}$.

Now, note that if $\lambda > (<) \bar{\lambda}$, then $\bar{\lambda}\, g(e^*(\bar{\lambda})) + (1-\bar{\lambda})\, b(e^*(\bar{\lambda})) > (<)\, b(e^*(1))$ from condition (C1*). Similarly, if $\lambda > (<) \bar{\bar{\lambda}}$, then $\bar{\bar{\lambda}}\, g(e^*(1)) + (1-\bar{\bar{\lambda}})\, b(e^*(0)) > (<)\, b(e^*(1))$ from condition (C3*). The four conclusions now immediately follow. It can be shown that

$$
\begin{aligned}
\underline{\psi}_2 &= \frac{b(e^*(1)) - P_0^{\psi_2}\theta g(e^*(P_0^{\psi_2}\theta)) - (1-P_0^{\psi_2}\theta) b(e^*(P_0^{\psi_2}\theta))}{P_0^{\psi_2}\theta[g(e^*(1)) - g(e^*(P_0^{\psi_2}\theta))] + (1-P_0^{\psi_2}\theta)[b(e^*(0)) - b(e^*(P_0^{\psi_2}\theta))]} \\
\overline{\psi}_2 &= \frac{P_0^{\psi_2}\theta g(e^*(P_0^{\psi_2}\theta)) + (1-P_0^{\psi_2}\theta) b(e^*(P_0^{\psi_2}\theta)) - b(e^*(1))}{P_0^{\psi_2}\theta[g(e^*(P_0^{\psi_2}\theta)) - g(e^*(1))] + (1-P_0^{\psi_2}\theta)[b(e^*(P_0^{\psi_2}\theta)) - b(e^*(0))]}.
\end{aligned}
$$

■

B Agent's experimentation decision

B.1 Base model as in Section 3

No information about the idea. If the agent has no information about ideas, the agent learns nothing. She neither learns the quality, nor does she update on her ability. She may, however, still want to experiment once as a gamble if her self-confidence is sufficiently high. Here, $\lambda = \pi\theta$. So, the agent experiments if

$$-c + V(\pi\theta) \geq V(0). \tag{C-NI}$$

Lemma 2 *Suppose the agent does not receive information about ideas.*

1. *If experimentation is not too costly with $c < V(\theta) - V(0)$, there exists a unique threshold on self-confidence N_0 that solves (C-NI) with an equality such that if the self-confidence is greater than (or equal to) N_0 then the agent experiments once and implements her idea by exerting effort $e^*(\pi_1\theta)$.*

2. *If the self-confidence is lower than N_0 or if experimentation is costly with $c \geq V(\theta) - V(0)$, the agent does not experiment and implements her outside option idea with effort $e^*(0)$.*

Full information about ideas. When the agent learns about the idea quality, her self-confidence also moves with this knowledge. Using Bayes' rule, it declines to $\frac{(1-\theta)\pi_{r-1}}{1-\pi_{r-1}\theta}$ if she observes a bad idea, and it moves to 1 with a good idea. The agent may now benefit from repeatedly experimenting if ideas are bad. In the process, she learns about their quality and own ability, and implement the preferred one. Note that λ is either 0 or 1 in this case, respectively for bad and good ideas.

Assuming that the agent wants to start experimenting, we want to determine when the agent stops experimenting with repeated bad ideas. Using the one-step look-ahead rule, if the agent finds it optimal to stop experimenting at a given self-confidence level π, then she also finds it optimal to stop after another round of experimentation. So, the agent experiments for levels that satisfy

$$-c + \pi\theta V(1) + (1-\pi\theta) V(0) \geq V(0) \iff \pi \geq F_0 := \frac{c}{\theta[V(1)-V(0)]}. \tag{C-FI}$$

Lemma 3 *Suppose the agent has full information about the idea quality.*

1. *If experimentation is not too costly with $c < \theta[V(1)-V(0)]$, there exists a unique threshold $F_0 := \frac{c}{\theta[V(1)-V(0)]}$ such that if the agent's previous idea was bad, she experiments again for any self-confidence $\pi_r \geq F_0$, and if the agent's previous idea was good, the agent implements her idea in the following round with an effort $e^*(1)$. However, she implements a bad idea with an effort of $e^*(0)$ for any self-confidence $\pi_r < F_0$.*

2. *If experimentation is costly with $c \geq \theta[V(1)-V(0)]$, she does not experiment for any belief and implements her outside option idea with effort $e^*(0)$.*

Proof.

Part 1: Existence and uniqueness of F_0 using one-step look-ahead rule

Denote by $\mathcal{V}^b(\pi_r)$ the value function of the agent at the beginning of round r with belief π_r when her last observed outcome is $q_{r-1} = b$. Then,

$$\mathcal{V}^b(\pi_r) = \max\left\{V(0)\ ,\ -c + \pi_r\theta V(1) + (1-\pi_r\theta)\mathcal{V}^b(\pi_{r+1})\right\}.$$

This Bellman equation reflects that an agent (1) who comes up with a good idea does not experiment any further, but (2) who comes up with a bad idea faces the same decision problem as she faced originally with a lower belief.

Using the one-step look-ahead rule, the agent experiments as long as condition (C-FI) is satisfied, which gives the threshold F_0. Note that this threshold exists, i.e., is a number between 0 and 1 if $c < \theta[V(1) - V(0)]$. Further, it is unique because $\pi\theta[V(1) - V(0)]$ linearly increases in π, and so it can cross c only once.

Part 2: Optimality of one-step look-ahead decision rule

If we show that our optimal stopping problem is monotone, Corollary of Theorem 3.3 (Chow, Robbins, and Siegmund (1971, p. 54)) readily establishes the optimality of the one-step look-ahead rule.[21] The problem is monotone if whenever the one-step look-ahead rule calls the agent to implement in round R, then so does it for all future rounds no matter what ideas are generated. Let

$$-c + \pi_R\theta V(1) + (1 - \pi_R\theta)V(0) < V(0). \tag{15}$$

Given our discussion in Part 1, $\pi_R < F_0$. We want to show that equation (15) is also true for a $\pi < \pi_R < F_0$. This is also immediate from the discussion and Figure 7. In addition, note that the benefit is bounded above by $g(e^*(1))$ and the cost of experimentation is fixed at c. Thus, it is optimal to implement the project with $q = b$ for $\pi < F_0$ and continue experimenting otherwise. ■

Comparing the no-information and full-information cases. We first compare the belief thresholds induced by the two policies.

Lemma 4 *Let $c < V(\theta) - V(0) < \theta[V(1) - V(0)]$. Both N_0 and F_0 exist with $N_0 > F_0$.*

Proof. Fix the parameters such that $c < V(\theta) - V(0)$. Since, $V(\theta) - V(0) < \theta[V(1) - V(0)]$, both N_0 and F_0 exist and are unique. Comparing the LHS of (C-NI) and (C-FI) establishes $F_0 < N_0$ as V is convex in λ so that $\pi\theta V(1) + (1 - \pi\theta)V(0) > V(\pi\theta.1 + (1 - \pi\theta).0) = V(\pi\theta)$. ■

[21] See also Ferguson (2006) for a description of the sufficient conditions when we have a maximization problem.

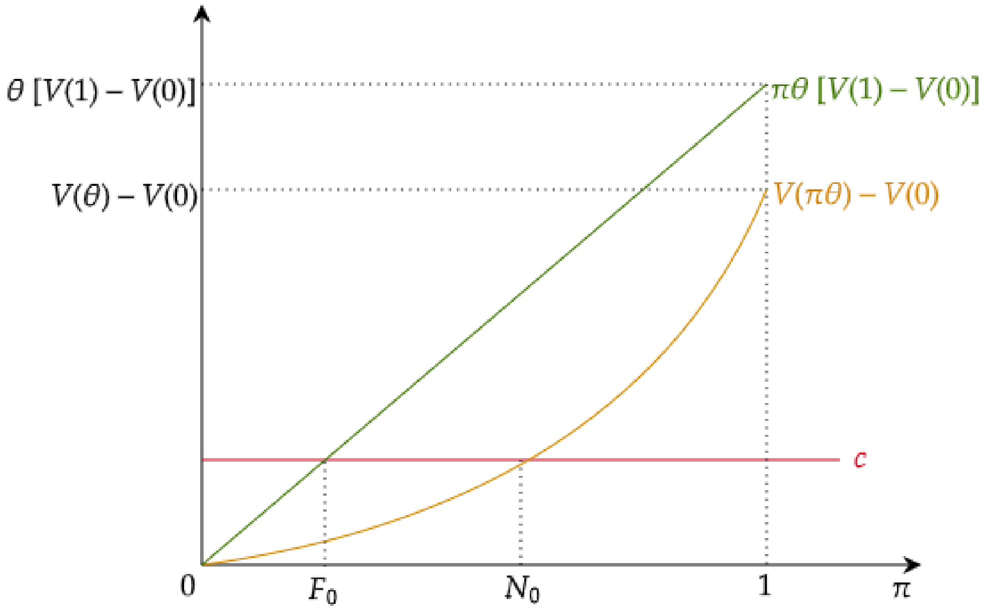

Figure 7: Comparing N_0 and F_0

Figure 7 illustrates why $N_0 > F_0$. It shows that for any belief the additional value of a final round of experimentation is lower under no information due to the absence of learning. As a result, for some region of self-confidence the agent does not experiment when there is no information, even though she would if she had full information. Accordingly, the cost condition associated with the no information policy, $c < V(\theta) - V(0)$, binds. Unless otherwise stated, we assume that this cost condition holds for the rest of the analysis.

Next, we compare the ex-ante expected utility of the agent in the two cases. As the agent pays for the additional cost of experimentation under full information, it is not immediately obvious if she would prefer it. However, our first proposition shows that this is so.

Proposition 7 *For any ex-ante self-confidence level $\pi_1 \geq F_0$, the agent (strictly) prefers to be fully informed about her ideas than having no information.*

Proof. Let $W^F(\pi_1)$ denote the ex-ante expected utility of the agent under the full-information case when her prior self-confidence is π_1. Suppose $F_j \leq \pi_1 < F_{j+1}$, then

$$W^F(\pi_1) = -c + \pi_1\theta \sum_{r=0}^{j} (1-\theta)^r V(1) + \left(1 - \pi_1\theta \sum_{r=0}^{j} (1-\theta)^r\right) V(0) - c\left[(1-\pi_1)j + \pi_1 \sum_{r=1}^{j} (1-\theta)^r\right] \quad (16)$$

For $F_0 \leq \pi_1 < N_0$, we know that $-c+\pi_1\theta V(1)+(1-\pi_1\theta)V(0) > V(0) > V(\pi_1\theta)-c$, and for $F_0 < N_0 \leq \pi_1$, $-c+\pi_1\theta V(1)+(1-\pi_1\theta)V(0) > V(\pi_1\theta)-c \geq V(0)$. Since both N_0 and F_0 exist, $-c + \pi_1\theta\, V(1) + (1-\pi_1\theta)\, V(0) > \max\{V(\pi_1\theta)-c, V(0)\}$. Therefore, it will be sufficient to show that $W^F(\pi_1) \geq -c + \pi_1\theta\, V(1) + (1-\pi_1\theta)\, V(0)$.

Proof by induction. The above is trivially true for $j = 0$. For $j = 1$, we can show that

$$-c + \pi_1\theta\, V(1) + (1-\pi_1\theta)\,[-c + \pi_2\theta\, V(1) + (1-\pi_2\theta)\, V(0)] > -c + \pi_1\theta\, V(1) + (1-\pi_1\theta)\, V(0),$$

which is true because $\pi_2 \geq F_0$ and from condition (C-FI), $-c+\pi_2\theta V(1)+(1-\pi_2\theta)V(0) > V(0)$. Now, suppose the statement is true for some $j = k-1$. We need to show that the statement is also true for $j = k$. To do so, we show that $W^F(\pi_1)$ under the latter case is larger than the former. The expression for the two case are easily obtained by replacing j with $k-1$ and k in equation (16) respectively. Taking the difference and simplifying, we get the required condition as

$$\pi_2\theta\,[V(1) - V(0)] > c \quad (17)$$

which is true since $\pi_2 > F_0$ and equation (17) holds from Lemma 4. ■

The reason is that higher self-confidence (particularly above F_0) increases the odds of generating a good idea. It makes the agent willing to experiment and pay the cost under full information policy. Such experimentation would not be possible when there is no information about ideas despite the agent believing that she can achieve success. Consequently, the agent is worse off when the supervisor does not provide information on ideas. We, therefore, seek to determine if and when the supervisor can engage in such beneficial truthful communication without commitment.

B.2 Learning-by-doing as in Section 5

Using the one-step look-ahead rule, it is easy to characterize the condition that yields P_0^{ψ}. Accordingly, the agent experiments with bad ideas if

$$-c + \psi\left(\pi\theta V(1) + (1-\pi\theta)V(0)\right) + (1-\psi)V(\pi\theta) \geq V(0). \tag{C-LbD}$$

Lemma 5 *Let $c < V(\theta) - V(0)$. There exists a unique threshold $F_0 \leq P_0^{\psi} \leq N_0$ associated with every $0 \leq \psi \leq 1$ which solves condition (C-LbD) with equality such that*

1. *if the agent's previous idea was bad, she experiments for any self-confidence $\pi_r \geq P_0^{\psi}$, but implements it with effort $e^*(0)$ for any belief $\pi_r < P_0^{\psi}$, and*
2. *if the agent's previous idea was good, she implements her idea in the following round with effort $e^*(1)$.*

Further, $P_0^{\psi} < P_0^{\psi'}$ for $\psi > \psi'$, i.e., the threshold P_0^{ψ} is decreasing in ψ.

Proof. Denote by $\mathcal{W}^b(\pi_r;\psi)$ (resp. $\mathcal{W}^{\varnothing}(\pi_r;\psi)$) the value function of the agent at the beginning of round r with belief π_r when her last observed outcome is $q_{r-1} = b$ (resp. $q_{r-1} = \varnothing$) and she learns her idea with probability ψ.

$$\mathcal{W}^b(\pi_r) = \max\left\{V(0), -c + \psi\left(\pi_r\theta V(1) + (1-\pi_r\theta)\mathcal{W}^b(\pi_{r+1})\right) + (1-\psi)\mathcal{W}^{\varnothing}(\pi_r)\right\}$$
$$\mathcal{W}^{\varnothing}(\pi_r) = \max\left\{V(\pi_r\theta), -c + \psi\left(\pi_r\theta V(1) + (1-\pi_r\theta)\mathcal{W}^b(\pi_{r+1})\right) + (1-\psi)\mathcal{W}^{\varnothing}(\pi_r)\right\}$$

where we suppress ψ in writing $\mathcal{W}(.)$ as it does not change with an agent working by herself.

If the agent experiments with an empty signal,

$$\begin{aligned} \mathcal{W}^{\varnothing}(\pi_r) &= -c + \psi\left(\pi_r\theta\, V(1) + (1-\pi_r\theta)\, \mathcal{W}^{b}(\pi_{r+1})\right) + (1-\psi)\mathcal{W}^{\varnothing}(\pi_r) \\ \Longleftrightarrow \mathcal{W}^{\varnothing}(\pi_r) &= -\frac{c}{\psi} + \pi_r\theta\, V(1) + (1-\pi_r\theta)\, \mathcal{W}^{b}(\pi_{r+1}), \end{aligned} \tag{18}$$

and as a result,

$$V(\pi_r\theta) < -\frac{c}{\psi} + \pi_r\theta\, V(1) + (1-\pi_r\theta)\, \mathcal{W}^{b}(\pi_{r+1}). \tag{19}$$

However, at this point $V(0) \leq V(\pi_r\theta)$ and so, if the agent experiments with an empty signal, then she does so with a bad idea as well.

Now, the proof of the first part is the same as that of Lemma 3. Note that in the application of the OSLA rule we will use the above – if the agent does not experiment with a bad idea, she does not do so with an empty signal as well. To show that P_0^{ψ} is decreasing in ψ note that the RHS of condition (C-LbD) is a convex combination of the RHS of conditions (C-NI) and (C-FI). ψ is the weight on the RHS of condition (C-FI). Naturally, as ψ increases, the belief threshold P_0^{ψ} moves closer to F_0. See Figure. ■