Interim Performance Evaluation of Career-Concerned Agents*

Moritz Janas,

Department of Economics, University of Gothenburg; moritz.janas@gu.se

November 24, 2026

Abstract

The timing of performance feedback is a central managerial decision. This paper studies interim performance evaluations (IPEs) in a two-period signaling model in which career-concerned agents privately know their ability. Revealing first-period performance allows agents to condition second-period effort on the first-period outcome. This flexibility creates two opposing forces: agents may stop exerting effort after an early success, while the option to continue after a failure can make first-period effort more attractive. The net effect depends on the cost of effort. Below a cost threshold, total effort is higher without IPE; above it, total effort is higher with IPE. I test the model at a cost above this threshold in a laboratory experiment with two task environments: an abstract task with randomly assigned ability and a quiz task in which ability reflects prior knowledge. The experiment also varies whether IPE is assigned exogenously or chosen by the principal. Total effort is significantly higher under IPE in both task environments, and the effect is similar under both assignment modes. The predicted drop in second-period effort after an early success is small and insignificant in the abstract task and reversed in the quiz task. Consequently, IPE raises total effort by more than the model predicts.

Keywords: Interim Performance Evaluation, Career Concerns, Signaling, Feedback, Effort Provision, Laboratory Experiment

JEL: C73, C91, D82, D83, M12, D91.

*I would like to thank Friedrich Breyer, Sebastian Fehrler, Urs Fischbacher, Susanne Goldlücke, Jeff Zeidel, Nick Zubanov, the research groups at the Thurgau Economic Institute, CBID, the University of Gothenburg, and participants at several seminars, workshops, and conferences for valuable comments. I gratefully acknowledge financial support from the Doctoral Fund at the University of Konstanz and from the Social Science Economics Laboratory at NYU Abu Dhabi.

1 Introduction

Successful management of organizations requires a precise understanding of employees' incentives. One of the primary tasks and goals of leadership is to align employees' efforts with the company's interests. From the employee's perspective, multiple incentives are at play in the workplace, one of which is career concerns (Holmström, 1999). These concerns can be modeled within a signaling game framework, in which the desire to signal one's privately known talent influences the willingness to exert effort (Spence, 1978). The prerequisites for such a signaling game include private information about the agent's competence, successful task execution as a function of ability and effort, and a reputation for high competence as part of the agent's utility function. In equilibrium, such a signaling game can lead to partial separation of more talented from less talented employees, benefiting the company through higher levels of effort provision. Information can influence such equilibria, and one of the tools available to managers is control over the flow of information within the company. This project focuses on the implications of one specific tool that influences information flow: an Interim Performance Evaluation (IPE).

The literature on IPEs typically assumes that the principal is better informed about the quality of the agent's work than the agent himself. Furthermore, most existing studies assume initial symmetric information about the agent's ability (Fuchs, 2015; Suvorov and Van de Ven, 2009; Zhang, 2019; Chen, 2015; Ederer, 2010; Gürtler and Harbring, 2010; Hansen, 2013), or do not account for different ability types (MacLeod, 2003; Fuchs, 2007; Maestri, 2012; Lang, 2019; Goltsman and Mukherjee, 2011; Lizzeri et al., 2002). In this paper, I also assume that the manager is better at judging the quality of the employee's work, but I deviate from the existing literature by assuming that employees have superior knowledge about their own talent.[1] For example, when an employee moves to a new position, the manager might be informed about the performance evaluations the employee received in a previous position or read about the employee's ability in a recommendation letter from a former supervisor. The employee, however, is more aware of whether the favorable evaluations were a result of ability or diligence at work. Similarly, the employee likely has a better understanding of the accuracy of the recommendation letter, such as whether the former supervisor was lenient in writing it. In the new work environment, the experienced manager may still be better informed about the quality of the employee's work, while the employee remains better informed about his own level of ability. The impact of an IPE in a career-concern model

[1] An exception in the IPE literature is Chen and Stephen Chiu (2013), who also assume that agents are privately informed about their own talent. However, their paper does not feature career concerns and assumes a different connection between first- and second-period output.

with these assumptions has not yet been explored.

To model such a situation, I explore a game that consists of two periods and involves two risk-neutral players: a principal (she) and an agent (he). The agent has a privately known type that determines his competence. In each period, the agent can choose whether to exert high effort at an exogenous cost. For a given effort provision, the probability of achieving success increases with the agent's type.[2] The principal's utility function positively depends on the successful production of high-quality output in each period. The principal observes the output quality for each period and updates her beliefs about the agent's type. As a shortcut for modeling career concerns, these beliefs (minus the potential costs of effort) determine the agent's payoff. Absent an IPE, the agent does not observe the output quality of the first period before making his second-period effort decision. It is publicly known whether the principal has implemented an Interim Performance Evaluation protocol in her company. An IPE informs the agent about the output quality in the first period.[3] Hence, the main assumptions of the model are: 1) a career-concerned agent, 2) a principal capable of Bayesian updating, 3) costly, unobservable effort provision, 4) an agent who is privately informed about his talent, and 5) a principal who is (without IPE) privately informed about the quality of the agent's output.

The model is solved using Perfect Bayesian equilibria, restricting attention to the equilibria with maximal signaling activity. The key mechanism is a conditional strategy that IPE enables: agents exert effort in the first period and only continue to exert effort in the second period if the first period did not result in success. Under NoIPE, agents cannot observe the first-period outcome before their second-period decision, so this conditioning is not possible. The conditional strategy creates two opposing forces. On the one hand, agents who succeed early stop exerting effort, because their first-period success has already sufficiently distinguished them from lower-ability agents and removed the signaling incentive for a second attempt. On the other hand, the option to try again only after a failure makes first-period effort more attractive, drawing in types who would not exert effort at all under NoIPE. Which force dominates depends on the cost of effort: below a threshold cost, NoIPE generates more total effort; above it, IPE does.

[2]Success can be understood as, for example, fulfilling a client's needs or producing a high-quality good.

[3]Although much literature addresses subjective performance evaluations, this study focuses solely on objective evaluations. While I acknowledge that the principal might have an incentive to misreport the agent's performance, this possibility is excluded by design, as it is not the focus of this project. Objective evaluations can be motivated by factors such as a lying-averse principal, a long-lived principal who wants to maintain a reputation for honesty, or the outsourcing the IPE to a third party to ensure credibility, as is common (see, e.g., Klaas et al., 2001).

To test the validity of the assumptions and causally identify behavioral implications potentially relevant to the game, I conduct a controlled laboratory experiment. In this experiment, I vary the type of task performed, the existence of an IPE, and its endogeneity. Agents can choose whether to exert monetarily costly effort in each of two periods to succeed in producing high-quality output. The chance of success depends on the effort exerted and the agent's type. The agents' monetary incentive is based on the belief about his type held by another subject who observes whether success occurred in a given period. In the experiment, the implementation of an IPE is varied, as it is either determined exogenously by a randomization device or chosen endogenously by the principal. This design allows for examining the role of reciprocity in the setup: a principal who chooses (not) to provide information that the agent might be interested in could influence effort provision through positive (negative) reciprocity.

To assess the effect of intrinsic motivation, half the subjects play in an *Abstract* context, and the other half play in a *Quiz* context. In the Abstract context, a type is randomly assigned to each agent, which determines the probability of success in a given period (given that the exogenous effort cost is paid). Intrinsic motivation is expected to be relatively low in this treatment. In the Quiz context, however, the agents' types are not induced but are based on their prior knowledge. The task consists of answering questions about the capitals of randomly selected countries. If an agent chooses to pay the exogenous cost and correctly answers at least 7 out of 10 questions, success is achieved. While the Abstract context closely tests the model, the Quiz context enhances external validity by using the subjects' true knowledge distribution as their competence type. It also potentially introduces elements of intrinsic motivation, curiosity about personal performance, and a non-monetary desire to demonstrate knowledge to anonymous others.

The experimental test of the model reveals that an IPE indeed increases effort provision. However, the predicted conditional strategy of exerting effort in the first period and exerting effort again only after a failure finds little support in the data. In the Abstract context, a first-period success lowers second-period effort only slightly, and the reduction is not statistically significant. In the Quiz context, participants adopt the opposite strategy: exert effort particularly when success was achieved in the first period. Thus, IPE increases agents' overall effort provision. The endogeneity of the IPE is varied by allowing either the principal or the computer to decide on its implementation. Assuming that agents are curious about their performance and exhibit reciprocal behavior, one would expect lower effort provision if the principal decides against the IPE. However, the data do not support the presence of this reciprocal mechanism.

The remainder of the paper is organized as follows. Section 2 reviews the related literature. Section 3 formalizes the model, and Section 4 characterizes the Perfect Bayesian Equilibria of the game with and without IPE. Section 5 describes the experimental design and develops testable hypotheses. Section 6 presents the experimental results and explores potential explanations for the deviating findings. Section 7 discusses the practical implications of the findings for managers and concludes.

2 Related Literature

The project builds on different strands of the literature: the consequences of performance evaluations, multi-period signaling games, and preferences for information.

Interim Performance Evaluation There is a vast and growing theoretical literature on the implementation and design of interim performance evaluations. More broadly, IPEs can be viewed as one particular form of organizational transparency: firms choose how information about employees' actions and outcomes flows within the organization, ranging from visibility to peers and managers to dashboards and formal interim evaluations. These information policies can discipline effort through peer monitoring and social pressure (Mas and Moretti, 2009) or by giving workers direct access to performance analytics that were previously available only to supervisors (Bernstein and Li, 2026). At the same time, greater transparency may backfire if it induces concealment that undermines learning and improvement (Bernstein, 2012), or if disclosure emphasizes easily monitored actions and thereby amplifies conformism driven by career concerns (Prat, 2005). Viewed through this lens, the principal's decision to reveal or conceal interim performance information is a choice of information structure (Kamenica and Gentzkow, 2011), for which credibility and skepticism become central when information is provided strategically (Milgrom and Roberts, 1986). There are two main strands within the theoretical IPE literature. The first strand focuses on subjective, non-verifiable evaluations, while the second, like the current paper, focuses on objective, verifiable information. With subjective IPEs, the honesty and credibility of the feedback provided by the principal become the focus of attention. In this literature, the main message is that bonus wages (e.g. MacLeod, 2003; Fuchs, 2007, 2015; Suvorov and Van de Ven, 2009; Maestri, 2012; Zhang, 2019), the ability to fire agents (Chen, 2015), or costly justification (Lang, 2019) are needed to provide a credible IPE. The other strand of the literature focuses on objective performance evaluations. The papers on objective IPEs differ in the specific assumptions about the game. Lizzeri et al. (2002) study the closest setting: a two-period moral-hazard model in which the agent learns his first-period outcome

only if the principal reveals it. Under a fixed incentive scheme, revelation raises total effort, yet a principal who designs the scheme optimally prefers not to reveal. The current paper replaces the explicit scheme with career concerns, where rewards are posterior beliefs about the agent's privately known ability, and IPE lowers total effort when effort is cheap. Chen and Stephen Chiu (2013) assumes that one final good is produced, and the IPE can increase the agent's motivation to exert effort in the second period if performance so far has been good, or decrease it if performance so far has been bad. However, this trade-off cannot occur in situations where two unrelated products are produced. IPEs in competitive environments, such as tournaments, create trade-offs in which agents who learn that they are more likely to be competent exert more effort than others (Ederer, 2010), or in which agents who learn that they are close to winning exert more effort than those who are leading or lagging far behind (Gürtler and Harbring, 2010; Goltsman and Mukherjee, 2011). A related strand of research examines optimal feedback provision in dynamic contests, identifying when concealing rather than disclosing interim performance information generates higher effort (Aoyagi, 2010; Ely et al., 2023). In these contest models, however, effort incentives arise from competition between agents, whereas the current paper studies the choice between concealment and revelation in a single-agent career-concern setting.

In career-concern models with IPE, if effort and talent are substitutes in the production function, a ratchet effect arises in which agents want principals to believe they did not exert effort, as this would increase the principal's belief about the agent's talent (Hansen, 2013). Under a complementary production function instead (as in Dewatripont et al., 1999; Suvorov and Van de Ven, 2009; Di Pei, 2016; Smolin, 2021), effort is crucial for success, and the probability of success increases with talent. Under complementary production functions, this ratchet effect cannot arise.

The assumption in the current paper that the principal is privately, or at least better, informed about the quality of the agent's production is the starting point of almost all studies on IPEs, while the assumption that an agent is better informed about his own ability is common to the signaling literature (see Riley (2001) for an overview of the earlier literature). What appears to be missing from the current literature is the combination of these two assumptions with the possibility that a principal provides objective performance evaluations. By assuming symmetric information about the agent's ability, agents' equilibrium strategies do not differ in the first period of a signaling game. The current paper therefore provides a rationale for observing different effort levels in the first period that does not solely arise from mixed-strategy play.

Performance feedback and effort A large experimental literature studies how performance feedback affects effort provision, with effects that vary across contexts and depend on whether the feedback is favorable or unfavorable. Kuhnen and Tymula (2012) find that the prospect of learning one's rank raises effort, but that agents who learn that they performed better than expected subsequently reduce effort, while those who learn that they performed worse increase it. Field and natural experiments on relative performance feedback document similarly heterogeneous effects (e.g., Azmat and Iriberri, 2010; Azmat et al., 2019; Eriksson et al., 2009); see Villeval (2020) for a survey. Evidence from firms shows that structured conversations about performance can raise performance on their own: in a field experiment, Manthei et al. (2023) find that biweekly reviews in which store managers report their activities to their supervisor increase profits. The authors view such reviews as a monitoring device that creates reputational concerns toward the supervisor, the motive that the present model formalizes as career concerns. Two features distinguish the current paper from this literature. First, most of these studies examine relative performance feedback, where effort responds to one's standing among peers, whereas the current paper studies feedback about an agent's own output in a non-competitive setting. Second, this literature is largely reduced-form, whereas the current paper derives the effort response to feedback from an explicit signaling model, providing a theoretical benchmark against which the experimental results can be assessed. Notably, the reduction in effort after favorable feedback found by Kuhnen and Tymula (2012) parallels the stopping effect predicted by the present model.

Repeated Signaling Career-concerns assumptions have been widely used since the seminal work of Holmström (1999) and Gibbons and Murphy (1992). Because the current paper assumes private information about the agent's talent, its career-concerns model builds on a signaling game (Spence, 1978) in a multiperiod setting. Depending on the players' action spaces, separation of types may or may not be possible. Kaya (2009) studies a repeated version of a signaling game with persistent binary types and observable history. Because off-equilibrium-path beliefs are unrestricted, signaling can occur even after full separation of types. Kaya therefore focuses on the least-cost separating equilibrium, in which the strong type signals just enough in every period, so that the weak type is not willing to mimic it. Such full separation can occur only if the outcome variable is continuous and unbounded. In the current paper, the type variable is continuous while the outcome variable, *Success*, is binary, which implies that perfect separation is not possible. This binary outcome variable represents scenarios in which a principal learns whether a client was satisfied with the agent's work. Additionally, the least-cost separating equilibrium characterized by Kaya (2009) depends on differences in the concavities of the payoff functions across types: when

the payoff functions have the same concavity, as in the current model with binary effort choice, a perfectly separating equilibrium cannot occur. Roddie (2012) derives the Riley (1979) separating equilibrium for repeated interactions but requires a positive probability of type change between periods to obtain an equilibrium with signaling efforts in later periods. Heinsalu (2018) also focuses on costly signaling in a multiperiod model in which effort is observed with noise. Because costs are convex in his model, there is a trade-off: agents want to smooth effort over time, but smoothing effort hinders separation possibilities. As in the early work of Holmström (1999), he finds that the more accurate beliefs are, the weaker are the incentives to exert costly effort. Noldeke and Van Damme (1990) develop a repeated version of the Spence (1978) model. They assume that at some point, the payoff from signaling (being hired at a given wage) is realized, and find that if periods are continuous (i.e., the time between periods approaches zero), the solution of their model coincides with the stable outcome of the one-shot game: the Riley (1979) outcome, which is the least-cost separating equilibrium.

Given the model's continuous type distribution, a binary effort choice with fixed costs, and the observation of a binary outcome, a perfectly separating equilibrium cannot exist in the current model. To derive a testable prediction, one must select among multiple equilibria. When excluding the pooling equilibrium in which no type exerts costly effort, the remaining equilibria in the current setup are semi-separating: observing success leads to higher beliefs about the agent's competence, and agents in different competence intervals follow different effort strategies.

Preferences for Information Another relevant strand of literature examines preferences for information and its timing. This literature is predominantly experimental; one exception is Golman and Loewenstein (2018), which provides a theoretical foundation for preferences for non-instrumental information. The experimental findings indicate that agents value receiving information early (Falk and Zimmermann, 2023; Nielsen, 2020; Ganguly and Tasoff, 2017) and that individuals have a strong desire to learn about one's relative position in real effort tasks (Alós-Ferrer et al., 2018). However, whether people react reciprocally by rewarding other players who provide feedback and punishing those who do not remains an open empirical question, with no empirical foundation thus far.[4]

This paper bridges these strands by developing a repeated signaling model of IPE and experimentally testing how agents respond when a principal either provides or withholds in-

[4]Reciprocal behavior, however, is well established in other contexts (e.g. Fehr et al., 1997, 2002; Engelmann and Fischbacher, 2009).

formation that may be intrinsically valuable to them.

3 The Model

A principal (she) and an agent (he) engage in a two-period game. The agent i is privately informed about his competence type $\theta_i \sim \mathcal{U}[0,1]$ and decides for each period $t \in \{1,2\}$ whether to exert effort by choosing one of two effort levels $e_t \in \{\varepsilon, 1\}$ where $\varepsilon > 0$.[5] Effort is unobservable to the principal but comes at publicly known cost $c > 0$ for the agent. In each period, the agent produces output that can be of high or low quality. The agent does not observe the quality of the output produced. The probability of *Success* ($S_t \in \{0,1\}$) in producing high-quality output depends on his competence and effort in the respective period:

$$Pr(S_t = 1) = 1 - Pr(S_t = 0) = \theta_i \cdot e_t \tag{1}$$

The agent is motivated by career concerns. His utility depends on the posterior belief about his competence type held by the principal (or an external observer) at the end of the game.[6] Incorporating the principal's belief directly into the agent's payoff can be interpreted as a reduced-form representation of the future benefits associated with being perceived as a higher type. Equivalently, the setup can be viewed as a psychological game with belief-dependent utility, following standard approaches in the signaling literature (see Battigalli and Dufwenberg, 2022, for a discussion). The agent's expected utility is given by

$$EU_i(\theta_i, e_1, e_2) = \mathbb{E}\left[\hat{\theta}(S_1, S_2) \mid \theta_i, e_1, e_2\right] - c \cdot \left(\mathbb{1}\{e_1 = 1\} + \mathbb{1}\{e_2 = 1\}\right), \tag{2}$$

where $\hat{\theta}(S_1, S_2) = \mathbb{E}[\theta_i \mid S_1, S_2]$ denotes the posterior belief about the agent's type at the end of the second period. The principal and the observer hold this common belief, since both observe the same outcomes and update by Bayes' rule.

[5] The modeling choice of $\varepsilon > 0$ instead of 0 as the lower effort level eliminates the need for assumptions about out-of-equilibrium beliefs, as, given the production function, neither success nor failure is proof of high or low effort provision.

[6] In the literature (see, e.g., Visser and Swank, 2007; Levy, 2007; Fehrler and Hughes, 2018; Fehrler and Janas, 2021), this observer may be the principal or a broader market. In line with the career-concern literature, I assume that the observer updates beliefs rationally.

The principal prefers success to failure in each period. Her utility is

$$U_P(S_1, S_2) = F(S_1, S_2),$$

where F is strictly increasing in each argument on $\{0,1\}^2$.

The observer aims to form accurate beliefs about the agent's competence:

$$U_O(S_1, S_2) = -\left(\hat{\theta}(S_1, S_2) - \theta_i\right)^2, \tag{3}$$

where θ_i denotes the agent's true type, which is unknown to the observer.

At the beginning of the game, the principal commits to one of two publicly observable regimes: interim performance evaluation (*IPE*) or no interim performance evaluation (*NoIPE*). Under *IPE*, the principal truthfully reveals the first-period outcome S_1 to the agent at the end of period 1. Under *NoIPE*, the agent does not learn S_1.

Timeline. The timing is as follows:

Stage 0 The principal publicly chooses *IPE* or *NoIPE*. The agent privately learns his type $\theta_i \sim U[0,1]$.

Stage 1 The agent privately chooses effort $e_1 \in \{\varepsilon, 1\}$ and generates outcome $S_1 \in \{0,1\}$.

Stage 2 Information revelation:

- **NoIPE:** The agent does not observe S_1.
- **IPE:** The agent observes S_1.

Stage 3 The agent privately chooses effort $e_2 \in \{\varepsilon, 1\}$ and generates outcome $S_2 \in \{0,1\}$.

Stage 4 The principal and the observer observe (S_1, S_2), and the observer forms the posterior belief $\hat{\theta}(S_1, S_2)$.[7]

Stage 5 Payoffs are realized.

4 Equilibrium Predictions

To derive predictions, I characterize the Perfect Bayesian equilibria (PBEs) of the game. A PBE consists of a strategy profile and a belief system such that beliefs are formed according

[7]Since no further actions are taken after period 2, whether the agent also observes S_2 does not affect equilibrium behavior.

to Bayes' rule whenever possible and the agent's strategy is sequentially optimal for each type at every information set.

As is standard in signaling games, the model admits multiple equilibria. When $\varepsilon = 0$, equilibrium selection depends on restrictions on off-path beliefs, since success occurs with zero probability under low effort. By contrast, when $\varepsilon > 0$, both success and failure occur with positive probability under any effort choice, implying that no observation is off the equilibrium path. Consequently, there exists a range of effort costs for which a pooling equilibrium in which all types choose low effort exists alongside equilibria with costly effort provision. In this pooling equilibrium, effort choices do not differ across types, so success is less informative about competence. As a result, posterior beliefs respond less strongly to observed success, which in turn weakens incentives to exert effort. Hence, as $\varepsilon \to 0$, the no-effort strategy profile constitutes an equilibrium for all $c > \frac{1}{6}$, even though equilibria with positive effort provision continue to exist for all $c < \frac{1}{2}$. Numerically, the lowest cost that supports this equilibrium falls as ε increases, and it tends to zero as $\varepsilon \to 1$, when effort no longer affects the distribution of outcomes.[8]

This pooling equilibrium is not useful for comparing the incentive effects of interim performance evaluation. To obtain meaningful predictions, I therefore restrict attention to non-pooling equilibria with active signaling. In particular, following the equilibrium-selection logic of the signaling literature (e.g., Cho and Kreps, 1987; Riley, 1979), I focus on equilibria that maximize signaling activity, that is, equilibria exhibiting the highest effort level for a given cost parameter. Because every outcome occurs with positive probability when $\varepsilon > 0$, refinements that discipline off-path beliefs place no restrictions here, so the selection operates directly on the set of equilibria rather than through belief restrictions.

Subsection 4.1 presents the equilibrium strategies under the *NoIPE* regime, and Subsection 4.2 presents those under the *IPE* regime. Subsection 4.3 compares the resulting predictions under the two regimes regarding expected total effort provision and expected total success. The equilibrium strategies under *NoIPE* and *IPE* are stated in Propositions 1 and 2, respectively, and the resulting differences in effort provision are summarized in Proposition 3. Appendices A2 and A3 present the derivations of the equilibrium strategies, and Appendix A4 presents further insights into the comparison summarized in Proposition 3.

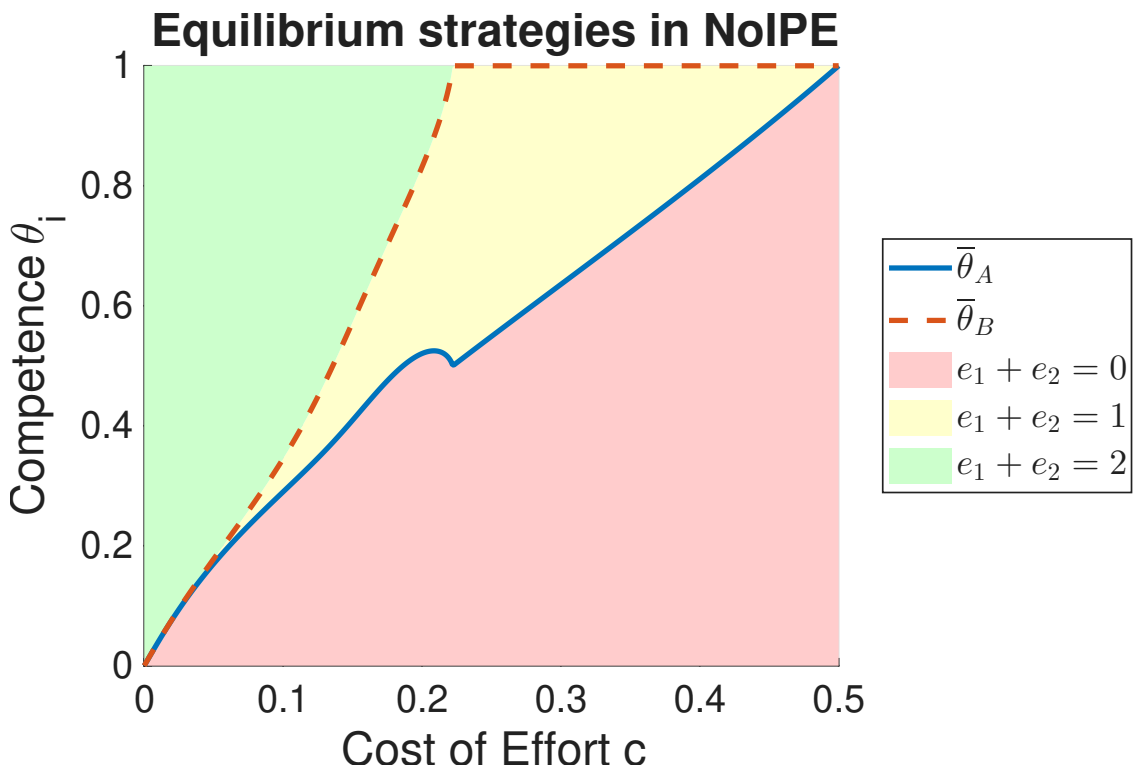

Figure 1: NoIPE: Competence thresholds $\bar{\theta}_A$ and $\bar{\theta}_B$ for $\varepsilon \to 0$ that determine the highest-effort equilibrium strategies.

4.1 Equilibria in the NoIPE regime

As in all signaling games, the opportunity to separate oneself from other types depends on the cost of signaling. Consequently, the equilibrium strategies vary with the cost levels. For very high costs, no type is willing to exert costly effort. For costs close to zero, (almost) every type exerts effort in each period.

For low costs, there is an equilibrium with two thresholds $\bar{\theta}_A(c) \leq \bar{\theta}_B(c)$: high types ($\theta_i > \bar{\theta}_B$) exert effort in both periods, middle types ($\bar{\theta}_A < \theta_i \leq \bar{\theta}_B$) exert effort exactly once, and low types ($\theta_i \leq \bar{\theta}_A$) do not exert effort. In this low-cost region, the timing of the one-time effort is uniquely pinned down. Because high types exert effort in both periods, they generate the outcomes $(S_1, S_2) = (1, 0)$ and $(S_1, S_2) = (0, 1)$ with identical likelihoods, and the rare successes generated by low effort are likewise symmetric across the two outcomes. Any asymmetry between the two one-success outcomes must therefore stem from the timing of the middle types. Suppose middle types exerted effort in the first period with probability above one half. The outcome $(1, 0)$ would then originate relatively more often from middle types and relatively less often from high types, and conversely for $(0, 1)$. Since middle types

[8] Details and derivation of this pooling equilibrium are outlined in Appendix A1.

lie below $\bar{\theta}_B$ and high types above it, the posterior belief after the crowded outcome $(1,0)$ would fall below the posterior belief after $(0,1)$, and every middle type would strictly prefer to move his effort to the second period. The symmetric argument rules out any probability below one half. Deviation incentives thus always point toward the less crowded period, which excludes asymmetric mixing and, in particular, both pure timing profiles. Hence, in the low-cost region, agents who exert effort exactly once randomize with probability 0.5 between first- and second-period effort.

For intermediate cost levels, exerting effort twice is no longer optimal for any type. Low types still exert no effort, while sufficiently high types exert effort exactly once. Without feedback provision, second-period effort cannot depend on the first-period outcome. In the limiting case ($\varepsilon = 0$), the timing of this one-time effort is not pinned down uniquely, because posterior beliefs after $(S_1, S_2) = (1,0)$ and $(S_1, S_2) = (0,1)$ are the same for any common interior mixing probability. However, for any $\varepsilon > 0$, even if arbitrarily small, inactive types can also generate success. This breaks the indifference under asymmetric mixing probabilities: if active types exert effort more often in one period than in the other, the posterior beliefs following the two one-success outcomes are no longer equal. The direction, however, is the opposite of the low-cost case. Because no type exerts effort twice, rare low-effort successes are the only other source of a one-success outcome, and they make up a relatively larger share of the outcome that active types use less often. Since these rare successes stem from a pool that includes all the inactive types below $\bar{\theta}_A$, the posterior after the less used period falls below the posterior after the crowded period, and the deviation incentive favors the crowded period. Concentrating effort in a single period is therefore self-sustaining. Hence, with $\varepsilon > 0$, the interior mixed equilibrium requires symmetric mixing, while pure timing equilibria also exist in which all active types exert effort only in the first period or only in the second period. In all three cases, the same threshold separates inactive from active types, and all active types exert effort exactly once. Total expected effort is therefore identical across these equilibria, even though the timing of effort differs.

Given thresholds $0 < \bar{\theta}_A(c) \leq \bar{\theta}_B(c) \leq 1$, Proposition 1 summarizes the selected active-signaling equilibria of the NoIPE regime:

Proposition 1. *The selected active-signaling PBE strategies under the NoIPE regime, in the limit as $\varepsilon \to 0$, depend on c and have the following threshold structure:*

- *For $0 \leq c < \frac{2}{9}$, agents with $\theta_i \leq \bar{\theta}_A(c)$ play $(e_1; e_2) = (0;0)$, agents with $\bar{\theta}_A(c) < \theta_i \leq \bar{\theta}_B(c)$ play $(e_1; e_2) = (1;0)$ or $(e_1; e_2) = (0;1)$ with 50% probability each, and agents with $\bar{\theta}_B(c) < \theta_i$ play $(e_1; e_2) = (1;1)$.*

- *For $\frac{2}{9} \leq c < \frac{1}{2}$, agents with $\theta_i \leq \bar{\theta}_A(c)$ play $(e_1; e_2) = (0;0)$, while agents with $\bar{\theta}_A(c) < \theta_i$*

exert effort exactly once. There are three timing equilibria: all active agents play $(e_1; e_2) = (1; 0)$*, all active agents play* $(e_1; e_2) = (0; 1)$*, or all active agents mix symmetrically between* $(e_1; e_2) = (1; 0)$ *and* $(e_1; e_2) = (0; 1)$ *with probability* 50% *each. In all three equilibria, the same threshold* $\bar{\theta}_A(c)$ *separates inactive from active types, so total expected effort is identical.*

- *For* $c \geq \frac{1}{2}$*, no costly high effort is provided in any period.*

Proof. See Appendix A2. □

Figure 1 illustrates the equilibrium thresholds across the range of c for which some effort is still exerted ($c < \frac{1}{2}$). The green area represents cost-competence combinations for which effort is exerted in both periods. The red area represents combinations for which no effort is exerted. The yellow area represents combinations for which agents exert effort exactly once. In the low-cost region ($c < \frac{2}{9}$), these agents randomize with probability 0.5 between exerting effort only in the first period and exerting effort only in the second period. In the medium-cost region ($\frac{2}{9} \leq c < \frac{1}{2}$), the yellow area represents the equilibria in which active agents exert effort only in the first period, only in the second period, or mix symmetrically between the two. Appendix A2 presents the derivations of the equilibrium strategies in detail and provides visualizations of the equilibrium belief-updating behavior.

4.2 Equilibria in the IPE regime

The key difference between IPE and NoIPE is that agents under IPE learn the outcome of the first period before deciding on their second-period effort. This information allows agents to condition second-period effort on first-period success. The value of this information is largest for agents who exert effort in the first period, since these agents learn whether their first-period effort has already generated a favorable signal. Agents who do not exert effort in the first period obtain less useful information from the evaluation, because in the limit as $\varepsilon \to 0$, choosing low effort in the first period implies $S_1 = 0$ with probability one.

As under NoIPE, more able types gain more from effort, because every reputational gain is weighted by the probability θ_i that effort produces success. Types therefore sort by ability into up to four groups: no effort, effort exactly once, the conditional strategy, and effort in both periods. The conditional strategy arises because a second success is worth less after an early success than after an early failure. An agent who succeeded in the first period is already pooled with relatively able types, so a second success raises the posterior only from $\hat{\theta}(1, 0)$ to $\hat{\theta}(1, 1)$. An agent who failed is pooled with all inactive types at $\hat{\theta}(0, 0)$, and a second-period success lifts him to $\hat{\theta}(0, 1)$. Hence a range of types tries again after a failure

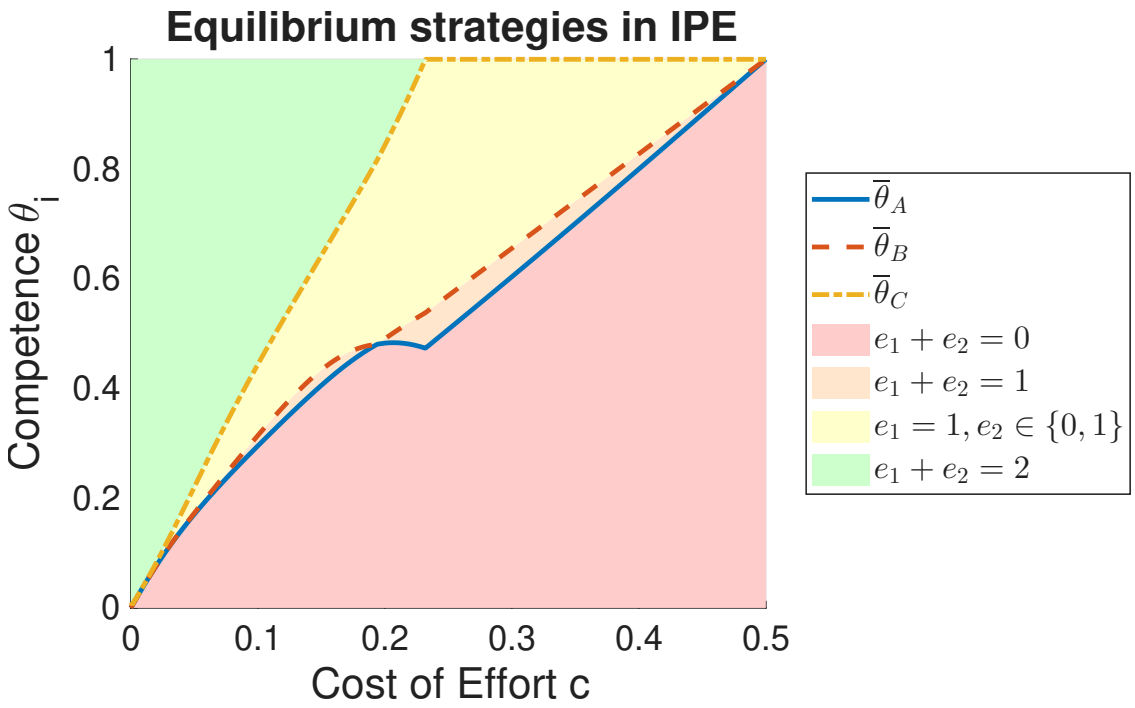

Figure 2: IPE: Competence thresholds $\bar{\theta}_A$, $\bar{\theta}_B$, and $\bar{\theta}_C$ for $\varepsilon \to 0$ that determine the selected active-signaling equilibrium strategies.

and stops after a success. At high costs, the cost exceeds even the largest possible gain after an early success, $1 - \hat{\theta}(1, 0)$, so no type exerts effort again after a success. At lower costs, the most able types find a second success worth its cost even after an early success and exert effort in both periods.

Proposition 2 outlines the selected active-signaling equilibrium strategies when IPE is implemented. Figure 2 visualizes the equilibrium thresholds as a function of c. In Figure 2, the green area represents cost-competence combinations for which agents exert effort in both periods. The red area represents cost-competence combinations for which agents exert no effort. The orange area represents cost-competence combinations for which agents exert effort exactly once: in the low-cost region, these types exert effort only in the second period, while in the medium- and high-cost regions, they exert effort only in the first period. The yellow area represents cost-competence combinations for which agents follow the conditional strategy: they exert effort in the first period and exert effort in the second period only if they do not achieve success in the first period. Appendix A3 provides the derivations of these equilibrium strategies in detail and visualizes the corresponding belief-updating behavior.

Let c_I^L and c_I^H denote the cost thresholds separating the low-, medium-, and high-cost regions under IPE. The cutoff c_I^L is defined as the cost at which one-period effort types switch from exerting effort in the second period to exerting it in the first period. At this

cost, $\bar{\theta}_A(c_I^L) = \bar{\theta}_B(c_I^L)$, so the set of one-effort types is empty when the switch occurs. The cutoff c_I^H is defined as the cost at which $\bar{\theta}_C(c_I^H) = 1$, so that above this cost no type exerts effort unconditionally in both periods. Numerically, $c_I^L \approx 0.194$ and $c_I^H \approx 0.232$.

The switch at c_I^L follows from who generates each one-success outcome, because one-time effort goes to the period whose success earns the higher posterior. Conditional types land in $(1,0)$ with early success and in $(0,0)$ or $(0,1)$ after an early failure. When costs are high and few types exert effort twice, this makes $(0,1)$ the weaker signal, since an early failure is bad news about ability, and one-time effort takes place in the first period. When costs are low, the most able types exert effort in both periods and reach both outcomes equally often. Relative to $(0,1)$, the outcome $(1,0)$ then contains more conditional types and fewer of the most able types, so $\hat{\theta}(0,1)$ rises slightly above $\hat{\theta}(1,0)$ and one-time effort moves to the second period. Each of these types still exerts effort exactly once; only the timing changes.

Proposition 2. *The selected active-signaling PBE strategies under the IPE regime, in the limit as $\varepsilon \to 0$, depend on c and have the following threshold structure:*

- *For $0 \leq c < c_I^L$, agents with $\theta_i \leq \bar{\theta}_A(c)$ play $(e_1; e_2) = (0; 0)$; agents with $\bar{\theta}_A(c) < \theta_i \leq \bar{\theta}_B(c)$ play $(e_1; e_2) = (0; 1)$; agents with $\bar{\theta}_B(c) < \theta_i \leq \bar{\theta}_C(c)$ play $e_1 = 1$ and then choose $e_2 = 0$ if $S_1 = 1$ and $e_2 = 1$ if $S_1 = 0$; agents with $\theta_i > \bar{\theta}_C(c)$ play $(e_1; e_2) = (1; 1)$.*
- *For $c_I^L \leq c < c_I^H$, agents with $\theta_i \leq \bar{\theta}_A(c)$ play $(e_1; e_2) = (0; 0)$; agents with $\bar{\theta}_A(c) < \theta_i \leq \bar{\theta}_B(c)$ play $(e_1; e_2) = (1; 0)$; agents with $\bar{\theta}_B(c) < \theta_i \leq \bar{\theta}_C(c)$ play $e_1 = 1$ and then choose $e_2 = 0$ if $S_1 = 1$ and $e_2 = 1$ if $S_1 = 0$; agents with $\theta_i > \bar{\theta}_C(c)$ play $(e_1; e_2) = (1; 1)$.*
- *For $c_I^H \leq c < \frac{1}{2}$, agents with $\theta_i \leq \bar{\theta}_A(c)$ play $(e_1; e_2) = (0; 0)$; agents with $\bar{\theta}_A(c) < \theta_i \leq \bar{\theta}_B(c)$ play $(e_1; e_2) = (1; 0)$; agents with $\theta_i > \bar{\theta}_B(c)$ play $e_1 = 1$ and then choose $e_2 = 0$ if $S_1 = 1$ and $e_2 = 1$ if $S_1 = 0$.*
- *For $c \geq \frac{1}{2}$, no costly high effort is provided in any period.*

Proof. See Appendix A3. □

4.3 Comparison of IPE and NoIPE

Theoretical Result 1: Whether IPE increases or decreases total effort depends on the cost of effort, and the effort ranking reverses at a single cost threshold.

Let $\mathcal{E}^r(c)$ denote expected total high effort under regime $r \in \{\text{NoIPE}, \text{IPE}\}$. Because types are uniformly distributed, expected total effort follows directly from the equilibrium thresholds by integrating each type's effort choice over $[0,1]$.

Under NoIPE,

$$\mathcal{E}^{NoIPE}(c) = \begin{cases} 2 - \bar{\theta}_A^N(c) - \bar{\theta}_B^N(c), & \text{if } 0 \le c < \frac{2}{9}, \\ 1 - \bar{\theta}_A^N(c), & \text{if } \frac{2}{9} \le c < \frac{1}{2}, \\ 0, & \text{if } c \ge \frac{1}{2}. \end{cases}$$

For $c < \frac{2}{9}$, types in $(\bar{\theta}_A^N, \bar{\theta}_B^N)$ exert effort exactly once and types above $\bar{\theta}_B^N$ exert effort in both periods; for $\frac{2}{9} \le c < \frac{1}{2}$ all active types exert effort exactly once.

Under IPE,

$$\mathcal{E}^{IPE}(c) = \begin{cases} 2 - \bar{\theta}_A^I(c) - \bar{\theta}_B^I(c) - \dfrac{(\bar{\theta}_C^I(c))^2 - (\bar{\theta}_B^I(c))^2}{2}, & \text{if } 0 \le c < c_I^H, \\ 1 - \bar{\theta}_A^I(c) + \dfrac{\left(1 - \bar{\theta}_B^I(c)\right)^2}{2}, & \text{if } c_I^H \le c < \frac{1}{2}, \\ 0, & \text{if } c \ge \frac{1}{2}. \end{cases}$$

The first expression applies throughout the low- and medium-cost IPE regimes, $0 \le c < c_I^H$. Although the equilibrium thresholds solve different systems in the two sub-ranges $0 \le c < c_I^L$ and $c_I^L \le c < c_I^H$, the effort accounting is identical: in both, the active population consists of one-effort types in $(\bar{\theta}_A^I, \bar{\theta}_B^I)$, conditional types in $(\bar{\theta}_B^I, \bar{\theta}_C^I)$ who exert expected effort $1+(1-\theta)$, and two-effort types above $\bar{\theta}_C^I$. At c_I^L, the two systems have the same solution, with $\bar{\theta}_A^I = \bar{\theta}_B^I$, so expected total effort is continuous at the switch. The second expression corresponds to the high-cost regime, in which $\bar{\theta}_C^I = 1$ and the conditional types span $(\bar{\theta}_B^I, 1)$; it is obtained from the first by setting $\bar{\theta}_C^I = 1$, and its two terms are first-period effort $1 - \bar{\theta}_A^I$ and the conditional second-period contribution $\frac{1}{2}(1 - \bar{\theta}_B^I)^2$.

In the NoIPE regime, the medium-cost range admits several timing equilibria. However, all of these equilibria induce the same total effort, since the same threshold separates inactive from active types and all active types exert effort exactly once. The comparison of total effort is therefore independent of which timing equilibrium is selected within this range. The following proposition summarizes the resulting comparison across regimes.

Proposition 3. *Expected total effort is strictly higher under IPE than under NoIPE for all $c \in [c_I^H, \frac{1}{2})$, and equals zero in both regimes for $c \ge \frac{1}{2}$. Numerically solving the equilibrium threshold systems on $(0, c_I^H)$ shows that the expected total effort is strictly higher under NoIPE than under IPE for all costs below a single crossing point $c^* \approx 0.135$, and strictly higher under*

IPE than under NoIPE above it, so the effort ranking reverses at a single cost threshold.

Proof. See Appendix A4 for the proof of the analytic statements and the numerical characterization of the crossing. □

Proposition 3 reflects two opposing effects of IPE on total effort. First, IPE lets agents condition second-period effort on first-period performance. This produces a *stopping effect*: an agent who has already generated a favorable signal and thus separated from the low types has weaker incentives to exert costly effort again, so effort falls after an early success. The stopping effect dominates when effort is cheap. Under NoIPE, high types are willing to exert effort in both periods; under IPE, many of these same agents economize on effort once a first-period success has separated them from lower types. As a result, NoIPE induces more total effort for $c < c^*$.

Second, IPE makes first-period effort more attractive through a *conditional-continuation effect*. By exerting effort early, an agent learns whether he has already separated from the lower types, and can still try again in the second period if the first attempt fails. This conditional strategy is most valuable when effort is costly enough that exerting effort unconditionally in both periods is no longer worthwhile. In this range NoIPE induces active agents to exert effort only once, whereas under IPE some agents exert effort in the first period and again only after a failure. This raises total effort under IPE relative to NoIPE for all costs above c^*, up to the point at which active signaling disappears. The two forces exactly balance at $c^* \approx 0.135$, where the effort ranking reverses.

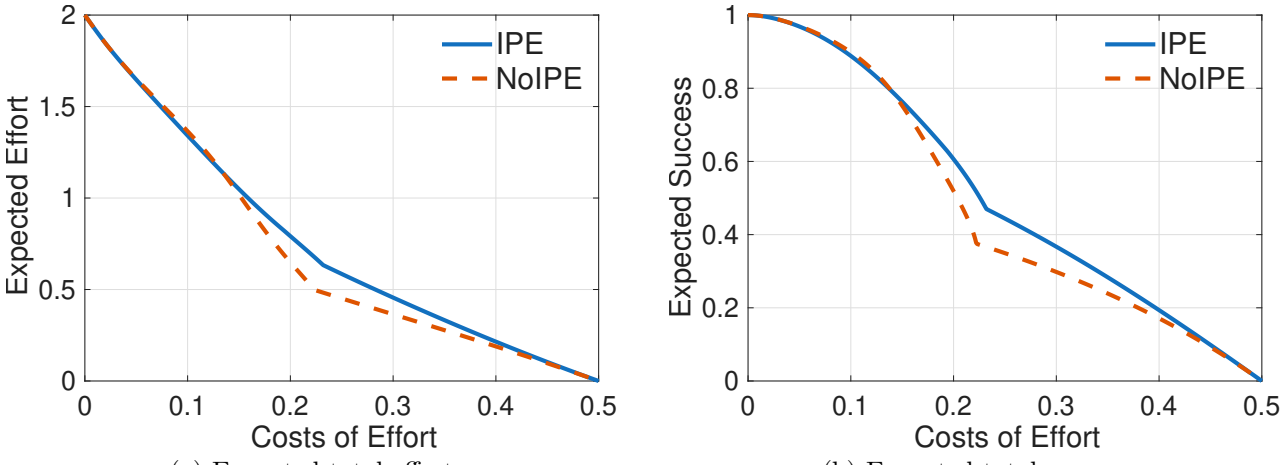

(a) Expected total effort

(b) Expected total success

Figure 3: Comparison of IPE and NoIPE. Panel (a) shows expected total high effort, $\mathcal{E}^r(c)$, under the selected active-signaling equilibrium in each regime. Panel (b) shows expected total success, $E[S_1] + E[S_2]$.

Figure 3 illustrates the comparison. Panel (a) plots expected total high effort $\mathcal{E}^r(c)$ under the selected active-signaling equilibrium in each regime. Consistent with Proposition 3, the two curves cross once: NoIPE lies above IPE for $c < c^*$ and IPE lies above NoIPE for $c^* < c < \frac{1}{2}$, with both vanishing at $c = \frac{1}{2}$. The magnitudes are asymmetric: the NoIPE advantage at low costs is small, so the two curves nearly coincide there, whereas the IPE advantage at higher costs is sizable. The maximal NoIPE advantage is about 0.024 units of expected total effort, attained near $c \approx 0.10$, while the maximal IPE advantage is about 0.18, attained near $c \approx 0.22$.

Panel (b) plots expected total success, $E[S_1] + E[S_2]$. The effort and output comparisons are related but not identical: because higher types are more likely to convert a given effort choice into success, reallocating effort across types affects expected success even when total-effort differences are small. For low costs the two regimes differ little, since both induce substantial effort by high types. For intermediate costs IPE generates higher expected total success, as the conditional strategy directs additional second-period effort toward agents who failed to produce success in the first period. For sufficiently high costs active signaling disappears in both regimes, and both expected effort and expected success converge to zero.

5 Experimental Design and Hypotheses

To test whether IPE can indeed cause higher effort provision, I conduct a laboratory experiment. There are several behavioral factors that could further differentiate the effects of NoIPE and IPE. In this experiment, I focus on two of them: intrinsic motivation to signal talent and curiosity about outcomes. If agents are intrinsically motivated to signal their talent, in addition to the extrinsic motivation driven by career concerns, their effort might increase in both periods. However, the model does not provide a clear theoretical prediction about how this would affect beliefs, equilibria, and therefore also the difference between IPE and NoIPE. To test the influence of intrinsic motivation, the experiment includes either an abstract task, in which agent types are randomly assigned, or quiz questions, in which the agent's *competence level* is based on their actual knowledge.

If agents are curious about the results of their effort, they might punish the principal for withholding information by reducing their effort. This assumption is tested by introducing another treatment variation: either the principal or a computer decides between IPE and NoIPE. Theoretical predictions do not differ depending on whether the principal or the computer chooses IPE or NoIPE, as reciprocity concerns are not included in the utility functions. By varying both the nature of the task and the endogeneity of IPE implementation, the experiment also tests the interaction effects between curiosity/reciprocity and

intrinsic motivation. This design allows for examining whether, for example, agents who are intrinsically motivated to signal their type are particularly likely to punish the principal for withholding information about their success.

5.1 Experimental Design

The treatments of the experiment vary along two dimensions: 1) whether the decision to implement NoIPE/IPE is made by the principal or by a computer, and 2) whether the task is abstract with randomly assigned types, or a quiz where subjects' competence is determined by correctly answering questions, without any type assignment by the experimenter. The decision-making authority for implementing IPE (whether by the computer or the principal) is varied within-subject. Half of the sessions begin with exogenous implementation by the computer, while the other half start with endogenous implementation by the principal. The context of the task (Abstract/Quiz) is varied between subjects. Principals and agents are randomly matched within their matching groups, and a stranger matching protocol is applied between rounds.

The experiment consists of 10 rounds, each containing two periods. At the start, participants are randomly assigned to the roles of either principal/observer or agent, with half of the subjects in each role.[9] In the Abstract treatments, each agent is privately assigned a random type ($Type \in \{0, 1, 2, ..., 100\}$) at the beginning of each round, with all types being equally likely. This type determines the potential probability of success in each period of the round. In the Quiz treatments, each subject answers 10 questions at the beginning of the experiment. Each question is of the form: "What is the capital of [country]?" with four options: the capital and the three largest cities other than the capital in that country. The 10 questions are randomly and independently drawn from a sample of 194 countries for each subject, with random positioning of the correct answer. After completing the quiz, each subject privately learns the number of correct answers, which, for the agents, determines their type for the remainder of the experiment. For principals, the quiz serves only to give an impression of the task's complexity.

At the start of each round, an agent is given a budget of 60 points. In each of the two periods, the agent can choose to pay 30 points to attempt to produce success. If the agent chooses not to pay, success is impossible in that period. In the Abstract treatments, if the agent pays the 30 points, the probability of success equals their type (in percent). In the Quiz treatments, agents also choose whether to pay 30 points and answer a new set of randomly

[9]Neutral wording is used throughout the experiment. For details, refer to the Instructions and Decision Screens in the Supplementary Material SM1 and SM2.

drawn capital questions in each period. Success occurs if the costs have been paid and at least 7 out of 10 questions are answered correctly.

The principal serves two roles: they act as the principal for one agent and as an observer for another agent within the same matching group. When acting as the principal, they can choose to implement IPE or not during the endogenous rounds. The principal earns 50 points for each successful outcome achieved by the agent they are matched with. As an observer, the principal must guess the type of another agent based on the observed successes across the two periods. Using the strategy method (Selten, 1967), the observer makes predictions for four scenarios: 1) no success, 2) success in the first period only, 3) success in the second period only, and 4) success in both periods. In the Abstract treatments, these predictions range from 0 to 100; in the Quiz treatments, they range from 0 to 10. Observers are incentivized for accuracy.[10] The agent's round earnings are determined by the observer's guess for the scenario that actually occurred. In the Abstract treatments, the entered number directly translates to the agent's points earned in that round (minus the costs, if paid), while in the Quiz treatment, the agent's earnings are 10 times the entered number (minus the costs, if paid).

At the end of each round, the agent learns his round payoff. The principal learns whether success occurred in each period. The observer does not learn the true type of the agent during the game but will be informed of it on the feedback screen at the end of the experiment. If IPE is implemented, the agent learns whether success occurred in both periods. If NoIPE is implemented, the agent learns only whether success occurred in the second period. At the end of the experiment, three [two] rounds are randomly selected to be payoff-relevant for the agent [principal]. For the observer, two rounds are randomly chosen as payoff-relevant. Each point is converted into Euros at the rate of 25 points = 1 Euro, and each subject receives a show-up fee of €3 in the Abstract treatments and €6 in the Quiz treatments.[11] At the end, subjects are asked to complete a short questionnaire about their socio-economic background and provide open-ended responses regarding their behavior in the experiment.

The experiment was programmed with zTree (Fischbacher, 2007) and conducted in two waves due to lab closures related to the Covid-19 pandemic. The first wave took place in the Lakelab at the University of Konstanz in September and October 2020. It comprised seven sessions of the Quiz treatment with 10 subjects each and eight sessions of the Abstract treatment, with six sessions involving 10 subjects and two sessions involving eight subjects.

[10] In the Abstract treatments, the observer's round payoff equals $100 - |\text{true type} - \text{entered guess}|$ points. In the Quiz treatments, the observer's round payoff equals $100 - 10 \cdot |\text{true type} - \text{entered guess}|$ points.

[11] The laboratory rule is to pay subjects an expected hourly wage. The higher show-up fee in the Quiz treatments stems from the longer experiment duration, because each agent answers 210 questions.

The second wave was conducted at the WiSo lab in Hamburg between May and July 2024, with matching group sizes ranging from 8 to 14, involving 88 subjects in the Abstract treatment and 88 subjects in the Quiz treatment. In total, 322 subjects in 33 matching groups participated in the experiment. Subjects were recruited via hroot (Bock et al., 2014) from the student subject pools at the Universities of Konstanz and Hamburg. The subjects had an average age of 24.75 years (SD = 4.85), 44.12% of them were male, and enrolled in various majors. The Quiz [Abstract] treatments lasted an average of 60 [50] minutes, including the instruction phase and the post-experimental questionnaire. The average payoff in the Quiz [Abstract] treatments was €14.45 (SD = 2.95) [€12.58 (SD = 2.47)].

5.2 Hypotheses

The experimental design allows for testing several hypotheses. Hypotheses 1-3 and 5 are derived from the model. Hypothesis 4 is exploratory and follows from behavioral considerations: the potential intrinsic value of (un)fulfilled curiosity, reciprocity, and the non-monetary utility from signaling competence in the Quiz task. The predictions are based on the implemented effort cost of 30 points, corresponding to the cost level of $c = 0.3$ in the theory section. This cost lies above the crossing point $c^* \approx 0.135$ and above the regime boundary $c_I^H \approx 0.232$, so the one-threshold equilibrium of Proposition 1 and the high-cost conditional-strategy equilibrium of Proposition 2 apply. At this cost the model predicts expected total effort of approximately 0.46 under IPE and 0.36 under NoIPE, and it prescribes the conditional strategy for all types above $\bar{\theta}_B \approx 0.655$ under IPE. Thus, the implemented cost of 30 points makes both the directional effort comparison (Hypothesis 2) and the stopping mechanism (Hypothesis 3) testable with sizable predicted differences.

Hypothesis 1: Total effort increases with the competence type.

Hypothesis 2: Under IPE, there is more total effort provision for a given competence type than without IPE.

Hypothesis 3: Under IPE, among agents who provide effort in Period 1, Period 2 effort is higher after NoSuccess than after Success.

Hypothesis 4: In the endogenous treatments, there is more accumulated effort when the principal chooses to implement IPE than when IPE is imposed exogenously; conversely, there is less accumulated effort when the principal chooses NoIPE than when NoIPE is imposed exogenously. This reciprocity effect is expected to be stronger in the Quiz task than in the Abstract task.

Hypothesis 5: The observer's belief about an agent's competence type increases with the number of Successes.

As the theoretical predictions arise from a specific equilibrium selection criterion, alternative selection criteria might lead to different predictions. For example, at the cost level $c = 0.3$, no one providing effort would be an equilibrium in both regimes. Consequently, it is not clear a priori whether the predicted equilibria will be observed in practice. Additionally, the predicted behavior of principals and agents involves a level of strategic reasoning that may not accurately reflect actual behavior. Prior research also raises questions about whether Bayesian updating accurately describes how people actually update beliefs (e.g., Ouwersloot et al., 1998; Zizzo et al., 2000), which would render a PBE invalid in practice. Finally, while there is evidence in the literature supporting the existence of curiosity (e.g., Alós-Ferrer et al., 2018) and reciprocity (e.g., Falk and Fischbacher, 2006), the extent to which these incentives interact as hypothesized remains uncertain. Therefore, it is not obvious that any of the hypotheses will hold in the laboratory setting.

6 Results

This section begins with the findings on effort provision across the different treatments. Section 6.2 then analyzes the observers' evaluation behavior. Section 6.3 discusses the principals' IPE choices, Section 6.4 provides an overview of further results from the Quiz task, and Section 6.5 explores potential reasons for deviations from the theoretical predictions.

To test for treatment differences, standard errors are clustered at the matching-group level. For comparisons within the Abstract treatments, the paired structure of the data can be taken into account, as agents participate under both IPE and NoIPE regimes, with both endogenous and exogenous implementation, and as different types. Thus, some regressions include subject fixed effects to focus specifically on within-subject differences.

6.1 Effort Provision of the Agents

Experimental Result 1 (Effort and Type): *Hypothesis 1 is confirmed: Effort significantly increases with the type of agent.*

Experimental Result 2 (Effort under IPE): *Hypothesis 2 is confirmed: IPE significantly increases overall effort provision.*

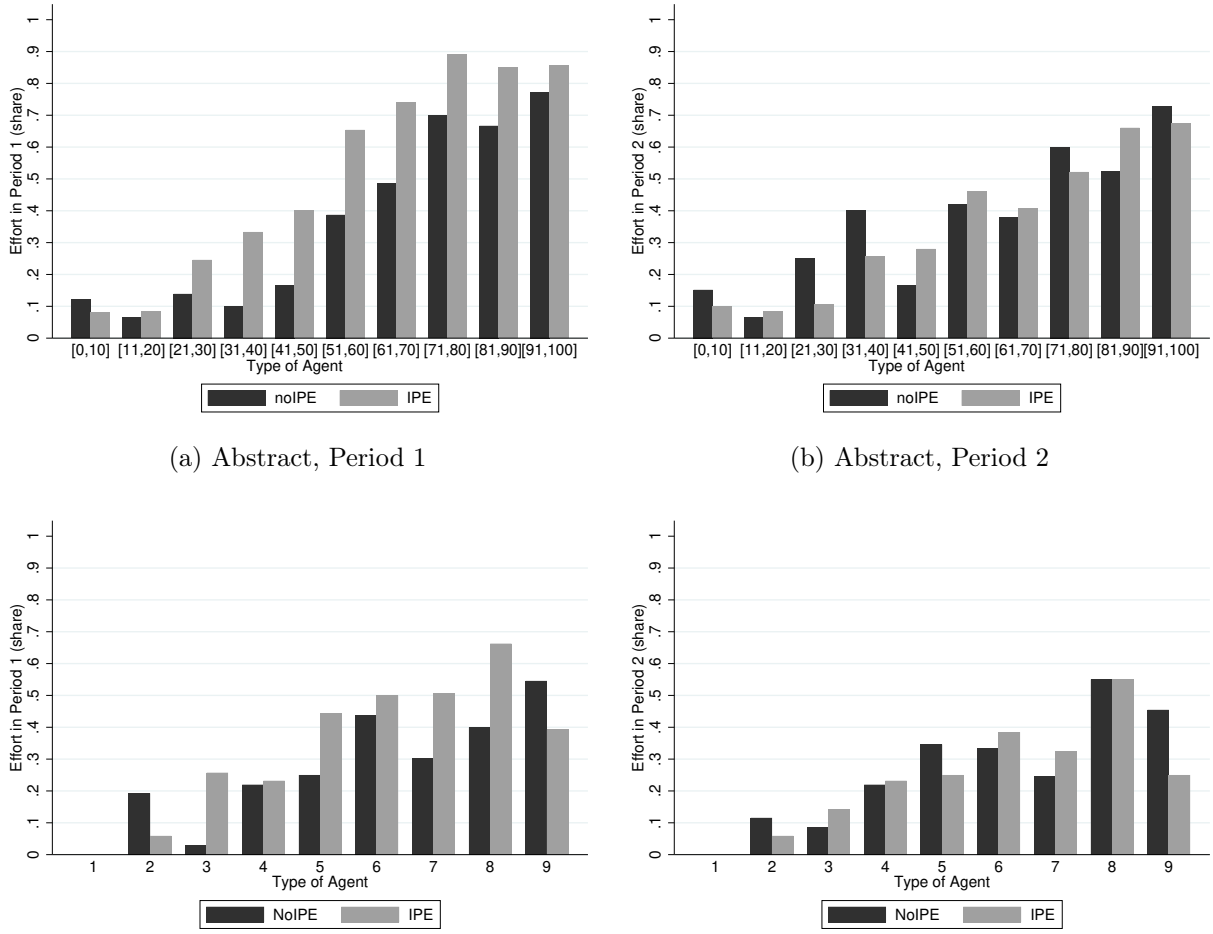

(a) Abstract, Period 1

(b) Abstract, Period 2

(c) Quiz, Period 1

(d) Quiz, Period 2

Figure 4: Effort provision in the first and second period with (IPE) and without (NoIPE) feedback about first period success.

Notes: In the Abstract contexts, type is exogenously given at the beginning of each round. In the Quiz context, type is based on the performance of an agent prior to the first round, and kept throughout the experiment.

For an initial overview, Figure 4 presents the share of effort provision in the first and second periods within the Quiz and Abstract contexts. In both contexts, exerting effort in a period is defined as paying the cost of 30 points to have a chance of Success. Table 1 provides statistical analysis of the influence of agents' types and IPE on overall effort provision. Tables 2 and 3 provide separate tests for first and second period effort provision. The panels in Figure 4 demonstrate that in both contexts, higher types exert more effort. This holds overall (Table 1), as well as for Period 1 (Table 2) and Period 2 (Table 3) separately, and thus confirms Hypothesis 1. Panels (a) and (c) of Figure 4 further show that, for given

Table 1: Total effort provision

	Abstract			Quiz	
	(1)	(2)	(3)	(4)	(5)
Type of Agent	0.016***	0.017***	0.016***	0.123***	0.123***
	(0.001)	(0.001)	(0.001)	(0.037)	(0.037)
IPE	0.144**	0.146***	0.132**	0.117***	0.121***
	(0.054)	(0.044)	(0.056)	(0.037)	(0.035)
Round			-0.045***		-0.061***
			(0.010)		(0.009)
const.	-0.096	-0.636***	0.156*	-0.094	0.239
	(0.086)	(0.075)	(0.084)	(0.178)	(0.180)
Subject FE	No	Yes	No	No	No
Obs.	820	820	820	790	790
Clusters	17	17	17	16	16
R^2	0.326	0.557	0.350	0.109	0.158

Notes: OLS regressions on total effort provision for agents. Columns (1)–(3) report results for the Abstract treatments and columns (4)–(5) for the Quiz treatments. *IPE* denotes rounds in which the agent received feedback after the first period. *Type of Agent* $\in [0, 100]$ in the Abstract treatments, and $\in [0, 10]$ in the Quiz treatments. *Round* $\in (1, 10)$ denotes the experimental round. Subject fixed effects are included only in column (2). Standard errors are clustered at the matching-group level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

types, first-period effort is higher under IPE than under NoIPE. Table 2 confirms that this increase in first period effort through IPE is highly significant. As indicated by the *Round* coefficient, effort tends to decrease over time; however, the significant effect of IPE on effort provision remains robust when controlling for the round in the experiment.

Experimental Result 3 (Second Period Effort): *Hypothesis 3 cannot be confirmed. In the Abstract treatments, receiving positive feedback decreases effort provision only insignificantly, and in the Quiz treatments, receiving positive feedback increases effort provision.*

For second-period effort, panels (b) and (d) of Figure 4 suggest that in both the Abstract and Quiz contexts, second-period effort does not appear to differ systematically between IPE and NoIPE. Table 3 reports the results of an OLS regression that tests this observation statistically. The regressions in column (1) for the Abstract context and column (7) for the Quiz context show that while the agent's type remains a significant predictor of effort provision

Table 2: Period 1 effort provision

	Abstract			Quiz	
	(1)	(2)	(3)	(4)	(5)
Type of Agent	0.010***	0.010***	0.010***	0.067***	0.067***
	(0.001)	(0.001)	(0.001)	(0.019)	(0.019)
IPE	0.153***	0.149***	0.147***	0.117***	0.119***
	(0.029)	(0.025)	(0.031)	(0.035)	(0.029)
Round			-0.024***		-0.036***
			(0.005)		(0.005)
const.	-0.124**	-0.516***	0.008	-0.077	0.121
	(0.044)	(0.044)	(0.046)	(0.093)	(0.096)
Subject FE	No	Yes	No	No	No
Obs.	820	820	820	790	790
Clusters	17	17	17	16	16
R^2	0.331	0.505	0.350	0.098	0.145

Notes: OLS regressions on the choice to pay 30 points to have a chance of success in the first period. Columns (1)–(3) report results for the Abstract treatments and columns (4)–(5) for the Quiz treatments. *IPE* denotes rounds in which the agent received feedback after the first period. *Round* $\in (1, 10)$ denotes the experimental round. Subject fixed effects are included only in column (2). Standard errors are clustered at the matching-group level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

in the second period, the influence of IPE is close to zero and statistically insignificant.

To dig deeper into second-period effort, one can differentiate second-period effort by the effort provision and outcome in the first period. The model admits three active-signaling equilibria for this cost range under NoIPE: low type agents do not provide any effort, and higher type agents either mix uniformly between first- or second-period effort, or provide effort in the first or the second period only. In the second period under IPE, however, the model predicts that for a large share of types, if first-period effort led to success, no effort should be provided in the second period. Conversely, if first-period effort did not lead to success, effort should be provided in the second period.

Columns (6) and (12) of Table 3 test this prediction within IPE rounds in which the agent exerted first-period effort. In the Abstract treatments, first-period success lowers second-period effort by 11.7 percentage points, and the effect is not significant. In the Quiz treatments, first-period success raises second-period effort by 25.1 percentage points. Columns (2) to (5) for the Abstract context and (8) to (11) for the Quiz context show that

Table 3: Period 2 Effort Decision

	All Data	$e_1 = 0$	$e_1 = 1$	$e_1 = 1$, $S_1 = 0$	$e_1 = 1$, $S_1 = 1$	IPE, $e_1 = 1$
Abstract	(1)	(2)	(3)	(4)	(5)	(6)
IPE	-0.009	-0.082**	-0.026	-0.107	0.002	
	(0.033)	(0.036)	(0.053)	(0.084)	(0.054)	
Type of Agent	0.007***	0.003**	0.005**	0.003	0.009***	0.006***
	(0.001)	(0.001)	(0.002)	(0.002)	(0.002)	(0.001)
Period 1 Success						-0.117
						(0.070)
const.	0.028	0.107**	0.285*	0.518***	-0.071	0.231*
	(0.050)	(0.043)	(0.159)	(0.175)	(0.169)	(0.129)
Obs.	820	448	372	116	256	257
Clusters	17	17	17	17	17	17
R^2	0.158	0.056	0.047	0.028	0.098	0.061
Quiz	(7)	(8)	(9)	(10)	(11)	(12)
IPE	0.000	-0.059*	-0.014	-0.071	0.068	
	(0.029)	(0.033)	(0.055)	(0.080)	(0.080)	
Type of Agent	0.056**	0.013	0.079***	0.049	0.091**	0.066**
	(0.019)	(0.014)	(0.027)	(0.039)	(0.035)	(0.029)
Period 1 Success						0.251***
						(0.081)
const.	-0.017	0.115	0.059	0.178	0.025	0.003
	(0.097)	(0.066)	(0.175)	(0.229)	(0.259)	(0.181)
Obs.	790	502	288	142	146	195
Clusters	16	16	16	16	16	16
R^2	0.064	0.013	0.075	0.032	0.099	0.129

Notes: OLS regressions on the choice to pay 30 points to have a chance of success in the second period in the Abstract ((1) - (6)) and the Quiz ((7) - (12)) treatments. *IPE* and *Period 1 Success* are dummy variables. *Type of Agent* $\in [0, 100]$ in the Abstract treatments, and $\in [0, 10]$ in the Quiz treatments. *IPE* denotes the rounds in which the agent received feedback after the first period. Columns (1) and (7) contain all data. Columns (2) and (8) contain only data from rounds in which no effort was provided in the first period. Columns (3) - (5) and (9) - (11) contain only data from rounds in which effort was provided in the first period. Columns (4) and (10) contain only data from rounds in which first period effort did not result in success, and columns (5) and (11) contain only data from rounds in which first period effort resulted in success. Columns (6) and (12) contain only data from rounds with IPE in which effort was provided in the first period. Standard errors are clustered at the matching group level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

the only case in which IPE is associated with (weakly) significantly less second-period effort is the subsample in which no effort was provided in the first period (and therefore no success could have occurred in the first period). These subsample estimates condition on the first-period choice, which is itself affected by IPE (so they compare endogenously selected groups and should not be read as causal effects of feedback). This applies in particular to columns (2) and (8): an agent who exerted no first-period effort knows that no success occurred regardless of the feedback regime, so IPE conveys no information in this subsample, and the negative coefficient is more plausibly a composition effect, as IPE draws marginal types into first-period effort and leaves a more negatively selected group with $e_1 = 0$. Columns (6) and (12), which compare second-period choices across realized first-period outcomes within IPE rounds with first-period effort, provide the most direct test of the conditional-strategy prediction, since the outcome is random conditional on the agent's type, which the regressions control for.

Overall, when combining effort provision across both periods, the regressions in Table 1 show that IPE increases effort provision in both the Abstract and Quiz contexts. This is in line with the model, which predicts increased total effort under IPE at a cost of 0.3 (30 points in the experiment). The empirical results show that this increase is driven by an increase in first period effort while second-period effort does not decrease.

Experimental Result 4 (Effort and Reciprocity): *Hypothesis 4 cannot be confirmed: the endogeneity of the IPE implementation does not influence effort provision in the experiment, neither in the Abstract nor in the Quiz treatments.*

Table 4 investigates the effect of the endogeneity of IPE on effort provision. While Hypothesis 4 predicts that agents would reciprocate by providing more effort if the principal chooses to provide feedback (or less effort if feedback is not provided), this behavior is not evident in the data. There is no significant difference in the effect of IPE when it is implemented by the principal instead of randomly by the computer. Columns (1) and (4) of Table 4 confirm that this lack of effect holds in both the Abstract and Quiz treatments. Columns (2) and (5), as well as (3) and (6), respectively, show that the lack of effect of endogeneity on effort provision holds for both the first and the second period, and in both contexts.

Table 4: Effort provision and endogenous IPE

	Abstract			Quiz		
	Total effort	Period 1	Period 2	Total effort	Period 1	Period 2
	(1)	(2)	(3)	(4)	(5)	(6)
Type of Agent	0.016***	0.010***	0.007***	0.123***	0.067***	0.056**
	(0.001)	(0.001)	(0.001)	(0.037)	(0.019)	(0.019)
Endogenous	-0.002	-0.044	0.041	0.010	-0.020	0.030
	(0.096)	(0.051)	(0.062)	(0.113)	(0.070)	(0.055)
IPE	0.144*	0.118***	0.026	0.101	0.086*	0.015
	(0.068)	(0.031)	(0.051)	(0.069)	(0.045)	(0.045)
Endogenous × *IPE*	-0.000	0.075	-0.075	0.024	0.059	-0.036
	(0.120)	(0.048)	(0.087)	(0.100)	(0.070)	(0.046)
const.	-0.096	-0.111**	0.015	-0.098	-0.071	-0.027
	(0.095)	(0.050)	(0.052)	(0.169)	(0.088)	(0.092)
Obs.	820	820	820	790	790	790
Clusters	17	17	17	16	16	16
R^2	0.326	0.332	0.160	0.110	0.099	0.065

Notes: OLS regressions for effort provision. Columns (1)–(3) report results for the Abstract treatments and columns (4)–(6) for the Quiz treatments. The dependent variables are total effort, period 1 effort, and period 2 effort, respectively. *IPE* denotes rounds in which the agent received feedback after the first period. *Endogenous* is a dummy variable indicating rounds in which the implementation of IPE was decided by the principal rather than assigned by the computer. Standard errors are clustered at the matching-group level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

6.2 Observers' Evaluation Behavior

Experimental Result 5 (Observer Behavior): *Hypothesis 5 is confirmed: observing more successes significantly increases the belief about an agent's type.*

For signaling incentives to be present, observers must form higher beliefs about an agent's type after observing success than after observing no success. Figure 5 displays the cumulative distributions of the observers' beliefs about an agent's type. The observer was informed whether the agent received feedback after the first period (IPE) or not (NoIPE). The graphs demonstrate that signaling incentives to provide effort are indeed present in the experiment: observing success significantly increases the beliefs reported by observers. The difference in beliefs between observing success twice and observing it only once is significant in both the

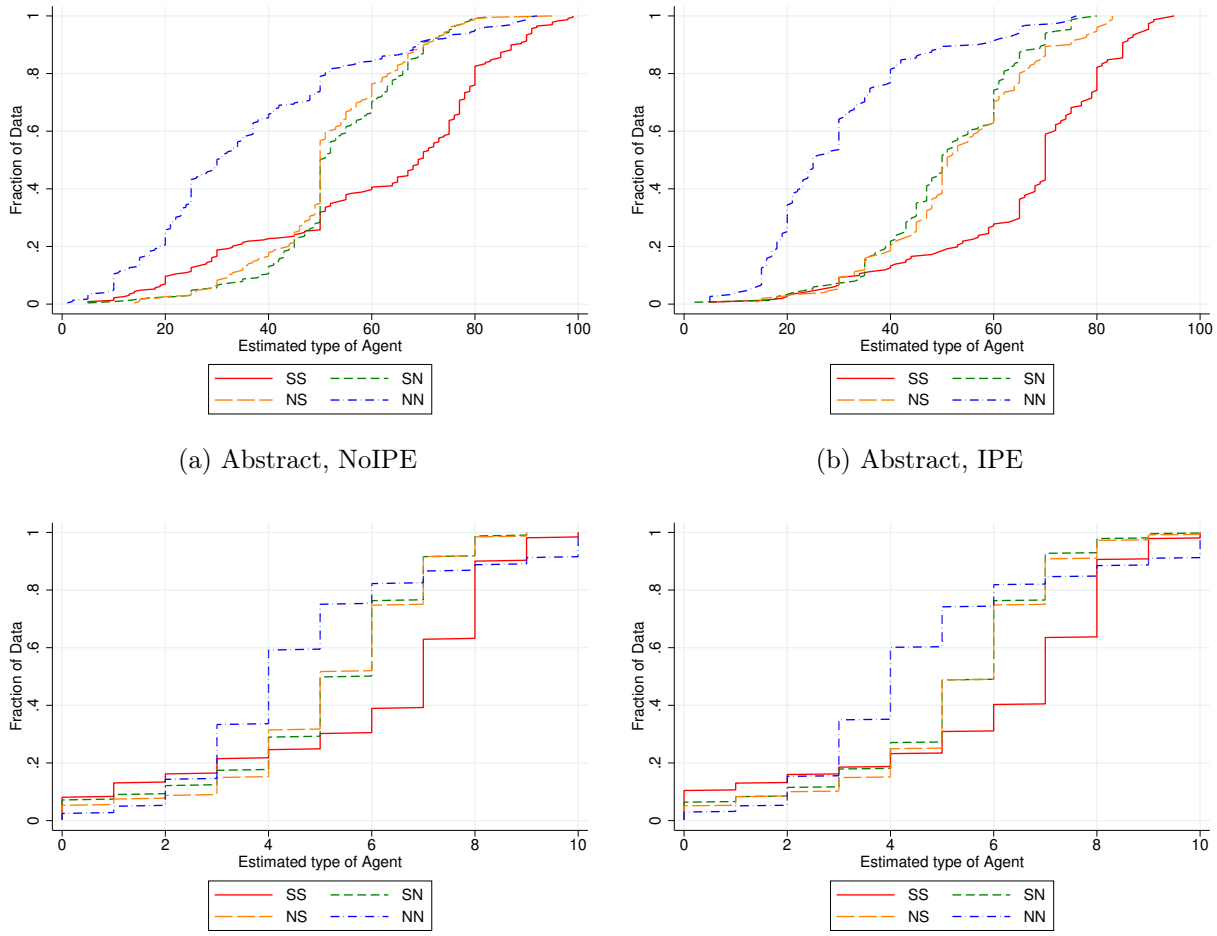

(a) Abstract, NoIPE

(b) Abstract, IPE

(c) Quiz, NoIPE

(d) Quiz, IPE

Figure 5: CDF of Observers' Evaluations

Notes: (SS),(SN),(NS),(NN) correspond to the beliefs about an agent who had (Success, Success), (Success, No Success), (No Success, Success), and (No Success, No Success), respectively. This evaluation was conducted via the strategy method, without feedback between the experimental rounds.

Quiz ($p < 0.05$) and Abstract ($p < 0.001$) treatments.[12] Similarly, the difference in beliefs between observing no success and observing one success is also significant for both the Quiz ($p < 0.01$) and Abstract ($p < 0.001$) treatments. However, the difference in beliefs between success in the first period only and success in the second period only is not significant: this holds true in both the Abstract and Quiz treatments, regardless of whether IPE or NoIPE was played ($p > 0.1$).

Table 5 presents OLS regressions analyzing treatment differences in observers' average

[12]The results stem from t-tests clustered at the principal level.

Table 5: Observers' Beliefs about Agent's Type

	Abstract				Quiz			
	SS	SN	NS	NN	SS	SN	NS	NN
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
IPE	1.310	-0.995	1.834	-4.665	0.113	0.099	0.221	-0.034
	(4.385)	(2.915)	(3.093)	(3.892)	(0.252)	(0.217)	(0.226)	(0.226)
endogenous	-1.485	-0.903	-1.076	-0.792	0.233	-0.004	0.089	0.070
	(1.800)	(1.338)	(1.485)	(1.864)	(0.326)	(0.306)	(0.259)	(0.308)
endogenous	3.860	2.419	0.896	-0.433	-0.291	-0.086	-0.234	0.023
$\times IPE$	(3.046)	(1.958)	(3.108)	(3.459)	(0.446)	(0.359)	(0.320)	(0.421)
const.	63.158***	50.963***	51.279***	35.552***	5.877***	5.089***	5.123***	4.591***
	(2.465)	(1.403)	(1.476)	(2.016)	(0.292)	(0.229)	(0.228)	(0.255)
Obs.	820	820	820	820	790	790	790	790
Clusters	82	82	82	82	79	79	79	79
R^2	0.004	0.001	0.003	0.008	0.001	0.000	0.002	0.000

Notes: OLS regressions of the observers' beliefs for the different possible success occurrences. *IPE* denotes the rounds where the observed agent received feedback after the first period. *endogenous* denotes the rounds where another principal (instead of the computer) chose whether IPE is played. Standard-errors are clustered at the principal level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

behavior. The results indicate that observers do not significantly adjust their assessments based on the existence or the endogeneity of IPE. The independence of the updating behavior on the endogeneity aligns with the finding in the previous section, where it was shown that agents also did not behave differently in exogenous versus endogenous rounds. The regressions in Table 5 also reveal that observers do not seem to base their beliefs on whether IPE was implemented. At the effort cost used in the experiment and under the symmetric mixing equilibrium of NoIPE, the equilibrium posteriors for a given outcome differ across regimes by at most about 0.05 on the unit type scale (five points in the Abstract task and half a point in the Quiz task), so the absence of a significant regime effect on beliefs is not sharply at odds with the theory. Under either pure timing equilibrium of NoIPE, the posterior after a success in the period that active types do not use lies about 0.3 below its IPE counterpart. The sharpest belief prediction is instead within the IPE regime: the conditional strategy makes a first-period success more informative than a second-period success, with limiting posteriors of approximately 0.82 after (Success, NoSuccess) and 0.78 after (NoSuccess, Success). The data do not show this asymmetry, as beliefs after the two one-success outcomes do not differ significantly under IPE. Under NoIPE, the ranking of the two one-success posteriors depends on which timing equilibrium is selected (higher after a first-period success in the period-1 equilibrium, equal under mixing, and lower in the period-2 equilibrium), so the absence of a difference there is consistent with the symmetric mixing equilibrium.

6.3 Principals' IPE Choices

In half of the rounds, principals had the option to decide whether the agent they were matched with received feedback about their first-period performance. Principals chose to provide feedback in 72.44% of the cases in the Abstract task and 70.13% in the Quiz task. Table 1 presents regressions of total effort on IPE. The results reveal that IPE significantly increases total effort provision. Since, for a given type, the probability of success rises with effort, and types are balanced across regimes by design, higher effort translates into a higher expected number of successes, which determine the principal's round payoff.[13]

In the post-experimental questionnaire, principals were asked the open-ended question: *"What motivated your decision on whether Player A receives feedback about the success in period 1?"* Responding to this question was voluntary and unincentivized, so definitive conclusions cannot be drawn from the responses. However, the answers might still offer valuable insights into the principals' motivations. Table B3 in Appendix B provides a detailed list of the responses. Notable categories include the desire to be kind to the agent, random decision-making, and the statement of mechanisms for why IPE does (or does not) lead to greater success. Most responses fall into the last category, proposing mechanisms such as reciprocity, changes in the agents' risk attitudes, and increased motivation when positive results are known. In the Quiz treatments, the last mechanism is indeed supported by the agents' behavior, as shown in Table 3.

6.4 Knowledge in the Quiz Task

To be a suitable task for the given research question, the Quiz task must meet the following criteria: i) the knowledge tested by the questions during the game should correlate with the participants' initial knowledge (which determines the type that observers estimate); ii) knowledge should remain relatively stable throughout the experiment; iii) knowledge should not be influenced by the presence or absence of IPE; and iv) the task should effectively differentiate between participants' knowledge levels, while being easy enough for participants to answer some questions, yet challenging enough to avoid a ceiling effect. The following analyses confirm that all four criteria are met.[14]

Table B1 in Appendix B presents a regression of the number of correct answers provided during the game on the number of correct answers given at the beginning of the game

[13]The payoff was either 0, 50, or 100 points, depending on whether success was realized zero, one, or two times in a round.

[14]It is important to note that, by experimental design, participants were required to answer the questions regardless of whether IPE was implemented or whether the cost was paid.

(which determines the agent's type). The strong and significant correlation between the agent's type and correct answers across both periods confirms that criterion (i) is fulfilled. The insignificance of the round coefficient suggests that criterion (ii) is also met. The *IPE* coefficient is weakly significant for the number of correct answers in period 2 and insignificant for period 1, suggesting that while criterion (iii) is not perfectly satisfied, it is arguably met to a sufficient extent.

Figure B1 in Appendix B plots the number of correct answers throughout the game. The graph demonstrates that criterion (iv) is also largely fulfilled, as there is no evident floor or ceiling effect. The cutoff of 7 out of 10 correct answers for success appears to be reasonably set, with 36.71% of instances reaching or exceeding this threshold.

6.5 Potential Reasons for Deviations from Hypothesized Behavior

The PBE analysis predicts that for the cost level in the experiment, overall effort should be higher with IPE than without. As shown above, the data is in line with that prediction. However, under IPE, lower second-period effort is anticipated from agents who, after learning of their first-period success, no longer need to exert effort to distinguish themselves from lower types. As shown in Table 3, this effect is small and insignificant in the Abstract task, and reversed in the Quiz task. In the Abstract task, potential reasons for this deviation could include a lack of understanding of the incentives or social preferences toward the principal. Data from the post-experimental questionnaire can be used to explore these possibilities.[15] High school math grades, used as a proxy for cognitive ability, can provide insight into participants' understanding of the experimental instructions and incentives.

In addition to reporting their math grades, participants responded to a 7-point Likert scale question: *"How much emphasis did you place on the payoffs of the other players in your decision-making?"* Table B2 in Appendix B regresses second-period effort for participants who learned of their first-period success on their questionnaire responses, with *Fairness* representing their answers to the fairness question. The analysis reveals that math grades do not significantly correlate with effort provision, suggesting that a lack of understanding of the incentives is unlikely to explain the deviation. However, for the Abstract treatments, responses to the fairness question significantly correlate with second-period effort provision after first-period success, indicating that social preferences may explain the deviation from theoretical predictions. In the Quiz treatments, where intrinsic motivation to succeed is likely higher, no such correlation is observed. Neither gender nor the endogeneity of IPE significantly correlates with second-period effort provision after first-period success. For the

[15] As the questionnaire was unincentivized and self-reported, these findings should be interpreted cautiously.

Quiz treatments, Table 3 reveals significantly higher second-period effort after communicated first-period success. As fairness does not seem to drive these observations (Table B2), a natural candidate is the nature of the private information in this task. In the Quiz task, an agent's type is his own quiz performance, which he knows only imperfectly, whereas the model assumes that the agent knows his type exactly. A communicated first-period success is then an informative signal about the agent's own ability, so an agent who updates upward also raises his assessment of the return to second-period effort, and a failure has the opposite effect. This channel, which the model excludes by assumption, generates exactly the observed positive response to early success and is consistent with the interpretation that communicated success increases confidence in one's ability to succeed again.

6.5.1 Best Reply Behavior

The game is interactive, and while the theoretical model assumes that the evaluating observer engages in Bayesian updating, it remains uncertain whether actual observer behavior aligns with the incentive structures predicted by theory. According to the theoretical predictions in Propositions 1 and 2, the following behavior is expected in the Abstract treatments with a cost level of 30 points: Without IPE, the equilibrium threshold at this cost is $\bar{\theta}_A \approx 0.636$, so agents with a type of 63 or lower should refrain from exerting effort in either period, and agents with a type of 64 or higher should exert effort exactly once. No agent is predicted to benefit from exerting effort in both periods.

Abstract Treatment Given observers' empirical beliefs, exerting no effort yields an average payoff of 35.17 points. Exerting effort only in the first period yields an expected payoff of $\theta_i \cdot 50.53 + (1-\theta_i) \cdot 35.17 - 30$ points, while exerting effort only in the second period results in an expected payoff of $\theta_i \cdot 50.76 + (1-\theta_i) \cdot 35.17 - 30$ points. Exerting effort in both periods yields an expected payoff of $\theta_i^2 \cdot 62.45 + \theta_i \cdot (1-\theta_i) \cdot 50.53 + (1-\theta_i) \cdot \theta_i \cdot 50.76 + (1-\theta_i)^2 \cdot 35.17 - 60$ points. These calculations reveal that, regardless of the agent's type, not exerting effort is the empirical best response.

With IPE, the equilibrium thresholds of Proposition 2 at this cost are $\bar{\theta}_A \approx 0.604$ and $\bar{\theta}_B \approx 0.655$: agents with a type of 60 or lower should not exert effort, agents with a type between 61 and 65 should exert effort only in the first period, and agents with a type of 66 or higher are predicted to use the conditional strategy of exerting effort in the first period and exerting effort in the second period only if the first period did not result in success. However, based on the observer's empirical beliefs (30.17 points for (No Success, No Success), 50.86 points for (Success, No Success), 53.01 points for (No Success, Success), and 65.87 points for (Success, Success)), the best empirical reaction remains to not exert effort, regardless of the

agent's type.

Using the actual choice data of the agents, one can determine what the observers' *correct beliefs* should be in the Abstract treatments. In the NoIPE rounds, no success in either period was associated with agents having an average type of 37.58. Success in the first period only came from agents with an average type of 70.41, while success in the second period only was associated with an average type of 61.21. Success in both periods resulted from agents with an average type of 82.21. In the IPE rounds, no success came from agents with an average type of 34.45, success in the first period only from agents with an average type of 68.22, success in the second period only from agents with an average type of 61.65, and success in both periods from agents with an average type of 82.39.

Comparing these actual averages with the observed beliefs reveals that observers were generally too pessimistic. While their estimates were relatively accurate after observing no success in either period, the true types of agents who achieved success in one or both periods were, on average, higher than the observers' beliefs reflected. Consequently, due to these low beliefs after success, the true incentives to exert effort in the experiment were lower than theory would predict.

Quiz Treatment In the Quiz treatments, without IPE, observers' beliefs were 4.62 (NN), 5.09 (SN), 5.15 (NS), and 5.96 (SS). With IPE, the average beliefs were 4.61 (NN), 5.13 (SN), 5.26 (NS), and 5.96 (SS). Given that exerting effort costs 30 points and an increase in beliefs by one unit results in a 10-point payoff increase, these belief differences were also insufficient to make exerting effort the best response.

The *correct* beliefs, based on actual behavior, would have been: without IPE, 5.10 (NN), 6.53 (SN), 6.36 (NS), and 7.71 (SS); with IPE, 5.30 (NN), 6.19 (SN), 6.42 (NS), and 7.23 (SS). Thus, also in the Quiz treatments, observers' beliefs were overly pessimistic compared to the correct beliefs derived from actual agent behavior.

6.5.2 Behavior over Time and Learning

The experiment consists of 10 rounds. At the end of each round, the principal learns how often success occurred, while the agent receives information about the first-period success (in IPE rounds), second-period success, and the observers' belief about the agent's type. The observer, however, does not receive feedback about the accuracy of their beliefs until the end of the experiment.

To examine how agents' effort provision changes over time, Figures B2 and B3 in Appendix B show the first- and second-period effort provision for the Abstract (Panels a) and Quiz (Panels b) treatments over the course of the experiment. These figures reveal several

trends: In the Quiz treatment, agents' willingness to pay 30 points declines over time. One intuitive explanation is that agents learn that the observers' beliefs are not differentiated enough to justify the 30-point cost. Additionally, agents may realize that the quiz questions are more difficult than initially expected, taking time to recognize the limits of their knowledge of countries' capitals. The figures also show that the difference in first-period effort between NoIPE and IPE in the Quiz treatments emerges in the middle and late stages of the experiment. While effort levels are similar in the first three rounds, the decline in first-period effort without IPE is sharper than with IPE.

For the Abstract treatments, the time trends also vary with the regime. In rounds with IPE, a clear downward slope in effort provision is observed. Without IPE, the effort pattern remains relatively stable but at a lower level.

7 Discussion and Conclusion

Investigating interim performance evaluations under a career-concern framework reveals a key theoretical result: whether IPE increases or decreases total effort provision depends on the cost of effort, and the two regimes' total-effort curves cross exactly once. At the effort cost implemented in the experiment, the model predicts that IPE strictly increases total effort provision. This prediction is confirmed by the data: both first-period and total effort provision are significantly higher under IPE. The main deviation from the model's prediction concerns the mechanism rather than the direction: the model predicts that agents who succeed early reduce effort provision in the second period, but this decline is at most weakly present in the Abstract task and reversed in the Quiz task. IPE therefore improves performance to a greater extent than the model predicts. This pattern holds across two different settings, one with low and one with high intrinsic motivation, and regardless of whether the IPE is implemented exogenously or chosen by the principal.

A second empirical lesson concerns the calibration of observers' beliefs. Comparing observers' stated beliefs with the posteriors implied by agents' actual choices shows that observers were systematically too pessimistic following success, so that the reputational reward for a good outcome fell short of its Bayesian benchmark. Compressed assessments are a well-documented feature of performance evaluations. Kusterer and Sliwka (2025) show experimentally that ratings respond less to performance when evaluators receive coarser information, and that paying evaluators for accuracy reduces the leniency that arises when ratings determine the worker's pay. Both conditions hold in the present design: a binary success outcome is a coarse signal of ability, and observers were paid for accuracy, which fits beliefs that were compressed and, on average, below the agents' true types. Given the beliefs

agents actually faced, withholding effort was in fact the monetary best response across the entire type range.

That agents nonetheless exerted substantial effort points to motives beyond the monetary return, such as intrinsic satisfaction from signaling competence, image concerns toward the observer and principal, overconfidence about one's ability in the Quiz task, or slow convergence toward the best response, consistent with the decline of effort over rounds shown in Figures B2 and B3. These forces sustain effort even when the pecuniary return to signaling is thin. For organizations, this cuts two ways: signaling incentives can operate even when evaluators reward success only weakly, but managers who presume well-calibrated reputational rewards may overstate the monetary force of their scheme.

More broadly, the results carry a practical message for any organization that controls the flow of performance information to its members. Implementing an interim evaluation is essentially costless to the principal, yet in the experiment it raised effort and, through the resulting successes, the principal's expected payoff. Field evidence on a related low-cost practice, regular performance review conversations, points in the same direction (Manthei et al., 2023). The natural concern that revealing early success leads high performers to coast (the stopping behavior that the equilibrium prescribes for high types at the implemented cost, and the source of the predicted effort loss at low costs) does not appear as a decline in either task context; in the knowledge-based task, second-period effort was in fact higher after an early success. Managers who tell employees that they performed well therefore have little reason to fear that this news will make them ease off.

The findings can also apply to other settings. Consider a professor hiring a new PhD student and aiming to maximize research output. In addition to well-known effects—such as identifying students who may not thrive in academia, providing guidance before final projects, and incurring time costs for evaluations—this paper introduces a new factor. If PhD students are better informed about their own ability or opportunity costs than their supervisor, the supervisor can better assess the quality of research output, and students value their supervisor's belief in their competence and career prospects, then an interim evaluation is predicted to increase effort provision. Concerns that withholding information reduces effort through reciprocal behavior, though, are not supported by the data.

Several limitations qualify these conclusions. The experiment implements a single effort cost, $c = 0.3$, which lies above the threshold c^* at which the effort ranking reverses; the data therefore speak to the region in which the model predicts IPE to dominate, but not to the low-cost region in which NoIPE is predicted to elicit more total effort. Testing the crossing itself, and locating c^* empirically, would require varying the cost of effort across treatments. The model also abstracts from richer features of real evaluation environments:

effort is binary, the horizon is two periods, and evaluations are objective and truthful by construction. Agents in the model also respond to career concerns only, whereas interim evaluations in firms usually coexist with explicit incentives. The two can interact: in the retail field experiment of Manthei et al. (2023), performance reviews raised profits only in stores whose managers did not also receive performance pay. As is standard in laboratory work, the subject pool consists of students, so external validity to workplace settings remains to be established. Finally, the stopping effect is absent in both tasks, and in the Quiz task part of the positive association between early success and later effort also appears in rounds without feedback. Separating the response to feedback from what agents already know about their own performance would require a design that varies the nature of agents' private information more directly.

These limitations suggest several directions for future work. Varying the effort cost across treatments would allow the two effort curves to be traced out and the predicted reversal at c^* to be tested directly, a sharper test of the theory than a single cost level permits. Extending the design to subjective or potentially misreported evaluations would link the present objective-feedback setting to the large literature on subjective performance evaluation. And because agents respond to feedback even when evaluators' rewards are weakly calibrated, manipulating the evaluators' own information or incentives—for example, by giving observers feedback on the accuracy of their past beliefs—could help separate the strategic component of the effort response from the intrinsic one.

A final direction concerns the time horizon. While the current model is easily applied to limited-time contracts in which employees may later seek a reference from their former principal, the implications for longer-term employment with recurring IPEs are less clear. The conditional-strategy logic suggests that, under recurring IPEs, agents would exert effort until their record is sufficiently convincing and stop thereafter. However, the point at which this occurs—how many successes are needed before the reputational benefit of further effort no longer justifies the cost—depends on both effort costs and the agent's type, making the multi-period case richer than a straightforward extension of the two-period model. Formally characterizing equilibria with recurring IPEs and testing the resulting predictions experimentally would be a particularly valuable next step.

References

Alós-Ferrer, C., J. García-Segarra, and A. Ritschel (2018). Performance curiosity. *Journal of Economic Psychology 64*, 1–17.

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

Azmat, G., M. Bagues, A. Cabrales, and N. Iriberri (2019). What you don't know... can't hurt you? a natural field experiment on relative performance feedback in higher education. *Management Science 65*(8), 3714–3736.

Azmat, G. and N. Iriberri (2010). The importance of relative performance feedback information: Evidence from a natural experiment using high school students. *Journal of Public Economics 94*(7-8), 435–452.

Battigalli, P. and M. Dufwenberg (2022). Belief-dependent motivations and psychological game theory. *Journal of Economic Literature 60*(3), 833–882.

Bernstein, E. and S. X. Li (2026). The performance effects of giving front-line employees direct access to performance data and thereby limiting the supervisor's feedback-intermediation role: evidence from a field experiment. *Management Science 72*(2), 805–835.

Bernstein, E. S. (2012). The transparency paradox: A role for privacy in organizational learning and operational control. *Administrative Science Quarterly 57*(2), 181–216.

Bock, O., I. Baetge, and A. Nicklisch (2014). hroot: Hamburg registration and organization online tool. *European Economic Review 71*, 117–120.

Chen, B. R. (2015). Subjective performance feedback, ability attribution, and renegotiation-proof contracts. *Journal of Economic Behavior & Organization 117*, 155–174.

Chen, B. R. and Y. Stephen Chiu (2013). Interim performance evaluation in contract design. *The Economic Journal 123*(569), 665–698.

Cho, I.-K. and D. M. Kreps (1987). Signaling games and stable equilibria. *The Quarterly Journal of Economics 102*(2), 179–221.

Dewatripont, M., I. Jewitt, and J. Tirole (1999). The economics of career concerns, part ii: Application to missions and accountability of government agencies. *The Review of Economic Studies 66*(1), 199–217.

Di Pei, H. (2016). Reputation with strategic information disclosure. *mimeo*.

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

Ely, J. C., G. Georgiadis, S. Khorasani, and L. Rayo (2023). Optimal feedback in contests. *The Review of Economic Studies 90*(5), 2370–2394.

Engelmann, D. and U. Fischbacher (2009). Indirect reciprocity and strategic reputation building in an experimental helping game. *Games and Economic Behavior 67*(2), 399–407.

Eriksson, T., A. Poulsen, and M. C. Villeval (2009). Feedback and incentives: Experimental evidence. *Labour Economics 16*(6), 679–688.

Falk, A. and U. Fischbacher (2006). A theory of reciprocity. *Games and Economic Behavior 54*(2), 293–315.

Falk, A. and F. Zimmermann (2023). Attention and dread: Experimental evidence on preferences for information. *Management Science*.

Fehr, E., U. Fischbacher, and S. Gächter (2002). Strong reciprocity, human cooperation, and the enforcement of social norms. *Human nature 13*, 1–25.

Fehr, E., S. Gächter, and G. Kirchsteiger (1997). Reciprocity as a contract enforcement device: Experimental evidence. *Econometrica: Journal of the Econometric Society*, 833–860.

Fehrler, S. and N. Hughes (2018). How transparency kills information aggregation: theory and experiment. *American Economic Journal: Microeconomics 10*(1), 181–209.

Fehrler, S. and M. Janas (2021). Delegation to a group. *Management Science 67*(6), 3714–3743.

Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics 10*, 171–178.

Fuchs, W. (2007). Contracting with repeated moral hazard and private evaluations. *American Economic Review 97*(4), 1432–1448.

Fuchs, W. (2015). Subjective evaluations: Discretionary bonuses and feedback credibility. *American Economic Journal: Microeconomics 7*(1), 99–108.

Ganguly, A. and J. Tasoff (2017). Fantasy and dread: The demand for information and the consumption utility of the future. *Management Science 63*(12), 4037–4060.

Gibbons, R. and K. J. Murphy (1992). Optimal incentive contracts in the presence of career concerns: Theory and evidence. *Journal of political Economy 100*(3), 468–505.

Golman, R. and G. Loewenstein (2018). Information gaps: A theory of preferences regarding the presence and absence of information. *Decision 5*(3), 143.

Goltsman, M. and A. Mukherjee (2011). Interim performance feedback in multistage tournaments: The optimality of partial disclosure. *Journal of Labor Economics 29*(2), 229–265.

Gürtler, O. and C. Harbring (2010). Feedback in tournaments under commitment problems: Experimental evidence. *Journal of Economics & Management Strategy 19*(3), 771–810.

Hansen, S. E. (2013). Performance feedback with career concerns. *The Journal of Law, Economics, & Organization 29*(6), 1279–1316.

Heinsalu, S. (2018). Dynamic noisy signaling. *American Economic Journal: Microeconomics 10*(2), 225–249.

Holmström, B. (1999). Managerial incentive problems: A dynamic perspective. *The Review of Economic Studies 66*(1), 169–182.

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

Kaya, A. (2009). Repeated signaling games. *Games and Economic Behavior 66*(2), 841–854.

Klaas, B. S., J. A. McClendon, and T. W. Gainey (2001). Outsourcing hr: The impact of organizational characteristics. *Human Resource Management 40*(2), 125–138.

Kuhnen, C. M. and A. Tymula (2012). Feedback, self-esteem, and performance in organizations. *Management Science 58*(1), 94–113.

Kusterer, D. J. and D. Sliwka (2025). Social preferences and the informativeness of subjective performance evaluations. *Management Science 71*(5), 4028–4048.

Lang, M. (2019). Communicating subjective evaluations. *Journal of Economic Theory 179*, 163–199.

Levy, G. (2007). Decision-making procedures for committees of careerist experts. *American Economic Review 97*(2), 306–310.

Lizzeri, A., M. A. Meyer, and N. Persico (2002). The incentive effects of interim performance evaluations. CARESS Working Paper 02-09, University of Pennsylvania.

MacLeod, W. B. (2003). Optimal contracting with subjective evaluation. *American Economic Review 93*(1), 216–240.

Maestri, L. (2012). Bonus payments versus efficiency wages in the repeated principal-agent model with subjective evaluations. *American Economic Journal: Microeconomics 4*(3), 34–56.

Manthei, K., D. Sliwka, and T. Vogelsang (2023). Talking about performance or paying for it? a field experiment on performance reviews and incentives. *Management Science 69*(4), 2198–2216.

Mas, A. and E. Moretti (2009). Peers at work. *American Economic Review 99*(1), 112–145.

Milgrom, P. and J. Roberts (1986). Relying on the information of interested parties. *The RAND Journal of Economics*, 18–32.

Nielsen, K. (2020). Preferences for the resolution of uncertainty and the timing of information. *Journal of Economic Theory 189*, 105090.

Noldeke, G. and E. Van Damme (1990). Signalling in a dynamic labour market. *The Review of Economic Studies 57*(1), 1–23.

Ouwersloot, H., P. Nijkamp, and P. Rietveld (1998). Errors in probability updating behaviour: Measurement and impact analysis. *Journal of Economic Psychology 19*(5), 535–563.

Prat, A. (2005). The wrong kind of transparency. *American economic review 95*(3), 862–877.

Riley, J. G. (1979). Informational equilibrium. *Econometrica: Journal of the Econometric Society*, 331–359.

Riley, J. G. (2001). Silver signals: Twenty-five years of screening and signaling. *Journal of Economic Literature 39*(2), 432–478.

Roddie, C. (2012). Signaling and reputation in repeated games, i: Finite games. *Available at SSRN 1994378*.

Selten, R. (1967). Die strategiemethode zur erforschung des eingeschränkt rationale verhaltens im rahmen eines oligopolexperiments. *Beiträge zur experimentellen Wirtschaftsforschung*, 136.

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

Spence, M. (1978). Job market signaling. In *Uncertainty in economics*, pp. 281–306. Elsevier.

Suvorov, A. and J. Van de Ven (2009). Discretionary rewards as a feedback mechanism. *Games and Economic Behavior 67*(2), 665–681.

Villeval, M. C. (2020). Performance feedback and peer effects. *Handbook of labor, human resources and population economics*, 1–38.

Visser, B. and O. H. Swank (2007). On committees of experts. *The Quarterly Journal of Economics 122*(1), 337–372.

Zhang, M. (2019). When the principal knows better than the agent: Subjective evaluations as an optimal disclosure mechanism. *Journal of Economics & Management Strategy 28*(4), 631–655.

Zizzo, D. J., S. Stolarz-Fantino, J. Wen, and E. Fantino (2000). A violation of the monotonicity axiom: Experimental evidence on the conjunction fallacy. *Journal of Economic Behavior & Organization 41*(3), 263–276.

A Theoretical Appendix

This appendix derives the equilibria stated in Propositions 1 and 2 and proves Proposition 3. Appendix A1 introduces the notation, discusses the pooling equilibrium, and collects results that hold in both regimes. Appendices A2 and A3 derive the selected active-signaling equilibria under NoIPE and IPE, and Appendix A4 proves Proposition 3.

A1 Preliminaries

Notation. As in the main text, I write effort profiles as $(e_1; e_2)$ and outcomes as (S_1, S_2). For an outcome $(S_1, S_2) = (a, b)$, let

$$\mu_{ab} \equiv \hat{\theta}(a, b) = E[\theta_i \mid S_1 = a, S_2 = b]$$

denote the common posterior belief of the principal and the observer. Since $\theta_i \sim \mathcal{U}[0, 1]$, Bayes' rule gives

$$\mu_{ab} = \frac{\int_0^1 \theta_i \Pr(S_1 = a, S_2 = b \mid \theta_i)\, d\theta_i}{\int_0^1 \Pr(S_1 = a, S_2 = b \mid \theta_i)\, d\theta_i} \tag{A.1}$$

for every outcome that occurs with positive probability.

Pooling equilibrium. Suppose that all types choose low effort in both periods, $(e_1; e_2) = (\varepsilon; \varepsilon)$. A type θ_i then succeeds in each period with probability $\varepsilon\theta_i$, so every outcome occurs with positive probability: no success with probability $(1 - \varepsilon\theta_i)^2$, a success in the first period only or in the second period only with probability $\varepsilon\theta_i(1 - \varepsilon\theta_i)$ each, and two successes with probability $\varepsilon^2\theta_i^2$. By (A.1), the posterior beliefs are

$$\mu_{00} = \frac{\frac{1}{2} - \frac{2\varepsilon}{3} + \frac{\varepsilon^2}{4}}{1 - \varepsilon + \frac{\varepsilon^2}{3}}, \qquad \mu_{10} = \mu_{01} = \frac{\frac{1}{3} - \frac{\varepsilon}{4}}{\frac{1}{2} - \frac{\varepsilon}{3}}, \qquad \mu_{11} = \frac{3}{4}.$$

A deviation to high effort in one period raises the success probability in that period from $\varepsilon\theta_i$ to θ_i. Since $\mu_{10} = \mu_{01}$, the gross gain from this deviation is the same for either period:

$$G_1(\theta_i; \varepsilon) = (1 - \varepsilon)\theta_i \left[(1 - \varepsilon\theta_i)(\mu_{10} - \mu_{00}) + \varepsilon\theta_i(\mu_{11} - \mu_{01})\right]. \tag{A.2}$$

The gross gain from deviating to high effort in both periods is

$$\begin{aligned} G_2(\theta_i; \varepsilon) = {} & \theta_i^2\mu_{11} + 2\theta_i(1 - \theta_i)\mu_{10} + (1 - \theta_i)^2\mu_{00} \\ & - \left[(\varepsilon\theta_i)^2\mu_{11} + 2\varepsilon\theta_i(1 - \varepsilon\theta_i)\mu_{10} + (1 - \varepsilon\theta_i)^2\mu_{00}\right]. \end{aligned} \tag{A.3}$$

Under NoIPE, these are the only deviations. The pooling strategy is therefore an equilibrium if and only if $c \geq \bar{c}^N_{\text{pool}}(\varepsilon)$, where

$$\bar{c}^N_{\text{pool}}(\varepsilon) \equiv \max\left\{\max_{\theta_i\in[0,1]} G_1(\theta_i;\varepsilon),\ \frac{1}{2}\max_{\theta_i\in[0,1]} G_2(\theta_i;\varepsilon)\right\}. \tag{A.4}$$

Under IPE, the agent observes S_1 before choosing e_2, so pooling also requires that no type wants to exert high second-period effort after either first-period outcome. After $S_1 = s$, high second-period effort raises the expected posterior by $(1-\varepsilon)\theta_i(\mu_{s1} - \mu_{s0})$. No type gains from these deviations if and only if $c \geq (1-\varepsilon)(\mu_{01} - \mu_{00})$ and $c \geq (1-\varepsilon)(\mu_{11} - \mu_{10})$. In that case, every type exerts low second-period effort after either outcome, and the gain from high first-period effort is $G_1(\theta_i;\varepsilon)$, as under NoIPE. Hence, under IPE, the pooling strategy is an equilibrium if and only if $c \geq \bar{c}^I_{\text{pool}}(\varepsilon)$, where

$$\bar{c}^I_{\text{pool}}(\varepsilon) \equiv \max\left\{\max_{\theta_i\in[0,1]} G_1(\theta_i;\varepsilon),\ (1-\varepsilon)(\mu_{01}-\mu_{00}),\ (1-\varepsilon)(\mu_{11}-\mu_{10})\right\}. \tag{A.5}$$

As $\varepsilon \to 0$, the posteriors converge to $\mu_{00} \to \frac{1}{2}$ and $\mu_{10} = \mu_{01} \to \frac{2}{3}$, while $\mu_{11} = \frac{3}{4}$ for every ε. The one-period gain converges to $G_1(\theta_i;\varepsilon) \to \theta_i/6$, whose maximum is $\frac{1}{6}$, and the largest two-period gain per unit of effort cost converges to

$$\frac{1}{2}\max_{\theta_i\in[0,1]}\left[\frac{3}{4}\theta_i^2 + \frac{4}{3}\theta_i(1-\theta_i) + \frac{1}{2}(1-\theta_i)^2 - \frac{1}{2}\right] = \frac{1}{2}\max_{\theta_i\in[0,1]}\left(\frac{\theta_i}{3} - \frac{\theta_i^2}{12}\right) = \frac{1}{8}.$$

Under IPE, the largest second-period gains converge to $\frac{1}{6}$ after a failure and to $\frac{1}{12}$ after a success. Hence

$$\bar{c}^N_{\text{pool}}(\varepsilon) \to \frac{1}{6} \qquad \text{and} \qquad \bar{c}^I_{\text{pool}}(\varepsilon) \to \frac{1}{6} \qquad \text{as } \varepsilon \to 0.$$

Numerically, both thresholds are strictly decreasing in ε. For example, $\bar{c}^N_{\text{pool}} \approx 0.137$ and $\bar{c}^I_{\text{pool}} \approx 0.153$ at $\varepsilon = 0.2$, and $\bar{c}^N_{\text{pool}} \approx 0.089$ and $\bar{c}^I_{\text{pool}} \approx 0.116$ at $\varepsilon = 0.5$. Under IPE, the binding deviation is high second-period effort after a first-period failure, so the pooling equilibrium requires a higher cost under IPE than under NoIPE. Both thresholds tend to zero as $\varepsilon \to 1$, because effort then no longer affects the distribution of outcomes and every deviation gain vanishes. In particular, the pooling equilibrium exists in both regimes for all $c \geq \frac{1}{6}$ and every $\varepsilon \in (0,1)$.

Selection. I call an equilibrium an active-signaling equilibrium if a positive mass of types exerts high effort in at least one period. For a wide range of costs, the pooling equilibrium coexists with active-signaling equilibria. Following the main text, at each cost c I select

the active-signaling equilibrium with the highest expected total high effort, and I select the pooling equilibrium only if no active-signaling equilibrium exists. Where several active-signaling equilibria attain the highest total effort, I report all of them.

Limit convention. The remaining results concern active-signaling equilibria in the limit $\varepsilon \to 0$. For a given strategy profile, the probability of each outcome is a polynomial in ε. If the limiting profile generates an outcome with positive probability, the posterior after this outcome converges to the posterior computed at $\varepsilon = 0$. From here on, I therefore write $e_t = 0$ for low effort, so that $(0;0)$ denotes the limit of $(\varepsilon;\varepsilon)$, and I compute likelihoods and posteriors at $\varepsilon = 0$. The perturbation matters in two places. First, some equilibria generate an outcome only through low-effort success, which has probability zero in the limit. The posterior after such an outcome is the limit of the Bayesian posteriors of the perturbed game, and I compute it as this limit. Where an incentive argument holds for every possible value of such a posterior, I use only the bound $\mu_{ab} \leq 1$. Second, the perturbation determines which timing profiles are equilibria in the one-threshold region under NoIPE (Appendix A2).

Strategies and expected utilities. In the limit, a type θ_i that exerts high effort in a period succeeds with probability θ_i, and a type that exerts low effort fails. Given posteriors $(\mu_{00}, \mu_{10}, \mu_{01}, \mu_{11})$, the expected utilities of the strategies $(0;0)$, $(1;0)$, $(0;1)$, and $(1;1)$ are

$$U_0(\theta_i) = \mu_{00}, \tag{A.6}$$

$$U_1(\theta_i) = \theta_i\mu_{10} + (1-\theta_i)\mu_{00} - c, \tag{A.7}$$

$$U_2(\theta_i) = \theta_i\mu_{01} + (1-\theta_i)\mu_{00} - c, \tag{A.8}$$

$$U_{12}(\theta_i) = \theta_i^2\mu_{11} + \theta_i(1-\theta_i)(\mu_{10}+\mu_{01}) + (1-\theta_i)^2\mu_{00} - 2c. \tag{A.9}$$

These strategies are available in both regimes, and under NoIPE the agent may also mix over them. Under IPE, the agent can in addition condition second-period effort on S_1. The central strategy of this kind is the conditional strategy $(1; 1-S_1)$: high effort in the first period, and high effort in the second period if and only if the first period ended without success. Its expected utility is

$$U_{\text{cond}}(\theta_i) = \theta_i\mu_{10} + (1-\theta_i)\left[\theta_i\mu_{01} + (1-\theta_i)\mu_{00} - c\right] - c. \tag{A.10}$$

The derivations below repeatedly use the following differences, which follow directly from (A.6)–(A.10):

$$U_1(\theta_i) - U_0(\theta_i) = \theta_i(\mu_{10} - \mu_{00}) - c, \tag{A.11}$$

$$U_{12}(\theta_i) - U_1(\theta_i) = \theta_i^2(\mu_{11} - \mu_{10}) + \theta_i(1-\theta_i)(\mu_{01} - \mu_{00}) - c, \tag{A.12}$$

$$U_{\text{cond}}(\theta_i) - U_1(\theta_i) = (1-\theta_i)\left[\theta_i(\mu_{01} - \mu_{00}) - c\right], \tag{A.13}$$

$$U_{12}(\theta_i) - U_{\text{cond}}(\theta_i) = \theta_i\left[\theta_i(\mu_{11} - \mu_{10}) - c\right]. \tag{A.14}$$

Exchanging the roles of the two periods gives the corresponding expressions for $U_2 - U_0$ and $U_{12} - U_2$.

Posterior building blocks. In the limit, every likelihood below equals 1, $1-\theta_i$, $(1-\theta_i)^2$, θ_i, $\theta_i(1-\theta_i)$, θ_i^2, or 0 on intervals of types, possibly multiplied by a mixing probability or, for outcomes that arise only through low-effort success, by a power of ε. A factor that is common to all intervals cancels in (A.1), and most posteriors below then take one of three forms. For $0 \le a \le b \le 1$, define

$$\begin{aligned} m_0(a,b) &\equiv \frac{\int_0^a \theta_i\, d\theta_i + \int_a^b \theta_i(1-\theta_i)\, d\theta_i + \int_b^1 \theta_i(1-\theta_i)^2\, d\theta_i}{\int_0^a d\theta_i + \int_a^b (1-\theta_i)\, d\theta_i + \int_b^1 (1-\theta_i)^2\, d\theta_i} \\ &= \frac{1 + 4a^3 + 4b^3 - 3b^4}{2\left(2 + 3a^2 + 3b^2 - 2b^3\right)}, \end{aligned} \tag{A.15}$$

$$m_1(a,b) \equiv \frac{\int_a^b \theta_i^2\, d\theta_i + \int_b^1 \theta_i^2(1-\theta_i)\, d\theta_i}{\int_a^b \theta_i\, d\theta_i + \int_b^1 \theta_i(1-\theta_i)\, d\theta_i} = \frac{1 - 4a^3 + 3b^4}{2\left(1 - 3a^2 + 2b^3\right)} \qquad (a < 1), \tag{A.16}$$

and for $0 \le x < 1$ define

$$m_2(x) \equiv \frac{\int_x^1 \theta_i^3\, d\theta_i}{\int_x^1 \theta_i^2\, d\theta_i} = \frac{3\left(1 - x^4\right)}{4\left(1 - x^3\right)} = \frac{3(1+x)(1+x^2)}{4(1+x+x^2)}, \tag{A.17}$$

with $m_2(1) \equiv \lim_{x \to 1} m_2(x) = 1$. The function $m_0(a,b)$ is the posterior after no success when types up to a never succeed, types in $(a,b]$ fail with probability $1-\theta_i$, and types above b fail with probability $(1-\theta_i)^2$. The function $m_1(a,b)$ is the posterior after an outcome with exactly one success that types in $(a,b]$ generate with probability θ_i and types above b generate with probability $\theta_i(1-\theta_i)$. The function $m_2(x)$ is the posterior after two successes that only types above x generate, with probability proportional to θ_i^2. The following special

cases appear repeatedly:

$$\begin{aligned} m_0(a,1) &= \frac{1+2a^3}{3(1+a^2)}, & m_1(a,1) &= \frac{2(1+a+a^2)}{3(1+a)}, \\ m_1(b,b) &= \frac{1+2b+3b^2}{2(1+2b)}, & m_1(0,a) &= \frac{1+3a^4}{2(1+2a^3)}. \end{aligned} \tag{A.18}$$

The function $m_1(a,1)$, the mean type among types above a weighted by their success probability, is strictly increasing in a, with derivative $2a(a+2)/[3(1+a)^2]$.

An upper bound on the cost of active signaling. The following bound applies to both regimes.

Lemma A1. *In every active-signaling equilibrium of either regime, the infimum $\underline{\theta}$ of the types who exert high effort satisfies*

$$c \le \underline{\theta}\Big(1-\frac{\underline{\theta}}{2}\Big) = \frac{1}{2} - \frac{(1-\underline{\theta})^2}{2} < \frac{1}{2}.$$

Hence, for $c \ge \frac{1}{2}$, no active-signaling equilibrium exists in either regime, and no costly high effort is provided.

Proof. Let $n(\theta_i)$ denote the expected number of periods in which type θ_i exerts high effort. In the limit, a success requires high effort in that period and then occurs with probability θ_i, so the probability that type θ_i succeeds at least once is at most $\theta_i n(\theta_i)$. Without any success the outcome is $(0,0)$, and every posterior is at most 1. The gain of type θ_i from its equilibrium strategy over $(0;0)$ is therefore at most

$$\theta_i n(\theta_i)(1-\mu_{00}) - n(\theta_i)c = n(\theta_i)\big[\theta_i(1-\mu_{00}) - c\big],$$

so every type with $n(\theta_i) > 0$ satisfies $c \le \theta_i(1-\mu_{00})$. All types below $\underline{\theta}$ generate $(0,0)$ with probability one, and all other types that generate this outcome lie above $\underline{\theta}$, so $\mu_{00} \ge \underline{\theta}/2$. Letting $\theta_i \downarrow \underline{\theta}$ along active types gives $c \le \underline{\theta}(1-\underline{\theta}/2)$. Since a positive mass of types is active, $\underline{\theta} < 1$, and the bound lies strictly below $\frac{1}{2}$. □

A2 NoIPE: derivation of Proposition 1

Proof of Proposition 1. Under NoIPE, the agent does not observe S_1 before choosing e_2. A strategy is therefore a probability distribution over the effort profiles $(0;0)$, $(1;0)$, $(0;1)$, and $(1;1)$, with expected utilities (A.6)–(A.9). I first derive the two-threshold equilibrium for

low costs and the one-threshold equilibria for medium costs. I then connect the two cost regions, compare the coexisting equilibria, and discuss high costs.

Low costs: the two-threshold equilibrium. For $0 < c < \frac{2}{9}$, the selected equilibrium has two thresholds $0 < \bar{\theta}_A(c) < \bar{\theta}_B(c) < 1$. Types $\theta_i \leq \bar{\theta}_A$ play $(0;0)$, types $\bar{\theta}_A < \theta_i \leq \bar{\theta}_B$ exert high effort exactly once, and types $\theta_i > \bar{\theta}_B$ play $(1;1)$.

The timing of effort by the middle types is uniquely pinned down. To see this, suppose that middle types exert effort in the first period with probability q and in the second period with probability $1-q$. Then

$$\mu_{10}(q) = \frac{\int_{\bar{\theta}_A}^{\bar{\theta}_B} q\theta_i^2 \, d\theta_i + \int_{\bar{\theta}_B}^{1} \theta_i^2(1-\theta_i) \, d\theta_i}{\int_{\bar{\theta}_A}^{\bar{\theta}_B} q\theta_i \, d\theta_i + \int_{\bar{\theta}_B}^{1} \theta_i(1-\theta_i) \, d\theta_i},$$
$$\mu_{01}(q) = \frac{\int_{\bar{\theta}_A}^{\bar{\theta}_B} (1-q)\theta_i^2 \, d\theta_i + \int_{\bar{\theta}_B}^{1} \theta_i^2(1-\theta_i) \, d\theta_i}{\int_{\bar{\theta}_A}^{\bar{\theta}_B} (1-q)\theta_i \, d\theta_i + \int_{\bar{\theta}_B}^{1} \theta_i(1-\theta_i) \, d\theta_i}.$$

The high types, who exert effort in both periods, contribute symmetrically to both one-success outcomes. The middle types, however, contribute with weights q and $1-q$. Both denominators remain strictly positive at $q \in \{0,1\}$, because the high types generate each one-success outcome with positive probability, so the displayed posteriors also cover the pure timing profiles. Cross-multiplying shows that $\mu_{10}(q) - \mu_{01}(q)$ has the same sign as

$$(2q-1)\left(\int_{\bar{\theta}_A}^{\bar{\theta}_B} \theta_i^2 \, d\theta_i \int_{\bar{\theta}_B}^{1} \theta_i(1-\theta_i) \, d\theta_i - \int_{\bar{\theta}_A}^{\bar{\theta}_B} \theta_i \, d\theta_i \int_{\bar{\theta}_B}^{1} \theta_i^2(1-\theta_i) \, d\theta_i\right).$$

The bracketed term is strictly negative. It equals $\int_{\bar{\theta}_A}^{\bar{\theta}_B} \theta_i \, d\theta_i \int_{\bar{\theta}_B}^{1} \theta_i(1-\theta_i) \, d\theta_i > 0$ times the difference between the θ_i-weighted mean type of the middle interval and the $\theta_i(1-\theta_i)$-weighted mean type of the high interval. The former lies strictly below $\bar{\theta}_B$, and the latter lies strictly above it. Hence $\mu_{10}(q) < \mu_{01}(q)$ for every $q > \frac{1}{2}$ and $\mu_{10}(q) > \mu_{01}(q)$ for every $q < \frac{1}{2}$. Since $U_1 - U_2 = \theta_i(\mu_{10} - \mu_{01})$, whenever the middle types use one period more often, the one-success outcome of that period carries the strictly lower posterior, and every middle type strictly prefers to deviate to the other period. This excludes all asymmetric mixing probabilities and, in particular, both pure timing profiles. At $q = \frac{1}{2}$ the two posteriors coincide and the timing indifference holds. Thus, in the low-cost region, the unique equilibrium timing is for middle types to mix symmetrically between first- and second-period effort.

Given this timing, the likelihoods are

$$\Pr(S_1 = 0, S_2 = 0 \mid \theta_i) = \begin{cases} 1 & \text{if } \theta_i \le \bar{\theta}_A, \\ 1 - \theta_i & \text{if } \bar{\theta}_A < \theta_i \le \bar{\theta}_B, \\ (1-\theta_i)^2 & \text{if } \theta_i > \bar{\theta}_B, \end{cases} \tag{A.19}$$

$$\Pr(S_1 = 1, S_2 = 0 \mid \theta_i) = \Pr(S_1 = 0, S_2 = 1 \mid \theta_i) = \begin{cases} 0 & \text{if } \theta_i \le \bar{\theta}_A, \\ \frac{1}{2}\theta_i & \text{if } \bar{\theta}_A < \theta_i \le \bar{\theta}_B, \\ \theta_i(1-\theta_i) & \text{if } \theta_i > \bar{\theta}_B, \end{cases} \tag{A.20}$$

$$\Pr(S_1 = 1, S_2 = 1 \mid \theta_i) = \begin{cases} 0 & \text{if } \theta_i \le \bar{\theta}_B, \\ \theta_i^2 & \text{if } \theta_i > \bar{\theta}_B. \end{cases} \tag{A.21}$$

Figure A1 illustrates these likelihoods for selected threshold values $\bar{\theta}_A$ and $\bar{\theta}_B$.

By (A.1) and (A.15), the posterior belief after no success is

$$\begin{aligned} \mu_{00} &= \frac{\int_0^{\bar{\theta}_A} \theta_i \, d\theta_i + \int_{\bar{\theta}_A}^{\bar{\theta}_B} \theta_i(1-\theta_i) \, d\theta_i + \int_{\bar{\theta}_B}^1 \theta_i(1-\theta_i)^2 \, d\theta_i}{\int_0^{\bar{\theta}_A} d\theta_i + \int_{\bar{\theta}_A}^{\bar{\theta}_B} (1-\theta_i) \, d\theta_i + \int_{\bar{\theta}_B}^1 (1-\theta_i)^2 \, d\theta_i} \\ &= m_0(\bar{\theta}_A, \bar{\theta}_B) = \frac{1 + 4\bar{\theta}_A^3 + 4\bar{\theta}_B^3 - 3\bar{\theta}_B^4}{2\left(2 + 3\bar{\theta}_A^2 + 3\bar{\theta}_B^2 - 2\bar{\theta}_B^3\right)}. \end{aligned} \tag{A.22}$$

The posterior belief after exactly one success is

$$\begin{aligned} \mu_{10} = \mu_{01} &= \frac{\frac{1}{2}\int_{\bar{\theta}_A}^{\bar{\theta}_B} \theta_i^2 \, d\theta_i + \int_{\bar{\theta}_B}^1 \theta_i^2(1-\theta_i) \, d\theta_i}{\frac{1}{2}\int_{\bar{\theta}_A}^{\bar{\theta}_B} \theta_i \, d\theta_i + \int_{\bar{\theta}_B}^1 \theta_i(1-\theta_i) \, d\theta_i} \\ &= \frac{1 - 2\bar{\theta}_A^3 - 2\bar{\theta}_B^3 + 3\bar{\theta}_B^4}{2 - 3\bar{\theta}_A^2 - 3\bar{\theta}_B^2 + 4\bar{\theta}_B^3}. \end{aligned} \tag{A.23}$$

Because the middle types generate each one-success outcome with only half their success probability, this posterior differs from $m_1(\bar{\theta}_A, \bar{\theta}_B)$. At $\bar{\theta}_B = 1$, it equals $m_1(\bar{\theta}_A, 1)$. Finally, by (A.17), the posterior belief after two successes is

$$\mu_{11} = \frac{\int_{\bar{\theta}_B}^1 \theta_i^3 \, d\theta_i}{\int_{\bar{\theta}_B}^1 \theta_i^2 \, d\theta_i} = m_2(\bar{\theta}_B) = \frac{3\left(1 - \bar{\theta}_B^4\right)}{4\left(1 - \bar{\theta}_B^3\right)}. \tag{A.24}$$

The two thresholds are determined by two indifference conditions. First, the lowest active type $\bar{\theta}_A$ is indifferent between no high effort and high effort once, $U_0(\bar{\theta}_A) = U_1(\bar{\theta}_A)$.

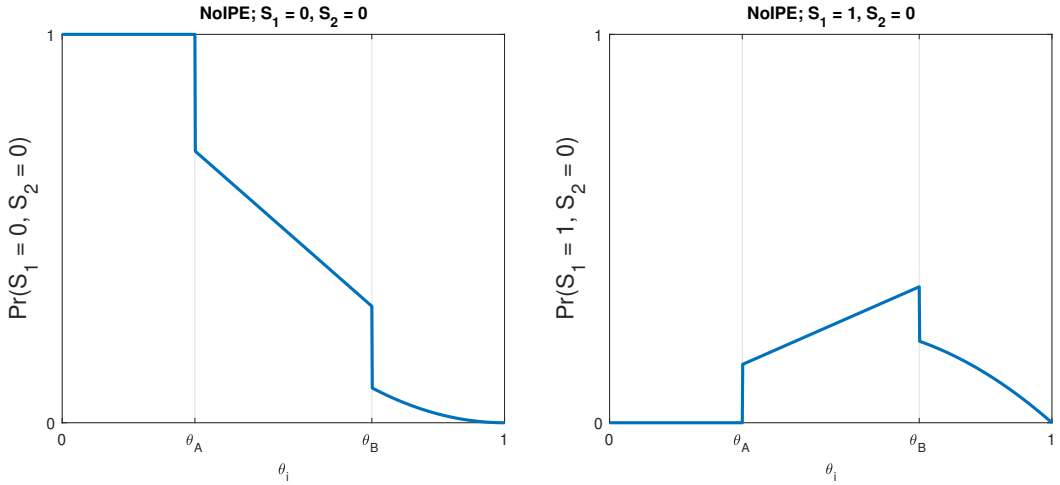

(a) $\Pr(S_1 = S_2 = 0 \mid \theta_i)$ (b) $\Pr(S_1 = 1, S_2 = 0 \mid \theta_i)$

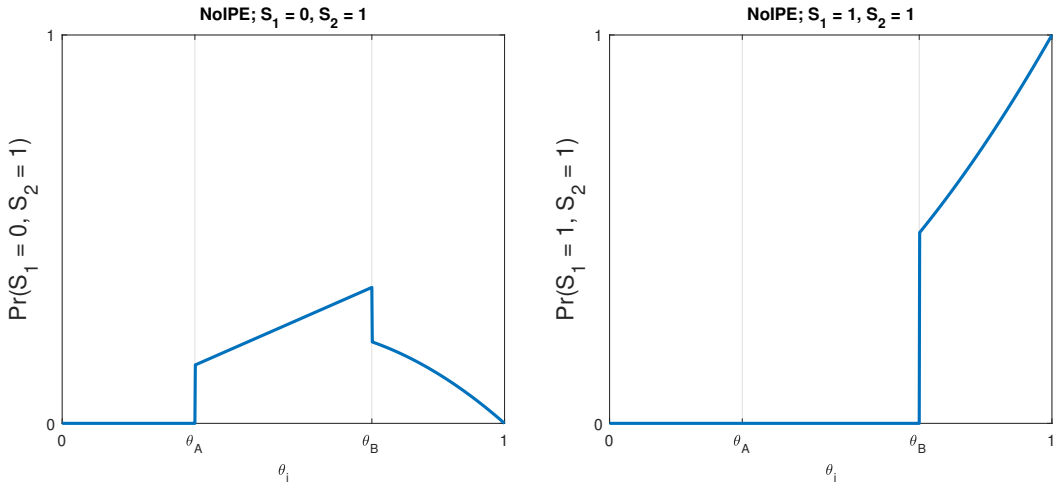

(c) $\Pr(S_1 = 0, S_2 = 1 \mid \theta_i)$ (d) $\Pr(S_1 = S_2 = 1 \mid \theta_i)$

Figure A1: NoIPE: likelihood that a type θ_i generates a given outcome (S_1, S_2) in the low-cost two-threshold equilibrium.

By (A.11), this is equivalent to

$$c = \bar{\theta}_A \left(\mu_{10} - \mu_{00}\right). \tag{A.25}$$

Second, type $\bar{\theta}_B$ is indifferent between high effort once and high effort twice, $U_1(\bar{\theta}_B) = U_{12}(\bar{\theta}_B)$. By (A.12) with $\mu_{10} = \mu_{01}$, this is equivalent to

$$c = \bar{\theta}_B \left[\bar{\theta}_B \left(\mu_{11} - \mu_{10}\right) + \left(1 - \bar{\theta}_B\right)\left(\mu_{10} - \mu_{00}\right)\right]. \tag{A.26}$$

Together with the posteriors (A.22)–(A.24), equations (A.25) and (A.26) define $\bar{\theta}_A(c)$ and $\bar{\theta}_B(c)$ in the low-cost region. The system has no closed-form solution. Solved numerically, it

has a unique solution with $0 < \bar{\theta}_A(c) < \bar{\theta}_B(c) < 1$ for every $c \in (0, \frac{2}{9})$ and no such solution for $c \geq \frac{2}{9}$. The boundary at $c = \frac{2}{9}$ is derived analytically below, after the one-threshold equilibria.

It remains to verify the incentive constraints of all types. The gain from one high effort over none, $U_1 - U_0 = \theta_i(\mu_{10} - \mu_{00}) - c$, increases in θ_i and equals zero at $\bar{\theta}_A$. Since $\mu_{10} = \mu_{01}$, the same holds for $U_2 - U_0$, and every type is indifferent between the two periods. The gain from a second high effort,

$$g(\theta_i) \equiv U_{12}(\theta_i) - U_1(\theta_i) = \theta_i(\mu_{10} - \mu_{00}) + \theta_i^2(\mu_{11} - 2\mu_{10} + \mu_{00}) - c,$$

is quadratic in θ_i, with $g(0) = -c < 0$ and $g(\bar{\theta}_B) = 0$ by (A.26). A quadratic that is negative at 0, zero at $\bar{\theta}_B \in (0,1)$, and non-negative at 1 is negative on $[0, \bar{\theta}_B)$ and non-negative on $[\bar{\theta}_B, 1]$. It therefore suffices that $g(1) \geq 0$. Using (A.26),

$$g(1) = \mu_{11} - \mu_{10} - c = (1 - \bar{\theta}_B)\Big[(1 + \bar{\theta}_B)(\mu_{11} - \mu_{10}) - \bar{\theta}_B(\mu_{10} - \mu_{00})\Big].$$

The bracket is strictly positive at the numerical solution for every $c \in (0, \frac{2}{9})$, so $g(1) > 0$ on this range. Hence types in $(\bar{\theta}_A, \bar{\theta}_B]$ prefer exactly one high effort, and types above $\bar{\theta}_B$ prefer two. Since $U_{12} - U_0 = g + (U_1 - U_0)$, types below $\bar{\theta}_A$ also prefer no high effort to two.

Medium costs: the one-threshold equilibria. For $\frac{2}{9} \leq c < \frac{1}{2}$, the selected equilibria have a single threshold $\bar{\theta}_A(c)$. Types $\theta_i \leq \bar{\theta}_A$ play $(0;0)$, and types $\theta_i > \bar{\theta}_A$ exert high effort exactly once. Let $q \in [0,1]$ denote the probability with which active types exert this effort in the first period. In the limit, the likelihoods are

$$\Pr(S_1 = 0, S_2 = 0 \mid \theta_i) = \begin{cases} 1 & \text{if } \theta_i \leq \bar{\theta}_A, \\ 1 - \theta_i & \text{if } \theta_i > \bar{\theta}_A, \end{cases} \tag{A.27}$$

$$\Pr(S_1 = 1, S_2 = 0 \mid \theta_i) = \begin{cases} 0 & \text{if } \theta_i \leq \bar{\theta}_A, \\ q\,\theta_i & \text{if } \theta_i > \bar{\theta}_A, \end{cases} \tag{A.28}$$

$$\Pr(S_1 = 0, S_2 = 1 \mid \theta_i) = \begin{cases} 0 & \text{if } \theta_i \leq \bar{\theta}_A, \\ (1 - q)\,\theta_i & \text{if } \theta_i > \bar{\theta}_A, \end{cases} \tag{A.29}$$

$$\Pr(S_1 = 1, S_2 = 1 \mid \theta_i) = 0. \tag{A.30}$$

The posterior belief after no success does not depend on q:

$$\mu_{00} = \frac{\int_0^{\bar\theta_A} \theta_i \, d\theta_i + \int_{\bar\theta_A}^1 \theta_i(1-\theta_i)\, d\theta_i}{\int_0^{\bar\theta_A} d\theta_i + \int_{\bar\theta_A}^1 (1-\theta_i)\, d\theta_i} = m_0(\bar\theta_A, 1) = \frac{1 + 2\bar\theta_A^3}{3\left(1+\bar\theta_A^2\right)}. \tag{A.31}$$

In a period that active types use with positive probability, the factor q or $1-q$ cancels in (A.1), so the posterior after a single success in that period is

$$\frac{\int_{\bar\theta_A}^1 \theta_i^2 \, d\theta_i}{\int_{\bar\theta_A}^1 \theta_i \, d\theta_i} = m_1(\bar\theta_A, 1) = \frac{2\left(1+\bar\theta_A+\bar\theta_A^2\right)}{3\left(1+\bar\theta_A\right)}. \tag{A.32}$$

With symmetric mixing, $\mu_{10} = \mu_{01} = m_1(\bar\theta_A, 1)$. The outcome $(1,1)$ and, in the pure timing profiles, a single success in the unused period have probability zero in the limit. Their posteriors follow from the perturbation and are derived below.

The marginal type $\bar\theta_A$ is indifferent between no high effort and high effort exactly once. By (A.11) and its counterpart for the second period, this requires $c = \bar\theta_A \left[m_1(\bar\theta_A, 1) - \mu_{00}\right]$ for every timing. Substituting (A.31) and (A.32) gives

$$c = \Gamma(\bar\theta_A), \qquad \Gamma(a) \equiv a\left[m_1(a,1) - m_0(a,1)\right] = \frac{a\left(4a^2+a+1\right)}{3(1+a)(1+a^2)}. \tag{A.33}$$

The function Γ satisfies $\Gamma(0) = 0$, $\Gamma(\frac{1}{2}) = \frac{2}{9}$, $\Gamma(1) = \frac{1}{2}$, and

$$\Gamma'(a) = \frac{3a^4 + 6a^3 + 12a^2 + 2a + 1}{3(1+a)^2(1+a^2)^2} > 0, \qquad \frac{1}{2} - \Gamma(a) = \frac{(1-a)\left(5a^2+4a+3\right)}{6(1+a)(1+a^2)}.$$

Hence (A.33) has a unique solution $\bar\theta_A(c) = \Gamma^{-1}(c) \in (0,1)$ for every $c \in (0, \frac{1}{2})$. This threshold increases in c, does not depend on the timing q, and tends to 1 only as $c \to \frac{1}{2}$.

Timing. With $\varepsilon > 0$, the payoff difference between $(1;\varepsilon)$ and $(\varepsilon;1)$ for an active type is $\theta_i(1-\varepsilon)(\mu_{10} - \mu_{01})$. Active types therefore prefer the period whose one-success outcome carries the higher posterior, and interior mixing requires $\mu_{10} = \mu_{01}$. Three cases arise.

If $q = \frac{1}{2}$, the outcomes $(1,0)$ and $(0,1)$ have identical likelihoods for every type and every ε, so $\mu_{10} = \mu_{01}$, and every active type is indifferent between the two periods.

If $q \in (0,1)$ and $q \neq \frac{1}{2}$, the likelihoods in the perturbed game are

$$\Pr(S_1 = 1, S_2 = 0 \mid \theta_i) = \begin{cases} \varepsilon\theta_i(1-\varepsilon\theta_i) & \text{if } \theta_i \leq \bar{\theta}_A, \\ q\,\theta_i(1-\varepsilon\theta_i) + (1-q)\,\varepsilon\theta_i(1-\theta_i) & \text{if } \theta_i > \bar{\theta}_A, \end{cases}$$

$$\Pr(S_1 = 0, S_2 = 1 \mid \theta_i) = \begin{cases} \varepsilon\theta_i(1-\varepsilon\theta_i) & \text{if } \theta_i \leq \bar{\theta}_A, \\ q\,(1-\theta_i)\varepsilon\theta_i + (1-q)\,(1-\varepsilon\theta_i)\theta_i & \text{if } \theta_i > \bar{\theta}_A. \end{cases}$$

Let $\bar{m} \equiv m_1(\bar{\theta}_A, 1)$. A first-order expansion of (A.1) in ε gives

$$\mu_{10} = \bar{m} + \frac{2\varepsilon(\kappa_0 - \kappa_1)}{q\left(1-\bar{\theta}_A^2\right)} + O(\varepsilon^2), \qquad \mu_{01} = \bar{m} + \frac{2\varepsilon(\kappa_0 - \kappa_1)}{(1-q)\left(1-\bar{\theta}_A^2\right)} + O(\varepsilon^2), \tag{A.34}$$

where

$$\kappa_0 \equiv \int_0^{\bar{\theta}_A} \theta_i(\theta_i - \bar{m})\,d\theta_i < 0, \qquad \kappa_1 \equiv \int_{\bar{\theta}_A}^1 \theta_i^2(\theta_i - \bar{m})\,d\theta_i = \int_{\bar{\theta}_A}^1 \theta_i(\theta_i - \bar{m})^2\,d\theta_i > 0.$$

The sign of κ_0 follows from $\bar{m} > \bar{\theta}_A$, and the second expression for κ_1 uses $\int_{\bar{\theta}_A}^1 \theta_i(\theta_i - \bar{m})\,d\theta_i = 0$, which is the definition of $\bar{m}$. The term κ_0 reflects the low-effort successes of inactive types, who all lie below $\bar{m}$. The term κ_1 reflects that a single success also requires a failure in the other period, which removes relatively able active types. Since $\kappa_0 - \kappa_1 < 0$, both posteriors fall below $\bar{m}$, by an amount inversely proportional to the share of active types who use the corresponding period. Hence, for every $q \neq \frac{1}{2}$ and all sufficiently small $\varepsilon > 0$, $\mu_{10} > \mu_{01}$ if $q > \frac{1}{2}$ and $\mu_{10} < \mu_{01}$ if $q < \frac{1}{2}$. The period that active types use more often carries the higher posterior, every active type strictly prefers it, and asymmetric mixing is not an equilibrium.

If $q = 1$, a single second-period success arises in the limit only through low-effort success. Its likelihood is $\varepsilon\theta_i + O(\varepsilon^2)$ for inactive types and $\varepsilon\theta_i(1-\theta_i) + O(\varepsilon^2)$ for active types, so by (A.16)

$$\mu_{01} \to \widetilde{\mu}_{01}^F(\bar{\theta}_A) \equiv \frac{\int_0^{\bar{\theta}_A} \theta_i^2\,d\theta_i + \int_{\bar{\theta}_A}^1 \theta_i^2(1-\theta_i)\,d\theta_i}{\int_0^{\bar{\theta}_A} \theta_i\,d\theta_i + \int_{\bar{\theta}_A}^1 \theta_i(1-\theta_i)\,d\theta_i} = m_1(0, \bar{\theta}_A),$$

while $\mu_{10} \to m_1(\bar{\theta}_A, 1)$. For all $\bar{\theta}_A \in (0,1)$, one has $m_1(0,\bar{\theta}_A) < m_1(\bar{\theta}_A, 1)$. Both posteriors are means of θ_i under densities proportional to $\theta_i w(\theta_i)$, where $w \equiv 1$ on $(0, \bar{\theta}_A)$ and $w(\theta_i) = 1-\theta_i$ above $\bar{\theta}_A$ for $m_1(0,\bar{\theta}_A)$, while $w(\theta_i) = \mathbb{1}\{\theta_i > \bar{\theta}_A\}$ for $m_1(\bar{\theta}_A, 1)$. The likelihood ratio of the latter density to the former is zero below $\bar{\theta}_A$ and increasing above it. Therefore, the density associated with $m_1(\bar{\theta}_A, 1)$ dominates in the monotone likelihood ratio order and has a strictly larger mean. Hence, for small ε, active types strictly prefer first-period effort, which sustains pure first-period timing. The case $q = 0$ is symmetric.

In sum, three timing profiles remain: all active types exert effort in the first period, all exert it in the second period, or all mix symmetrically. They share the threshold $\Gamma^{-1}(c)$, and in each of them all active types exert high effort exactly once.

No type exerts high effort twice. The outcome $(1,1)$ has probability zero in the limit. In the perturbed game, its likelihood is $\varepsilon\theta_i^2 + O(\varepsilon^2)$ for active types and $O(\varepsilon^2)$ for inactive types, so $\mu_{11} \to m_2(\bar{\theta}_A)$. The argument below uses only $\mu_{11} \leq 1$. Write $u \equiv m_1(\bar{\theta}_A, 1)$ and $z \equiv m_0(\bar{\theta}_A, 1)$, so that $c = \bar{\theta}_A(u-z)$. For an active type, let μ_s and $\mu_{s'}$ denote the one-success posteriors of the period it uses and of the other period. In all three timing profiles, $\mu_s = u$ and $\mu_{s'} \leq u$. By (A.12) and its counterpart for the second period, the gain from adding high effort in the other period is

$$\theta_i^2(\mu_{11} - \mu_s) + \theta_i(1-\theta_i)(\mu_{s'} - \mu_{00}) - c \leq R(\theta_i) \equiv \theta_i^2(1-u) + \theta_i(1-\theta_i)(u-z) - c.$$

Since

$$2u - z - 1 = \frac{\bar{\theta}_A^2\left(2\bar{\theta}_A^2 - \bar{\theta}_A + 5\right)}{3(1+\bar{\theta}_A)(1+\bar{\theta}_A^2)} > 0,$$

the bound $R(\theta_i) = \theta_i(u-z) - \theta_i^2(2u-z-1) - c$ is a strictly concave quadratic. Its maximum over all real θ_i equals $(u-z)^2/[4(2u-z-1)] - c$, which is non-positive if and only if $u - z \leq 4\bar{\theta}_A(2u-z-1)$. Direct computation gives

$$4\bar{\theta}_A(2u-z-1) - (u-z) = \frac{(2\bar{\theta}_A - 1)\left(4\bar{\theta}_A^4 + 10\bar{\theta}_A^2 + 3\bar{\theta}_A + 1\right)}{3(1+\bar{\theta}_A)(1+\bar{\theta}_A^2)},$$

so $R \leq 0$ on $[0,1]$ whenever $\bar{\theta}_A \geq \frac{1}{2}$. If $\bar{\theta}_A < \frac{1}{2}$, then

$$R(1) = 1 - u - c = \frac{(1-2\bar{\theta}_A)\left(1 + \bar{\theta}_A + \bar{\theta}_A^2\right)}{3\left(1+\bar{\theta}_A^2\right)} > 0.$$

Hence $R \leq 0$ for all types if and only if $\bar{\theta}_A \geq \frac{1}{2}$, that is, by the monotonicity of Γ, if and only if $c \geq \Gamma(\frac{1}{2}) = \frac{2}{9}$. For these costs, no type gains from high effort in both periods, whatever the value of μ_{11}. In particular, types below $\bar{\theta}_A$ do not gain from it, since their gain over $(0;0)$ is the sum of a gain bounded by $R \leq 0$ and the negative gain from one high effort in the period that active types use.

The remaining constraints follow directly. The gain from one high effort over none, $\theta_i(u-z) - c$, increases in θ_i and equals zero at $\bar{\theta}_A$. In the pure timing profiles, high effort in the unused period yields $\theta_i(\mu_{s'} - z) - c \leq \theta_i(u-z) - c$, which is negative for inactive types and below the gain from the used period for active types. Hence, for $\frac{2}{9} \leq c < \frac{1}{2}$, all three timing profiles with threshold $\Gamma^{-1}(c)$ are equilibria.

Connection between the two regions. At $\bar{\theta}_B = 1$, the low-cost posteriors become $\mu_{00} = m_0(\bar{\theta}_A, 1)$, $\mu_{10} = \mu_{01} = m_1(\bar{\theta}_A, 1)$, and $\mu_{11} = m_2(1) = 1$, and conditions (A.25) and (A.26) become $c = \Gamma(\bar{\theta}_A)$ and $c = 1 - m_1(\bar{\theta}_A, 1)$. The first right-hand side increases and the second decreases in $\bar{\theta}_A$, and both equal $\frac{2}{9}$ at $\bar{\theta}_A = \frac{1}{2}$. The two-threshold system therefore reaches the boundary $\bar{\theta}_B = 1$ only at $c = \frac{2}{9}$, with $\bar{\theta}_A = \frac{1}{2}$, and the numerical solution converges to this point: $\bar{\theta}_B(c) \to 1$ and $\bar{\theta}_A(c) \to \frac{1}{2}$ as $c \uparrow \frac{2}{9}$. The two-threshold equilibrium thus connects continuously to the one-threshold equilibrium with $\bar{\theta}_A = \frac{1}{2}$. At this point, the top type is exactly indifferent between one and two high efforts, since $R(1) = 0$ at $\bar{\theta}_A = \frac{1}{2}$.

Selection. For $c \in [\frac{2}{9}, \frac{1}{2})$, the two-threshold system has no solution with $0 < \bar{\theta}_A < \bar{\theta}_B < 1$, and the three one-threshold timing equilibria are the selected active-signaling equilibria. They share the threshold $\Gamma^{-1}(c)$ and induce the same total effort $1 - \Gamma^{-1}(c)$. For $c < \frac{2}{9}$, the one-threshold structure can also be an equilibrium when μ_{11} takes its limiting value $m_2(\bar{\theta}_A) < 1$. With symmetric timing, this is the case for all $c \geq 0.112$ (approximately); for lower costs, types close to 1 deviate to high effort in both periods. Its total effort $1 - \Gamma^{-1}(c)$ is strictly below the total effort $2 - \bar{\theta}_A(c) - \bar{\theta}_B(c)$ of the two-threshold equilibrium at every $c \in (0, \frac{2}{9})$ (for example, 0.54 versus 0.65 at $c = 0.2$), so the selection rule picks the two-threshold equilibrium. Numerically, I find no other active-signaling equilibria in which the number of high-effort periods is non-decreasing in θ_i.

High costs. For $c \geq \frac{1}{2}$, Lemma A1 rules out active signaling, and the pooling equilibrium of Appendix A1 exists. Hence no costly high effort is provided.

Combining the low-cost, medium-cost, and high-cost arguments establishes the threshold structure stated in Proposition 1. □

A3 IPE: derivation of Proposition 2

Proof of Proposition 2. Under IPE, the agent observes S_1 before choosing e_2, so a strategy specifies first-period effort and a second-period effort after each first-period outcome. In addition to $(0;0)$, $(1;0)$, $(0;1)$, and $(1;1)$, the agent can play the conditional strategy $(1; 1 - S_1)$, with expected utility (A.10). This strategy is central under IPE: the agent exerts effort in the first period and tries again in the second period only if the first attempt did not succeed.

Second-period choices. After a first-period outcome $S_1 = s$, high second-period effort raises the expected posterior from μ_{s0} to $\theta_i \mu_{s1} + (1 - \theta_i)\mu_{s0}$ at cost c. It is therefore optimal

if and only if

$$\theta_i \geq t_F \equiv \frac{c}{\mu_{01} - \mu_{00}} \quad \text{after } S_1 = 0, \qquad \theta_i \geq t_S \equiv \frac{c}{\mu_{11} - \mu_{10}} \quad \text{after } S_1 = 1. \tag{A.35}$$

Both conditions are threshold rules in θ_i. In each region below, $t_S > t_F$, so every type that exerts high second-period effort after a success also does so after a failure. The reverse conditional strategy $(1; S_1)$, which exerts second-period effort only after a success, is therefore never optimal. In the limit, a type with $e_1 = 0$ reaches $S_1 = 1$ with probability zero. Its choice after $S_1 = 1$ follows (A.35) and affects neither payoffs nor posteriors in the limit, so I describe each type's strategy by the effort it exerts on the path.

Low- and medium-cost candidate strategies. For low and medium costs, the selected equilibrium has three thresholds,

$$0 < \bar{\theta}_A(c) < \bar{\theta}_B(c) < \bar{\theta}_C(c) < 1.$$

Low types do not exert high effort, intermediate types exert high effort exactly once, higher intermediate types use the conditional strategy, and the highest types exert high effort in both periods. For low costs, the one-effort types exert effort only in the second period:

$$(e_1; e_2)(\theta_i) = \begin{cases} (0; 0) & \text{if } \theta_i \leq \bar{\theta}_A, \\ (0; 1) & \text{if } \bar{\theta}_A < \theta_i \leq \bar{\theta}_B, \\ (1; 1 - S_1) & \text{if } \bar{\theta}_B < \theta_i \leq \bar{\theta}_C, \\ (1; 1) & \text{if } \theta_i > \bar{\theta}_C. \end{cases} \tag{A.36}$$

For medium costs, the one-effort types exert effort only in the first period:

$$(e_1; e_2)(\theta_i) = \begin{cases} (0; 0) & \text{if } \theta_i \leq \bar{\theta}_A, \\ (1; 0) & \text{if } \bar{\theta}_A < \theta_i \leq \bar{\theta}_B, \\ (1; 1 - S_1) & \text{if } \bar{\theta}_B < \theta_i \leq \bar{\theta}_C, \\ (1; 1) & \text{if } \theta_i > \bar{\theta}_C. \end{cases} \tag{A.37}$$

In terms of (A.35), the low-cost strategy prescribes first-period effort for types above $\bar{\theta}_B$, with $t_F = \bar{\theta}_A$ and $t_S = \bar{\theta}_C$. The medium-cost strategy prescribes first-period effort for types above $\bar{\theta}_A$, with $t_F = \bar{\theta}_B$ and $t_S = \bar{\theta}_C$.

The difference between the low- and medium-cost regions is therefore the timing of effort

by types in $(\bar{\theta}_A, \bar{\theta}_B]$, who exert effort in the period whose one-success outcome carries the higher posterior belief. For low costs, the most able types exert effort in both periods and generate $(S_1, S_2) = (1, 0)$ and $(S_1, S_2) = (0, 1)$ equally often, while conditional types generate $(1, 0)$ with probability θ_i and $(0, 1)$ only with probability $\theta_i(1 - \theta_i)$. Relative to $(0, 1)$, the outcome $(1, 0)$ therefore contains more conditional types and fewer of the most able types, and μ_{01} slightly exceeds μ_{10}. As costs approach c_I^L, fewer types exert effort in both periods. The types who switch to the conditional strategy lie near $\bar{\theta}_C$, above μ_{10}, and their likelihood of generating $(1, 0)$ rises from $\theta_i(1 - \theta_i)$ to θ_i, which raises μ_{10} relative to μ_{01}. At the cutoff between the two regions, the two posteriors coincide, and the one-effort types switch from second-period effort to first-period effort.

Likelihoods. The principal observes (S_1, S_2) and knows that IPE is implemented. Under the low-cost strategy (A.36), the likelihoods are

$$\Pr(S_1 = 0, S_2 = 0 \mid \theta_i) = \begin{cases} 1 & \text{if } \theta_i \le \bar{\theta}_A, \\ 1 - \theta_i & \text{if } \bar{\theta}_A < \theta_i \le \bar{\theta}_B, \\ (1 - \theta_i)^2 & \text{if } \theta_i > \bar{\theta}_B, \end{cases} \tag{A.38}$$

$$\Pr(S_1 = 1, S_2 = 0 \mid \theta_i) = \begin{cases} 0 & \text{if } \theta_i \le \bar{\theta}_B, \\ \theta_i & \text{if } \bar{\theta}_B < \theta_i \le \bar{\theta}_C, \\ \theta_i(1 - \theta_i) & \text{if } \theta_i > \bar{\theta}_C, \end{cases} \tag{A.39}$$

$$\Pr(S_1 = 0, S_2 = 1 \mid \theta_i) = \begin{cases} 0 & \text{if } \theta_i \le \bar{\theta}_A, \\ \theta_i & \text{if } \bar{\theta}_A < \theta_i \le \bar{\theta}_B, \\ \theta_i(1 - \theta_i) & \text{if } \theta_i > \bar{\theta}_B, \end{cases} \tag{A.40}$$

$$\Pr(S_1 = 1, S_2 = 1 \mid \theta_i) = \begin{cases} 0 & \text{if } \theta_i \le \bar{\theta}_C, \\ \theta_i^2 & \text{if } \theta_i > \bar{\theta}_C. \end{cases} \tag{A.41}$$

Under the medium-cost strategy (A.37), the likelihoods of $(S_1, S_2) = (0, 0)$ and $(S_1, S_2) = (1, 1)$ are unchanged, while the likelihoods of the one-success outcomes become

$$\Pr(S_1 = 1, S_2 = 0 \mid \theta_i) = \begin{cases} 0 & \text{if } \theta_i \le \bar{\theta}_A, \\ \theta_i & \text{if } \bar{\theta}_A < \theta_i \le \bar{\theta}_C, \\ \theta_i(1-\theta_i) & \text{if } \theta_i > \bar{\theta}_C, \end{cases} \tag{A.42}$$

$$\Pr(S_1 = 0, S_2 = 1 \mid \theta_i) = \begin{cases} 0 & \text{if } \theta_i \le \bar{\theta}_B, \\ \theta_i(1-\theta_i) & \text{if } \theta_i > \bar{\theta}_B. \end{cases} \tag{A.43}$$

Figure A2 illustrates these likelihoods for selected threshold values.

Posterior beliefs in the low-cost region. By (A.1) and the likelihoods above, the posterior belief after no success is

$$\begin{aligned} \mu_{00} &= \frac{\int_0^{\bar{\theta}_A} \theta_i \, d\theta_i + \int_{\bar{\theta}_A}^{\bar{\theta}_B} \theta_i(1-\theta_i) \, d\theta_i + \int_{\bar{\theta}_B}^{1} \theta_i(1-\theta_i)^2 \, d\theta_i}{\int_0^{\bar{\theta}_A} d\theta_i + \int_{\bar{\theta}_A}^{\bar{\theta}_B} (1-\theta_i) \, d\theta_i + \int_{\bar{\theta}_B}^{1} (1-\theta_i)^2 \, d\theta_i} \\ &= m_0(\bar{\theta}_A, \bar{\theta}_B) = \frac{1 + 4\bar{\theta}_A^3 + 4\bar{\theta}_B^3 - 3\bar{\theta}_B^4}{2\left(2 + 3\bar{\theta}_A^2 + 3\bar{\theta}_B^2 - 2\bar{\theta}_B^3\right)}. \end{aligned} \tag{A.44}$$

This is the same function as under NoIPE in (A.22), because in both regimes types up to $\bar{\theta}_A$ never succeed, the one-effort types fail with probability $1-\theta_i$, and all higher types fail twice with probability $(1-\theta_i)^2$. The posterior belief after a success only in the first period is

$$\mu_{10} = \frac{\int_{\bar{\theta}_B}^{\bar{\theta}_C} \theta_i^2 \, d\theta_i + \int_{\bar{\theta}_C}^{1} \theta_i^2(1-\theta_i) \, d\theta_i}{\int_{\bar{\theta}_B}^{\bar{\theta}_C} \theta_i \, d\theta_i + \int_{\bar{\theta}_C}^{1} \theta_i(1-\theta_i) \, d\theta_i} = m_1(\bar{\theta}_B, \bar{\theta}_C) = \frac{1 - 4\bar{\theta}_B^3 + 3\bar{\theta}_C^4}{2\left(1 - 3\bar{\theta}_B^2 + 2\bar{\theta}_C^3\right)}. \tag{A.45}$$

The posterior belief after a success only in the second period is

$$\mu_{01} = \frac{\int_{\bar{\theta}_A}^{\bar{\theta}_B} \theta_i^2 \, d\theta_i + \int_{\bar{\theta}_B}^{1} \theta_i^2(1-\theta_i) \, d\theta_i}{\int_{\bar{\theta}_A}^{\bar{\theta}_B} \theta_i \, d\theta_i + \int_{\bar{\theta}_B}^{1} \theta_i(1-\theta_i) \, d\theta_i} = m_1(\bar{\theta}_A, \bar{\theta}_B) = \frac{1 - 4\bar{\theta}_A^3 + 3\bar{\theta}_B^4}{2\left(1 - 3\bar{\theta}_A^2 + 2\bar{\theta}_B^3\right)}. \tag{A.46}$$

Finally, the posterior belief after two successes is

$$\mu_{11} = \frac{\int_{\bar{\theta}_C}^{1} \theta_i^3 \, d\theta_i}{\int_{\bar{\theta}_C}^{1} \theta_i^2 \, d\theta_i} = m_2(\bar{\theta}_C) = \frac{3\left(1 - \bar{\theta}_C^4\right)}{4\left(1 - \bar{\theta}_C^3\right)}. \tag{A.47}$$

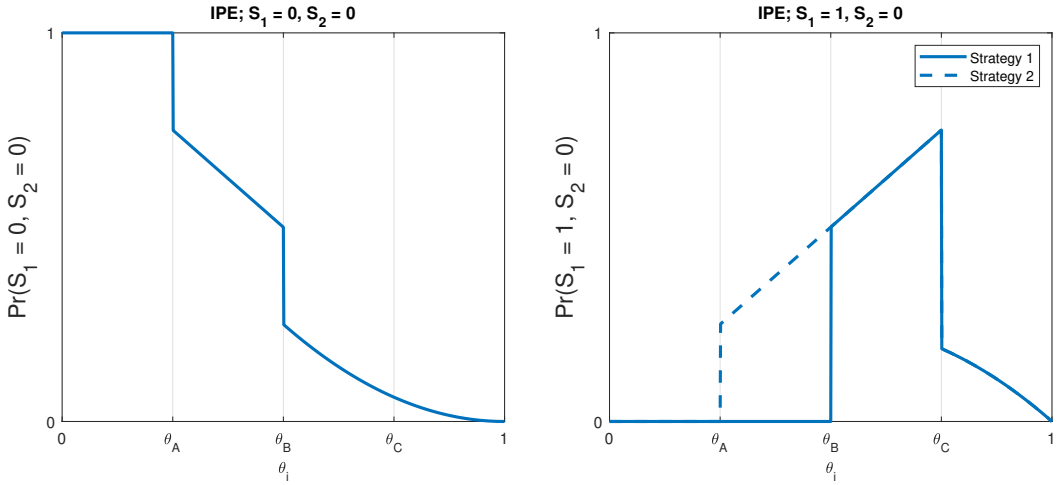

(a) $\Pr(S_1 = S_2 = 0 \mid \theta_i)$ (b) $\Pr(S_1 = 1, S_2 = 0 \mid \theta_i)$

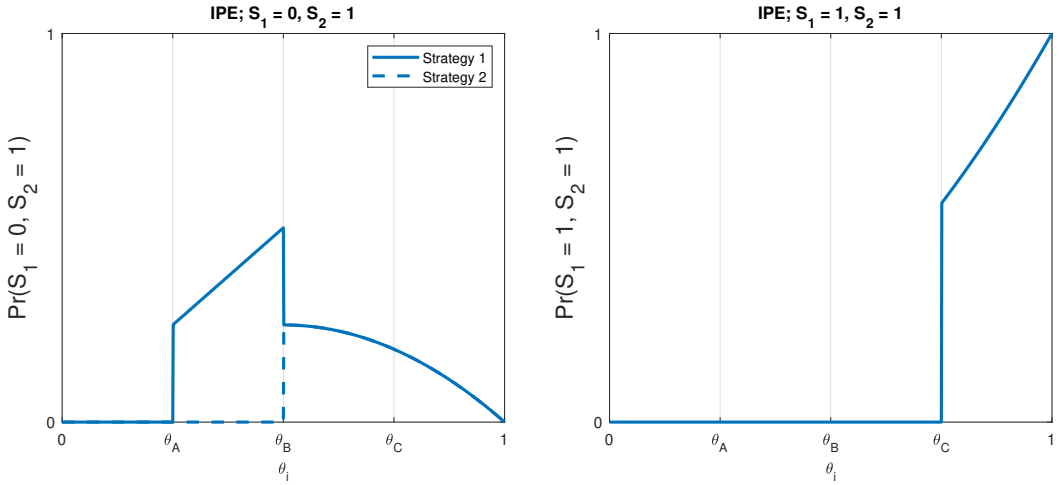

(c) $\Pr(S_1 = 0, S_2 = 1 \mid \theta_i)$ (d) $\Pr(S_1 = S_2 = 1 \mid \theta_i)$

Figure A2: IPE: likelihood that a type θ_i generates a given outcome (S_1, S_2). The solid lines (Strategy 1 in the legend) correspond to the low-cost strategy in (A.36), and the dashed lines (Strategy 2) correspond to the medium-cost strategy in (A.37).

Posterior beliefs in the medium-cost region. In the medium-cost region, $\mu_{00} = m_0(\bar{\theta}_A, \bar{\theta}_B)$ and $\mu_{11} = m_2(\bar{\theta}_C)$ are given by the same formulas as in the low-cost region. The posterior belief after a success only in the first period is

$$\mu_{10} = \frac{\int_{\bar{\theta}_A}^{\bar{\theta}_C} \theta_i^2 \, d\theta_i + \int_{\bar{\theta}_C}^{1} \theta_i^2 (1-\theta_i) \, d\theta_i}{\int_{\bar{\theta}_A}^{\bar{\theta}_C} \theta_i \, d\theta_i + \int_{\bar{\theta}_C}^{1} \theta_i (1-\theta_i) \, d\theta_i} = m_1(\bar{\theta}_A, \bar{\theta}_C) = \frac{1 - 4\bar{\theta}_A^3 + 3\bar{\theta}_C^4}{2\left(1 - 3\bar{\theta}_A^2 + 2\bar{\theta}_C^3\right)}. \tag{A.48}$$

The posterior belief after a success only in the second period is

$$\mu_{01} = \frac{\int_{\bar{\theta}_B}^1 \theta_i^2(1-\theta_i)\,d\theta_i}{\int_{\bar{\theta}_B}^1 \theta_i(1-\theta_i)\,d\theta_i} = m_1(\bar{\theta}_B, \bar{\theta}_B) = \frac{1+2\bar{\theta}_B+3\bar{\theta}_B^2}{2\left(1+2\bar{\theta}_B\right)}. \tag{A.49}$$

Threshold equations in the low-cost region. In the low-cost region, the thresholds are determined by three indifference conditions. First, the lowest active type $\bar{\theta}_A$ is indifferent between no high effort and effort only in the second period, $U_0(\bar{\theta}_A) = U_2(\bar{\theta}_A)$, which is equivalent to

$$c = \bar{\theta}_A\left(\mu_{01} - \mu_{00}\right). \tag{A.50}$$

This is the condition $t_F = \bar{\theta}_A$ in (A.35). Second, type $\bar{\theta}_B$ is indifferent between effort only in the second period and the conditional strategy, $U_2(\bar{\theta}_B) = U_{\text{cond}}(\bar{\theta}_B)$. Using (A.8) and (A.10), this is equivalent to

$$c = \frac{\bar{\theta}_B}{1-\bar{\theta}_B}\left[\mu_{10} - \bar{\theta}_B\mu_{01} - (1-\bar{\theta}_B)\mu_{00}\right]. \tag{A.51}$$

Third, type $\bar{\theta}_C$ is indifferent between the conditional strategy and high effort in both periods, $U_{\text{cond}}(\bar{\theta}_C) = U_{12}(\bar{\theta}_C)$. By (A.14), this is equivalent to

$$c = \bar{\theta}_C\left(\mu_{11} - \mu_{10}\right), \tag{A.52}$$

which is the condition $t_S = \bar{\theta}_C$. Equations (A.50)–(A.52), together with the posterior beliefs (A.44)–(A.47), define the low-cost thresholds $(\bar{\theta}_A(c), \bar{\theta}_B(c), \bar{\theta}_C(c))$.

This low-cost system applies as long as the one-effort types weakly prefer second-period effort to first-period effort, that is, as long as

$$\mu_{01} \geq \mu_{10}.$$

The cutoff c_I^L is the cost at which this inequality binds. At $\mu_{10} = \mu_{01}$, condition (A.51) reduces to $c = \bar{\theta}_B(\mu_{01} - \mu_{00})$, which together with (A.50) gives $\bar{\theta}_B = \bar{\theta}_A$. Hence the set of one-effort types is empty at c_I^L. The low- and medium-cost posteriors then coincide, since $m_1(\bar{\theta}_B, \bar{\theta}_C) = m_1(\bar{\theta}_A, \bar{\theta}_C)$ and $m_1(\bar{\theta}_A, \bar{\theta}_B) = m_1(\bar{\theta}_B, \bar{\theta}_B)$ at $\bar{\theta}_A = \bar{\theta}_B$, and the low- and medium-cost systems have the same solution. Solved numerically, the low-cost system has a unique solution with $0 < \bar{\theta}_A < \bar{\theta}_B < \bar{\theta}_C < 1$ for every $c \in (0, c_I^L)$, and $c_I^L \approx 0.194$, with $\bar{\theta}_A = \bar{\theta}_B \approx 0.481$ and $\bar{\theta}_C \approx 0.816$ at this cost.

Threshold equations in the medium-cost region. In the medium-cost region, the one-effort types exert effort only in the first period. The thresholds are again determined by three indifference conditions. First, $U_0(\bar{\theta}_A) = U_1(\bar{\theta}_A)$, which by (A.11) gives

$$c = \bar{\theta}_A \left(\mu_{10} - \mu_{00}\right). \tag{A.53}$$

Second, $U_1(\bar{\theta}_B) = U_{\text{cond}}(\bar{\theta}_B)$, which by (A.13) gives

$$c = \bar{\theta}_B \left(\mu_{01} - \mu_{00}\right), \tag{A.54}$$

that is, $t_F = \bar{\theta}_B$. Third, $U_{\text{cond}}(\bar{\theta}_C) = U_{12}(\bar{\theta}_C)$, which by (A.14) gives

$$c = \bar{\theta}_C \left(\mu_{11} - \mu_{10}\right), \tag{A.55}$$

that is, $t_S = \bar{\theta}_C$. Equations (A.53)–(A.55), together with $\mu_{00} = m_0(\bar{\theta}_A, \bar{\theta}_B)$, (A.48), (A.49), and $\mu_{11} = m_2(\bar{\theta}_C)$, define the medium-cost thresholds. This system applies as long as the one-effort types weakly prefer first-period effort to second-period effort,

$$\mu_{10} \geq \mu_{01},$$

and as long as $\bar{\theta}_C(c) < 1$. Solved numerically, the medium-cost system has a unique solution with $0 < \bar{\theta}_A < \bar{\theta}_B < \bar{\theta}_C < 1$ for every $c \in (c_I^L, c_I^H)$, with $\mu_{10} > \mu_{01}$, and $\bar{\theta}_C(c) \to 1$ as $c \uparrow c_I^H$. The cutoff c_I^H is defined by

$$\bar{\theta}_C(c_I^H) = 1.$$

At this cost, $\mu_{10} = m_1(\bar{\theta}_A, 1)$ and $\mu_{11} = m_2(1) = 1$, so (A.55) becomes $c_I^H = 1 - m_1(\bar{\theta}_A, 1)$. Together with (A.53) and (A.54), this gives $c_I^H \approx 0.232$, with $\bar{\theta}_A \approx 0.473$ and $\bar{\theta}_B \approx 0.538$.

High-cost active-signaling region. For $c \in [c_I^H, \frac{1}{2})$, no type exerts high effort after a first-period success. The selected equilibrium is obtained from the medium-cost equilibrium by setting $\bar{\theta}_C = 1$. The strategy has two thresholds:

$$(e_1; e_2)(\theta_i) = \begin{cases} (0;0) & \text{if } \theta_i \leq \bar{\theta}_A, \\ (1;0) & \text{if } \bar{\theta}_A < \theta_i \leq \bar{\theta}_B, \\ (1;1-S_1) & \text{if } \theta_i > \bar{\theta}_B. \end{cases} \tag{A.56}$$

The likelihoods are

$$\Pr(S_1 = 0, S_2 = 0 \mid \theta_i) = \begin{cases} 1 & \text{if } \theta_i \le \bar{\theta}_A, \\ 1 - \theta_i & \text{if } \bar{\theta}_A < \theta_i \le \bar{\theta}_B, \\ (1-\theta_i)^2 & \text{if } \theta_i > \bar{\theta}_B, \end{cases} \tag{A.57}$$

$$\Pr(S_1 = 1, S_2 = 0 \mid \theta_i) = \begin{cases} 0 & \text{if } \theta_i \le \bar{\theta}_A, \\ \theta_i & \text{if } \theta_i > \bar{\theta}_A, \end{cases} \tag{A.58}$$

$$\Pr(S_1 = 0, S_2 = 1 \mid \theta_i) = \begin{cases} 0 & \text{if } \theta_i \le \bar{\theta}_B, \\ \theta_i(1-\theta_i) & \text{if } \theta_i > \bar{\theta}_B, \end{cases} \tag{A.59}$$

$$\Pr(S_1 = 1, S_2 = 1 \mid \theta_i) = 0. \tag{A.60}$$

By (A.1), the posterior beliefs after the outcomes that occur in the limit are

$$\mu_{00} = m_0(\bar{\theta}_A, \bar{\theta}_B) = \frac{1 + 4\bar{\theta}_A^3 + 4\bar{\theta}_B^3 - 3\bar{\theta}_B^4}{2\left(2 + 3\bar{\theta}_A^2 + 3\bar{\theta}_B^2 - 2\bar{\theta}_B^3\right)}, \tag{A.61}$$

$$\mu_{10} = \frac{\int_{\bar{\theta}_A}^1 \theta_i^2 \, d\theta_i}{\int_{\bar{\theta}_A}^1 \theta_i \, d\theta_i} = m_1(\bar{\theta}_A, 1) = \frac{2\left(1 + \bar{\theta}_A + \bar{\theta}_A^2\right)}{3\left(1 + \bar{\theta}_A\right)}, \tag{A.62}$$

$$\mu_{01} = \frac{\int_{\bar{\theta}_B}^1 \theta_i^2(1-\theta_i) \, d\theta_i}{\int_{\bar{\theta}_B}^1 \theta_i(1-\theta_i) \, d\theta_i} = m_1(\bar{\theta}_B, \bar{\theta}_B) = \frac{1 + 2\bar{\theta}_B + 3\bar{\theta}_B^2}{2\left(1 + 2\bar{\theta}_B\right)}. \tag{A.63}$$

The posterior μ_{10} is the same function of $\bar{\theta}_A$ as the NoIPE one-success posterior (A.32), because every active type, whether it plays $(1;0)$ or the conditional strategy, generates $(1,0)$ exactly by succeeding in the first period. The outcome $(1,1)$ has probability zero in the limit. In the perturbed game, its likelihood is $\varepsilon\theta_i^2 + O(\varepsilon^2)$ for active types, so $\mu_{11} \to m_2(\bar{\theta}_A)$. The incentive argument below uses only $\mu_{11} \le 1$.

The thresholds in the high-cost region are determined by $U_0(\bar{\theta}_A) = U_1(\bar{\theta}_A)$ and $U_1(\bar{\theta}_B) = U_{\text{cond}}(\bar{\theta}_B)$. By (A.11) and (A.13), these conditions become

$$c = \bar{\theta}_A \left(\mu_{10} - \mu_{00}\right), \tag{A.64}$$

and

$$c = \bar{\theta}_B \left(\mu_{01} - \mu_{00}\right). \tag{A.65}$$

Solved numerically, this system has a unique solution with $0 < \bar{\theta}_A < \bar{\theta}_B < 1$ for every

$c \in [c_I^H, \frac{1}{2})$, with $\mu_{10} > \mu_{01}$. In addition, no type may want to exert high effort after a first-period success. By (A.35) and $\mu_{11} \leq 1$, this is ensured by

$$c \geq 1 - \mu_{10} = 1 - m_1(\bar{\theta}_A, 1).$$

At c_I^H, the condition holds with equality, as shown above. Above c_I^H, the numerical solution $\bar{\theta}_A(c)$ increases in c, and $m_1(\bar{\theta}_A, 1)$ increases in $\bar{\theta}_A$, so $1-\mu_{10}$ falls while c rises. The condition therefore holds for all $c \in [c_I^H, \frac{1}{2})$, whatever the value of μ_{11}.

Global incentive constraints. The threshold equations make the threshold types indifferent. It remains to show that each type's prescribed strategy is sequentially optimal. In every region, the prescribed second-period choices coincide with the threshold rules (A.35), with t_F and t_S as stated above and $t_S \geq 1$ in the high-cost region. It therefore suffices to check the first-period choice, given optimal second-period choices. Let $V_s(\theta_i) \equiv \mu_{s0} + \max\{0, \theta_i(\mu_{s1} - \mu_{s0}) - c\}$ denote the continuation value after $S_1 = s$. The gain from high first-period effort is $D(\theta_i) \equiv \theta_i [V_1(\theta_i) - V_0(\theta_i)] - c$. Since $t_F < t_S$,

$$D(\theta_i) = \begin{cases} U_1(\theta_i) - U_0(\theta_i) = \theta_i(\mu_{10} - \mu_{00}) - c & \text{if } \theta_i \leq t_F, \\ f(\theta_i) \equiv U_{\text{cond}}(\theta_i) - U_2(\theta_i) & \text{if } t_F < \theta_i \leq t_S, \\ h(\theta_i) \equiv U_{12}(\theta_i) - U_2(\theta_i) & \text{if } \theta_i > t_S, \end{cases} \tag{A.66}$$

where

$$f(\theta_i) = -\theta_i^2(\mu_{01} - \mu_{00}) + \theta_i(\mu_{10} - \mu_{00} + c) - c, \qquad h(\theta_i) = f(\theta_i) + \theta_i [\theta_i(\mu_{11} - \mu_{10}) - c].$$

Since $\mu_{01} > \mu_{00}$, f is strictly concave, and $f(1) = \mu_{10} - \mu_{01}$. Moreover, $h(\theta_i) > f(\theta_i)$ for $\theta_i > t_S$.

Medium- and high-cost regions. Here first-period effort is prescribed for types above $\bar{\theta}_A$, and $t_F = \bar{\theta}_B$. On $[0, \bar{\theta}_B]$, $D = U_1 - U_0$ increases in θ_i and equals zero at $\bar{\theta}_A$ by (A.53) and (A.64). On $(\bar{\theta}_B, 1]$, $D \geq f$. By (A.54) and (A.65), $f(\bar{\theta}_B) = \bar{\theta}_B(\mu_{10} - \mu_{01}) \geq 0$, and $f(1) = \mu_{10} - \mu_{01} \geq 0$ in these regions. Concavity gives $f \geq 0$ on $[\bar{\theta}_B, 1]$, so all prescribed first-period choices are optimal.

Low-cost region. Here first-period effort is prescribed for types above $\bar{\theta}_B$, with $t_F = \bar{\theta}_A$ and $t_S = \bar{\theta}_C$. For $\theta_i \leq \bar{\theta}_A$, the gain is $\theta_i(\mu_{10} - \mu_{00}) - c \leq \theta_i(\mu_{01} - \mu_{00}) - c \leq 0$, because $\mu_{10} \leq \mu_{01}$ and $\theta_i \leq t_F$. For $\bar{\theta}_A < \theta_i \leq \bar{\theta}_C$, the gain is $f(\theta_i)$, with $f(0) = -c < 0$ and $f(\bar{\theta}_B) = 0$ by (A.51). Here $f(1) = \mu_{10} - \mu_{01} \leq 0$, so concavity alone does not determine the sign of f

between $\bar{\theta}_B$ and $\bar{\theta}_C$. If $f(\bar{\theta}_C) \geq 0$, concavity implies $f \leq 0$ on $[0, \bar{\theta}_B]$ and $f \geq 0$ on $[\bar{\theta}_B, \bar{\theta}_C]$. For types above $\bar{\theta}_C$, the gain is $h(\theta_i)$, a quadratic with $h(0) = -c < 0$ and $h(\bar{\theta}_C) = f(\bar{\theta}_C)$. A quadratic that is negative at 0 and non-negative at $\bar{\theta}_C$ and at 1 is non-negative on $[\bar{\theta}_C, 1]$, so it suffices in addition that $h(1) = \mu_{11} - \mu_{01} - c \geq 0$. At the numerical solution, $f(\bar{\theta}_C) > 0$ and $\mu_{11} - \mu_{01} - c > 0.04$ for every $c \in (0, c_I^L)$. Hence all prescribed first-period choices are optimal in the low-cost region as well.

Selection. For some costs, other active-signaling equilibria coexist with the equilibrium derived above. First, for c between approximately 0.071 and c_I^H, the two-threshold strategy (A.56) is also an equilibrium when μ_{11} takes its limiting value $m_2(\bar{\theta}_A)$. In numerical solutions, it induces strictly less total effort than the three-threshold equilibrium at every such cost (for example, 0.72 versus 0.79 at $c = 0.2$). Second, for costs above roughly 0.15, there are equilibria in which no type uses the evaluation: every active type exerts high effort exactly once, either always in the first period (an equilibrium for $c \geq 0.156$) or always in the second period. These equilibria generate the same outcomes as the corresponding pure timing equilibria under NoIPE. Their threshold is therefore $\Gamma^{-1}(c)$ from (A.33), and their total effort is $1 - \Gamma^{-1}(c)$. In numerical solutions, this effort is strictly below the total effort of the equilibrium derived above at every $c \in (0, \frac{1}{2})$ (for example, 0.36 versus 0.46 at $c = 0.3$); for $c \geq c_I^H$, this also follows from Step 3 of the proof of Proposition 3. Numerically, I find no other active-signaling equilibria in which the types playing each strategy form intervals ordered from $(0; 0)$ through one high effort and the conditional strategy to $(1; 1)$. The selection rule therefore picks the equilibrium derived above in all three regions.

High costs. For $c \geq \frac{1}{2}$, Lemma A1 rules out active signaling, and the pooling equilibrium of Appendix A1 exists. Hence no costly high effort is provided.

Combining the low-, medium-, and high-cost threshold systems establishes the selected active-signaling equilibrium structure stated in Proposition 2. □

A4 Proof of Proposition 3

Throughout, I work with the selected active-signaling equilibria characterized in Propositions 1 and 2 in the limit $\varepsilon \to 0$. I write $\bar{\theta}_A^N \leq \bar{\theta}_B^N$ for the NoIPE thresholds and $\bar{\theta}_A^I \leq \bar{\theta}_B^I \leq \bar{\theta}_C^I$ for the IPE thresholds. As in the main text, $\mathcal{E}^{NoIPE}(c)$ and $\mathcal{E}^{IPE}(c)$ denote expected total high effort, that is, the integral over types of the expected number of periods with high effort, and

$$\Delta(c) \equiv \mathcal{E}^{IPE}(c) - \mathcal{E}^{NoIPE}(c).$$

Proposition 3 states that $\Delta(c) > 0$ for all $c \in [c_I^H, \frac{1}{2})$ and $\Delta(c) = 0$ for $c \geq \frac{1}{2}$, and reports that, in numerical solutions, Δ is negative below $c^* \approx 0.135$ and positive above it. The proof uses the following property of the failure posterior (A.15).

Lemma A2. *For all* $0 < a < b < 1$,

$$m_0(a, b) < m_0(a, 1).$$

In words, when types above b get a second attempt and fail twice only with probability $(1 - \theta_i)^2$, fewer able types remain among those who never succeed, and the posterior after no success falls.

Proof. A direct computation gives

$$m_0(a, 1) - m_0(a, b) = \frac{(1 - b)^2 \, \Pi(a, b)}{6\,(1 + a^2)\,(2 + 3a^2 + 3b^2 - 2b^3)},$$

where

$$\Pi(a, b) = 3a^2\big(1 + 2b + 3b^2\big) + \big(1 + 2b + 9b^2\big) - 4a^3\big(1 + 2b\big).$$

The denominator is positive on $(0, 1)^2$, since $2 + 3a^2 + 3b^2 - 2b^3 = 2 + 3a^2 + b^2(3 - 2b) > 0$, so it suffices to show $\Pi(a, b) > 0$. Since $0 < a < 1$, $a^3 < a^2$, hence $4a^3(1 + 2b) < 4a^2(1 + 2b)$ and

$$\Pi(a, b) > a^2\big(9b^2 - 2b - 1\big) + \big(1 + 2b + 9b^2\big).$$

If $9b^2 - 2b - 1 \geq 0$, the right-hand side is positive. If $9b^2 - 2b - 1 < 0$, then $a^2 < b^2$ implies $a^2(9b^2 - 2b - 1) > b^2(9b^2 - 2b - 1)$, so

$$\Pi(a, b) > b^2\big(9b^2 - 2b - 1\big) + \big(1 + 2b + 9b^2\big) = 9b^4 - 2b^3 + 8b^2 + 2b + 1 > 0.$$

In either case $\Pi(a, b) > 0$, which proves the claim. The difference vanishes only in the limit $b \to 1$. □

Proof of Proposition 3. The proof proceeds in five steps. Step 0 expresses expected total effort through the thresholds, Step 1 treats prohibitive costs, Step 2 establishes the continuity of Δ, Step 3 proves that IPE induces strictly more effort for $c \in [c_I^H, \frac{1}{2})$, and Step 4 compares the two regimes numerically for $c \in (0, c_I^H)$.

Step 0: Expected total effort as a function of the thresholds. Integrating the equilibrium effort choices of Propositions 1 and 2 over $\theta_i \sim \mathcal{U}[0, 1]$ gives closed forms for

total effort in terms of the thresholds. Under NoIPE, types in $(\bar{\theta}_A^N, \bar{\theta}_B^N]$ exert high effort once and types above $\bar{\theta}_B^N$ exert it in both periods, so

$$\mathcal{E}^{NoIPE}(c) = \begin{cases} 2 - \bar{\theta}_A^N - \bar{\theta}_B^N, & 0 \le c < \frac{2}{9}, \\ 1 - \bar{\theta}_A^N, & \frac{2}{9} \le c < \frac{1}{2}, \\ 0, & c \ge \frac{1}{2}. \end{cases} \tag{A.67}$$

Under IPE, types in $(\bar{\theta}_A^I, \bar{\theta}_B^I]$ exert high effort once, types in $(\bar{\theta}_B^I, \bar{\theta}_C^I]$ play the conditional strategy $(1; 1 - S_1)$ with expected effort $1 + (1 - \theta_i)$, and types above $\bar{\theta}_C^I$ exert high effort in both periods. Since

$$\int_{\bar{\theta}_B^I}^{\bar{\theta}_C^I} \big[1 + (1 - \theta_i)\big]\, d\theta_i = 2(\bar{\theta}_C^I - \bar{\theta}_B^I) - \tfrac{1}{2}\big[(\bar{\theta}_C^I)^2 - (\bar{\theta}_B^I)^2\big],$$

collecting terms yields

$$\mathcal{E}^{IPE}(c) = \begin{cases} 2 - \bar{\theta}_A^I - \bar{\theta}_B^I - \dfrac{(\bar{\theta}_C^I)^2 - (\bar{\theta}_B^I)^2}{2}, & 0 \le c < c_I^H, \\ 1 - \bar{\theta}_A^I + \dfrac{(1 - \bar{\theta}_B^I)^2}{2}, & c_I^H \le c < \frac{1}{2}, \\ 0, & c \ge \frac{1}{2}. \end{cases} \tag{A.68}$$

The first line of (A.68) applies to both the low-cost region ($c < c_I^L$) and the medium-cost region ($c_I^L \le c < c_I^H$) of Proposition 2, which differ only in the timing of the one-effort types. The second line is obtained from the first by setting $\bar{\theta}_C^I = 1$, the defining feature of the high-cost region; indeed, $2 - \bar{\theta}_A^I - \bar{\theta}_B^I - \frac{1}{2}\big[1 - (\bar{\theta}_B^I)^2\big] = 1 - \bar{\theta}_A^I + \frac{1}{2}(1 - \bar{\theta}_B^I)^2$.

Step 1: Prohibitive costs. For $c \ge \frac{1}{2}$, Lemma A1 rules out active signaling in both regimes, so $\mathcal{E}^{NoIPE}(c) = \mathcal{E}^{IPE}(c) = 0$ and $\Delta(c) = 0$. This establishes the claim of Proposition 3 for $c \ge \frac{1}{2}$.

Step 2: Continuity of Δ. Within each cost region, the thresholds solve the corresponding system of indifference conditions from Appendices A2 and A3, which consists of rational equations in the thresholds and c. At the computed solutions, the Jacobian of each system is non-singular, so by the implicit function theorem the thresholds are continuously differentiable functions of c within each region. At the region boundaries, the effort expressions agree: at $c = \frac{2}{9}$, $\bar{\theta}_B^N \to 1$ and $\bar{\theta}_A^N \to \frac{1}{2}$, so the two branches of (A.67) coincide; at $c = c_I^L$, the

low- and medium-cost IPE systems have the same solution, and (A.68) uses one expression on both sides; and at $c = c_I^H$, $\bar{\theta}_C^I \to 1$, which is the identity at the end of Step 0. Hence $\mathcal{E}^{NoIPE}$, $\mathcal{E}^{IPE}$, and Δ are continuous on $(0, \frac{1}{2})$, with $\Delta(c) \to 0$ as $c \downarrow 0$, where both efforts tend to 2, and as $c \uparrow \frac{1}{2}$, where both efforts tend to 0.

Step 3: IPE induces strictly more effort for $c \in [c_I^H, \frac{1}{2})$. Consider $c \in [c_I^H, \frac{1}{2})$. IPE is in its high-cost region with $\bar{\theta}_C^I = 1$, and NoIPE is in its one-threshold region because $c_I^H > \frac{2}{9}$, which is shown at the end of this step. Subtracting the relevant lines of (A.67) and (A.68),

$$\Delta(c) = \big(\bar{\theta}_A^N - \bar{\theta}_A^I\big) + \tfrac{1}{2}\big(1 - \bar{\theta}_B^I\big)^2. \tag{A.69}$$

The second term is strictly positive, because $\bar{\theta}_B^I < 1$ for $c < \frac{1}{2}$. It remains to sign the first term.

In both regimes, the marginal active type is indifferent between no high effort and high effort once. By (A.33) and by (A.64) with the posteriors (A.61) and (A.62),

$$\text{(NoIPE)} \qquad c = \bar{\theta}_A^N \big[m_1(\bar{\theta}_A^N, 1) - m_0(\bar{\theta}_A^N, 1)\big] = \Gamma(\bar{\theta}_A^N), \tag{A.70}$$

$$\text{(IPE)} \qquad c = \bar{\theta}_A^I \big[m_1(\bar{\theta}_A^I, 1) - m_0(\bar{\theta}_A^I, \bar{\theta}_B^I)\big]. \tag{A.71}$$

Two features are essential. First, the success posterior $m_1(\cdot, 1)$ is the same function of the lowest active type in both regimes, because every active type generates the outcome with a single success in the period it uses: under NoIPE all active types exert effort once, and under IPE the $(1;0)$ types do so directly, while the conditional types generate $(1,0)$ by succeeding in the first period and then stopping. Second, the failure posteriors are ordered by Lemma A2. Applying the lemma with $a = \bar{\theta}_A^I$ and $b = \bar{\theta}_B^I$, which is legitimate since $0 < \bar{\theta}_A^I < \bar{\theta}_B^I < 1$, gives

$$c = \bar{\theta}_A^I \big[m_1(\bar{\theta}_A^I, 1) - m_0(\bar{\theta}_A^I, \bar{\theta}_B^I)\big] > \bar{\theta}_A^I \big[m_1(\bar{\theta}_A^I, 1) - m_0(\bar{\theta}_A^I, 1)\big] = \Gamma(\bar{\theta}_A^I).$$

Hence $\Gamma(\bar{\theta}_A^I) < c = \Gamma(\bar{\theta}_A^N)$, and since Γ is strictly increasing (Appendix A2),

$$\bar{\theta}_A^I < \bar{\theta}_A^N.$$

Substituting into (A.69), both terms are positive, so

$$\Delta(c) = \big(\bar{\theta}_A^N - \bar{\theta}_A^I\big) + \tfrac{1}{2}\big(1 - \bar{\theta}_B^I\big)^2 > 0 \qquad \text{for all } c \in [c_I^H, \tfrac{1}{2}).$$

It remains to show that $c_I^H > \frac{2}{9}$. At $c = c_I^H$, the IPE thresholds satisfy (A.71) with $0 < \bar{\theta}_A^I < \bar{\theta}_B^I < 1$, so the argument above gives $\Gamma(\bar{\theta}_A^I) < c_I^H$, that is, $\bar{\theta}_A^I < \Gamma^{-1}(c_I^H)$. This step

uses only Lemma A2 and the function Γ, and does not depend on the NoIPE region. At the same cost, $c_I^H = 1 - m_1(\bar{\theta}_A^I, 1)$ (Appendix A3). Since $m_1(a, 1)$ is strictly increasing in a,

$$c_I^H = 1 - m_1(\bar{\theta}_A^I, 1) > \varphi(c_I^H), \qquad \varphi(c) \equiv 1 - m_1\big(\Gamma^{-1}(c), 1\big).$$

The function φ is strictly decreasing, and $\varphi(\frac{2}{9}) = \frac{2}{9}$, because $\Gamma(\frac{1}{2}) = \frac{2}{9}$ and $m_1(\frac{1}{2}, 1) = \frac{7}{9}$. If $c_I^H \leq \frac{2}{9}$, then $\varphi(c_I^H) \geq \frac{2}{9} \geq c_I^H$, which contradicts the display. Hence $c_I^H > \frac{2}{9}$. Numerically, $c_I^H \approx 0.232$.

Thus expected total effort is strictly higher under IPE on the whole high-cost range, which establishes the analytic dominance statement of Proposition 3. Both terms in (A.69) reflect the conditional-continuation effect described after Proposition 3. Because conditional types try again after a first-period failure, the failure posterior is lower for a given lowest active type (Lemma A2), so the lowest type willing to signal is lower under IPE ($\bar{\theta}_A^I < \bar{\theta}_A^N$, the extensive-margin effect). In addition, the conditional types add the expected second-period effort $\frac{1}{2}(1 - \bar{\theta}_B^I)^2$.

Step 4: Numerical comparison for $c \in (0, c_I^H)$. For $c < \frac{2}{9}$, both regimes are described by the first lines of (A.67) and (A.68), and

$$\Delta(c) = \underbrace{\big(\bar{\theta}_A^N + \bar{\theta}_B^N\big) - \big(\bar{\theta}_A^I + \bar{\theta}_B^I\big)}_{\text{threshold margin}} - \underbrace{\tfrac{1}{2}\big[(\bar{\theta}_C^I)^2 - (\bar{\theta}_B^I)^2\big]}_{\text{early-stopping loss} \geq 0}. \tag{A.72}$$

The early-stopping loss records the second-period effort that conditional types forgo after a first-period success, a possibility that exists only under IPE. It is the stopping effect described after Proposition 3, and it makes NoIPE the more effort-intensive regime at low costs: when effort is cheap, a large mass of high types exerts effort in both periods under NoIPE, whereas under IPE many of these types stop after an early success. The threshold margin works in the opposite direction: in the numerical solutions, $\bar{\theta}_A^I + \bar{\theta}_B^I < \bar{\theta}_A^N + \bar{\theta}_B^N$ for every $c < \frac{2}{9}$. At low costs, the two terms are of similar size and nearly offset. The thresholds solve coupled polynomial systems without closed-form solutions, so I evaluate Δ numerically on $(0, c_I^H)$. The difference is negative on $(0, c^*)$ and positive on (c^*, c_I^H), with a unique sign change at

$$c^* \approx 0.135.$$

Close to zero, the NoIPE advantage is very small, for example $\Delta(0.04) \approx -2.1 \times 10^{-3}$ and $\Delta(0.01) \approx -2.3 \times 10^{-6}$. This is why the two total-effort curves are visually indistinguishable at low costs in Figure 3. The largest NoIPE advantage is $\max_c(-\Delta(c)) \approx 0.024$, attained near

$c \approx 0.10$, whereas the largest IPE advantage is $\max_c \Delta(c) \approx 0.18$, attained near $c \approx 0.22$. At the experimental cost $c = 0.3$, the model predicts strictly higher total effort under IPE ($\mathcal{E}^{IPE} \approx 0.46$ versus $\mathcal{E}^{NoIPE} \approx 0.36$).

Conclusion. Steps 1 and 3 establish the analytic statements of Proposition 3. By Step 2, Δ is continuous on $(0, \frac{1}{2})$, and by Steps 3 and 4 it is negative for small c and positive on $[c_I^H, \frac{1}{2})$, so it changes sign at least once in $(0, c_I^H)$. The numerical solutions in Step 4 show that this sign change is unique, so the two cost ranges of Proposition 3 partition the active region $(0, \frac{1}{2})$ at $c^* \approx 0.135$, and the effort ranking reverses exactly once. □

B Further Experimental Results

B1 Answers in Quiz treatments

Table B1: Agent's Answers in Quiz

	#Correct Period 1	#Correct Period 2
	(1)	(2)
Agent Type	0.291***	0.193***
	(0.058)	(0.051)
Round	-0.034	-0.027
	(0.022)	(0.020)
IPE	-0.150	-0.195*
	(0.109)	(0.117)
Effort Period 1	0.552**	
	(0.210)	
Effort Period 2		0.472**
		(0.199)
const.	4.331***	4.820***
	(0.319)	(0.310)
Obs.	790	790
Clusters	79	79
R^2	0.152	0.081

Notes: OLS regressions of the number of correct answers in the quiz task in the first period of a round (1) and the second period of a round (2). The maximum of correct answers in a period (and the maximum value of type) is 10, and 10 rounds were played. Standard-errors are clustered at the agent level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

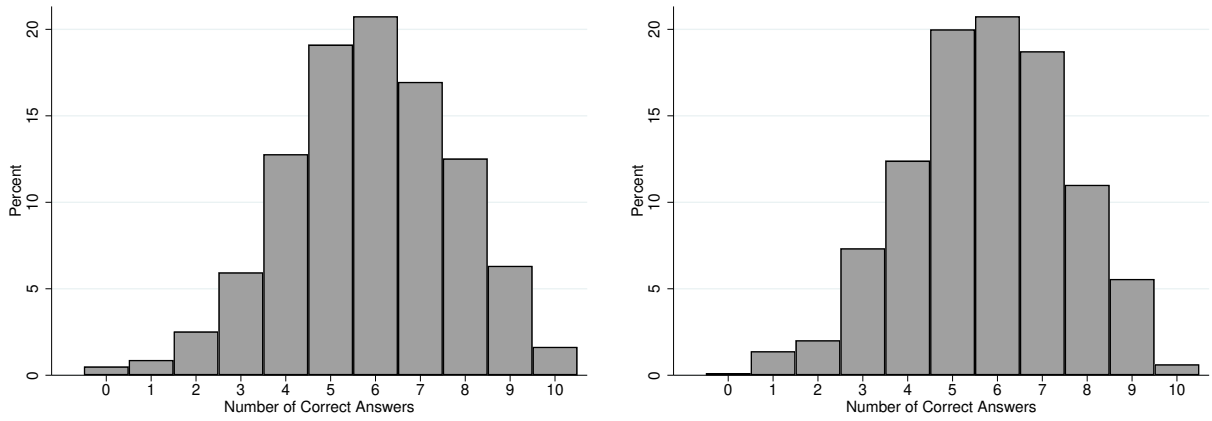

(a) First Period (b) Second Period

Figure B1: Distribution of number of correct Quiz answers by the agents in the first (a) and second (b) period of a round.

B2 Effort after Success in IPE

Table B2: Second Period Effort after Success (IPE)

	Abstract		Quiz	
	(1)	(2)	(3)	(4)
Type of Agent	0.009***	0.009***	0.087**	0.080**
	(0.002)	(0.002)	(0.034)	(0.036)
Fairness	0.053***	0.050***	-0.017	-0.019
	(0.016)	(0.017)	(0.032)	(0.033)
Math Grade	0.036	0.041	0.071	0.075
	(0.036)	(0.034)	(0.053)	(0.054)
Male		-0.086		0.078
		(0.060)		(0.096)
endogenous		-0.104		-0.062
		(0.099)		(0.073)
const.	-0.384*	-0.275	0.024	0.081
	(0.195)	(0.224)	(0.286)	(0.282)
Obs.	174	165	97	97
Clusters	17	17	16	16
R^2	0.137	0.155	0.147	0.157

Notes: OLS regressions of the second period effort provision in a round in the Abstract ((1) and (2)) and the Quiz ((3) and (4)) treatments. Only data of agents who were successful in the first period and played in the IPE treatment is used. *Fairness* $\in [1, 7]$ denotes the answer to above mentioned question. *Math Grade* is self reported and ranges from 1 (good) to 6 (bad). *endogenous* denotes the rounds where another principal (instead of the computer) chose whether IPE is played. *Male* is a dummy for the participant's self-reported gender. Participants who stated to not be either male or female are excluded from columns (2) and (4). Standard-errors are clustered at the matching-group level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

B3 Behavior over Time

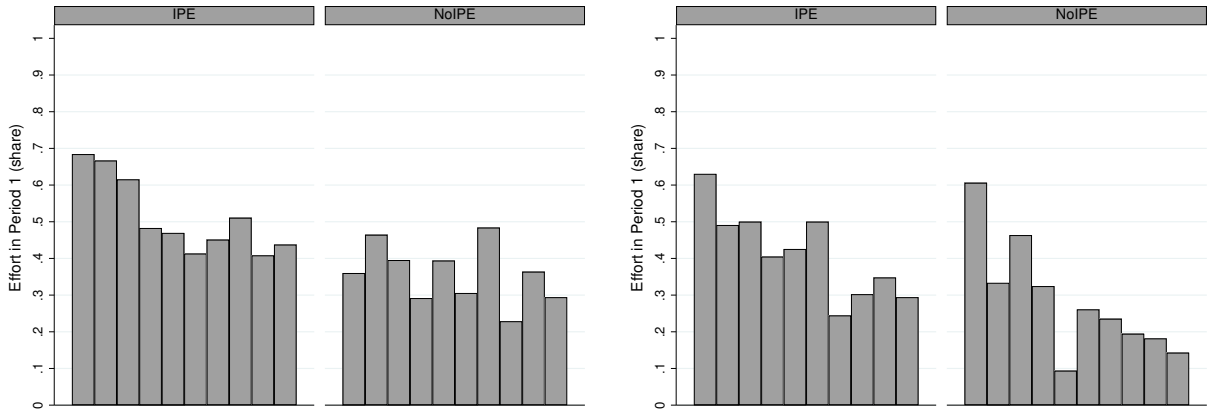

(a) Abstract Treatments (b) Quiz Treatments

Figure B2: First period effort provision over time.

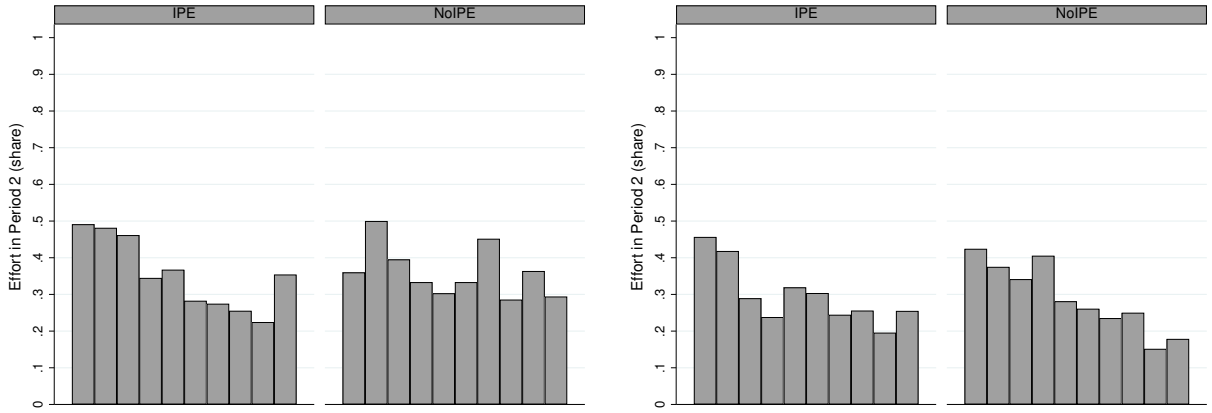

(a) Abstract Treatments (b) Quiz Treatments

Figure B3: Second period effort provision over time.

B4 Explanation for IPE Choice of Principals

In the post-experimental questionnaire, principals answered the non-incentivized free-form question: *"What motivated your decision whether Player A receives feedback about the success in period 1?"*. In Table B3, I provide translations of all non-empty answers, sorted into categories that seem to stand out. For some answers, the categorization is debatable. For this reason, I transparently provide the verbatim content in translation, and not just counts by category based on my own judgment.

Table B3: Principal's Answers to Open IPE-Choice Question

Context	**Participant's answer about motivation to (not) provide feedback**
Category: IPE is better for Agent	
Quiz	I thought that if member A receives feedback and learns that there was no success, he or she might not be willing to invest the 30 points.
Quiz	I thought that if member A gets no information, the probability is higher that he or she will keep paying for success.
Quiz	I thought it was only fair to inform member A in every round (1 to 5) about the success in period 1.
Quiz	I found no reason for the decision, so I always decided that member A should receive feedback.
Quiz	I saw hardly any point here; I always let A know about success because I support information sharing.
Quiz	I always chose feedback because it seemed like the better option and had no effect on my payoff.
Quiz	I found it fair that A receives my assessment.
Quiz	In the end it did not matter; the other person could also learn that they had, of course, not been successful.
Quiz	There is no reason, neither for A nor for me, to withhold the feedback. One could say that, since the As were usually so bad and perhaps did not realize it, B might have an advantage in letting them believe they still had a chance, so that they pay the 30 points and B has a chance as well, but anyway.

Table B3: Principal's Answers to Open IPE-Choice Question (continued)

Context	**Participant's answer about motivation to (not) provide feedback**
Quiz	I am a curious person myself and would want to know how I performed. So I transferred this to the other person.
Quiz	I always chose feedback because I think member A can recognize in this way whether the answer was wrong or correct and then make fewer mistakes afterwards.
Quiz	So that A can better assess in the second period whether period 2 is likely to lead to success or not.
Quiz	I always preferred feedback because then member A can adjust the decision of whether to invest the 30 points accordingly.
Quiz	I always chose feedback because I personally would have wanted it that way if I had been A.
Quiz	I thought he should always receive feedback so that he can assess how to proceed.
Abstract	I always allowed feedback because I think that with feedback A has more information to judge whether it is worth investing in the next period. In a sense, this gives an additional probability assessment. Perhaps it increases the investment rate.
Abstract	I usually made feedback visible to member A. I think it would have been more worthwhile for member A to try to achieve success in at least one of the rounds. If I had been member A, I would only have paid the cost in the second round if I had known that I had had no success in the first round despite a high type percentage. Therefore I wanted member A to receive feedback.
Abstract	A would surely like to be informed about it.
Abstract	Since a new person was assigned each time, I automatically chose feedback. The more information, the more precise the assessment of the chances of success can be.
Abstract	If he or she had more information, he or she could make better decisions. Therefore always with feedback.
Abstract	I thought that I myself would also like to know the feedback, which is why I almost always decided to communicate it.

Table B3: Principal's Answers to Open IPE-Choice Question (continued)

Context	Participant's answer about motivation to (not) provide feedback
Abstract	I always chose that he receives feedback so that he knows whether his strategy worked or whether he should proceed differently.
Abstract	Information is always helpful, so I generally allowed feedback. Whether this affected player A's decision, I cannot judge.
Category: IPE is (not) better for Principal	
Quiz	No feedback → higher chance that more periods are played, even if only marginally.
Quiz	I wanted it so that it would weigh in favor of success.
Quiz	At first I thought feedback is good, so the person can also choose the chance in the second round, but then I thought it could also be discouraging if someone sees they had no success. So I experimented a bit and tried both, since the partners changed anyway.
Quiz	I hoped for higher chances of success this way, also for A through direct feedback.
Quiz	I wrote down how the participants decided and acted more or less intuitively.
Quiz	I gave feedback only in the first round; after that I wanted to continue without it.
Quiz	In the first five rounds I had the impression that I received more points when A did not get feedback.
Quiz	I always gave feedback so that the people can see how bad they are at answering quiz questions.
Abstract	I consistently chose no feedback because all A-participants had decided terribly anyway.
Abstract	In favor of feedback, because then period 2 will probably still be paid for if no success occurs, so that there is at least one success overall.
Abstract	With feedback, A is less willing to pay for the chance.
Abstract	Better adjustment of my own estimate.
Abstract	Since if member A succeeds, then I do as well. If one receives information, there is a better chance of finding the "right" solution.

Table B3: Principal's Answers to Open IPE-Choice Question (continued)

Context	Participant's answer about motivation to (not) provide feedback
Abstract	I felt that if member A receives feedback, he or she is more likely to invest, so I often selected feedback.
Abstract	If successful, then more willing to bet on success again in round 2.
Abstract	I usually gave feedback because it increased the chance that the 30 points would be paid in the second stage.
Abstract	Generally positive feedback so that A is also motivated to invest for success.
Abstract	I gave A feedback four times because I thought it might motivate A to invest the 30 points again in period 2 if A sees that period 1 was already successful. The one time I decided against feedback, I simply wanted to test what effect that would have on the outcome.
Abstract	I thought it was only fair and hoped to increase the probability of success that way, because teamwork makes the dream work and all that.
Abstract	I thought that if A received feedback, A would be less willing to take risks if A had unsuccessfully paid the 30 points in period 1. It was in my interest that A pays the 30 points whenever possible, so I did not want A to see the feedback.
Abstract	So that he sees from a decision not to invest that I assign 0 points to it.
Abstract	Previous results; possible motivators.
Abstract	On member A's success.
Abstract	Two times yes, more often no, without better winning chances.
Category: Random / no clear strategy / intuition	
Quiz	Gut feeling.
Quiz	Nothing.
Quiz	Nothing in particular.
Quiz	I had no strategy.
Quiz	Random.
Quiz	Whether it makes a difference.
Quiz	In fact, I regarded this as unimportant.

Table B3: Principal's Answers to Open IPE-Choice Question (continued)

Context	**Participant's answer about motivation to (not) provide feedback**
Abstract	Nothing.
Abstract	That was rather random, since I did not see higher or lower chances of success with or without feedback.
Abstract	Just like my answer to the previous question.
Abstract	Just like that.
Abstract	I gave feedback randomly, and I think one really could not do much with feedback.
Abstract	I still do not know whether feedback made any difference at all, so I just did it randomly.
Abstract	To be honest, I could not think of anything that spoke for or against feedback, so I decided relatively randomly in favor of feedback.
Abstract	It depended somewhat on the round payoff; otherwise I simply decided arbitrarily whether with or without feedback.
Abstract	Pure chance principle.
Category: Always feedback	
Quiz	Always feedback.
Quiz	I always let the person receive feedback.

Supplementary Material

SM1 Decision Screens

Translations of the German texts (from top to bottom of each screen) are provided in the figure notes.

SM1 Abstract Treatments

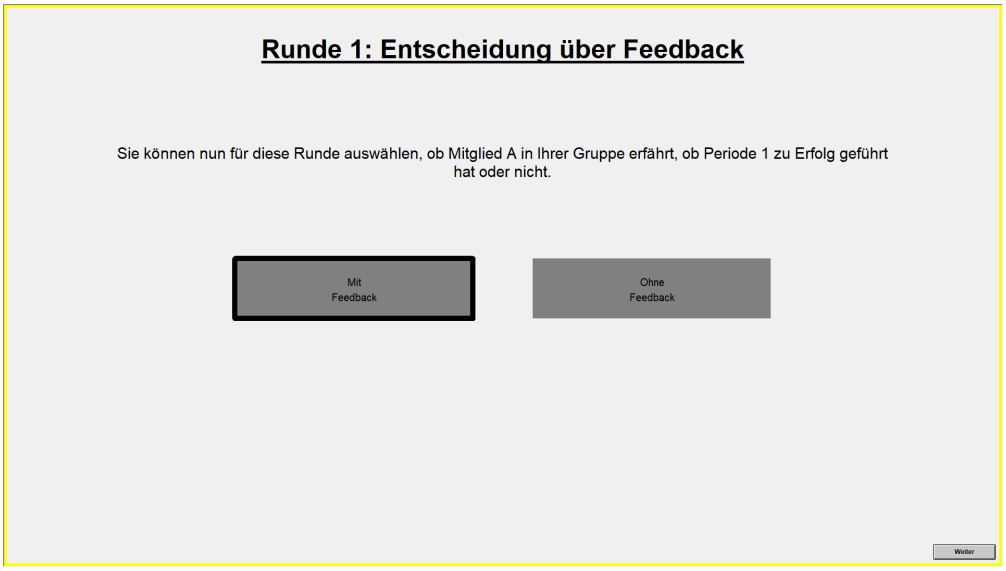

Figure SM1: Endogenous Treatment, Principal.

Notes: "Round 1: Decision about Feedback. You can choose for this round, whether Member A of your group learns whether or not Success was achieved in Period 1. With feedback. Without feedback."

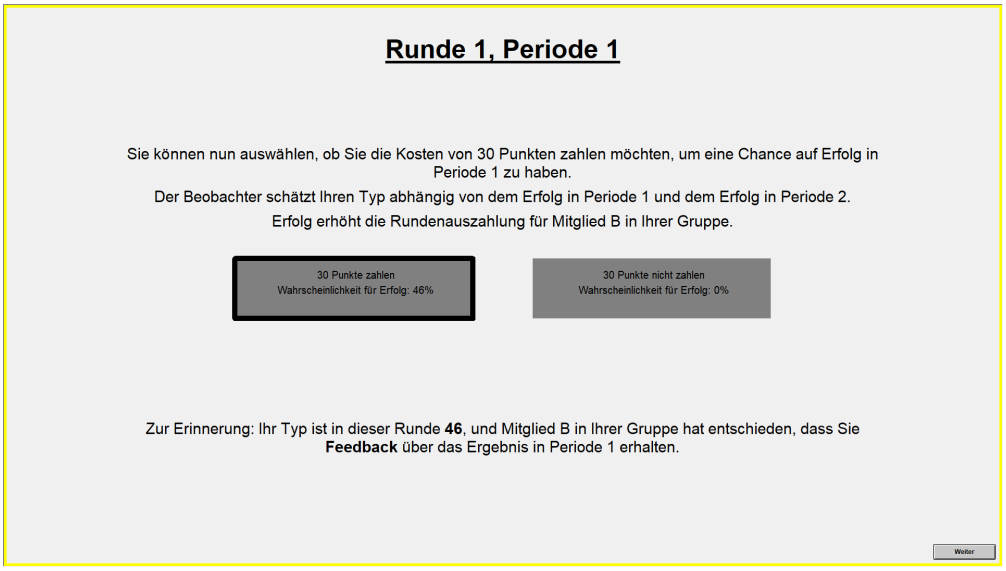

Figure SM2: Endogenous Treatment, Agent.

Notes: "Round 1, Period 1. You can now choose whether you want to pay the costs of 30 points in order to have a chance for Success in Period 1. The observer guesses your type dependent on your Success in Period 1 and your Success in Period 2. Success increases the roundpayoff of Member B of your group. Paying 30 points. Probability for Success: 46%. Not paying 30 points. Probability for Success: 0%. Reminder: In this round your type is 46, and Member B of your Group decided, that you receive feedback about the Success in Period 1."

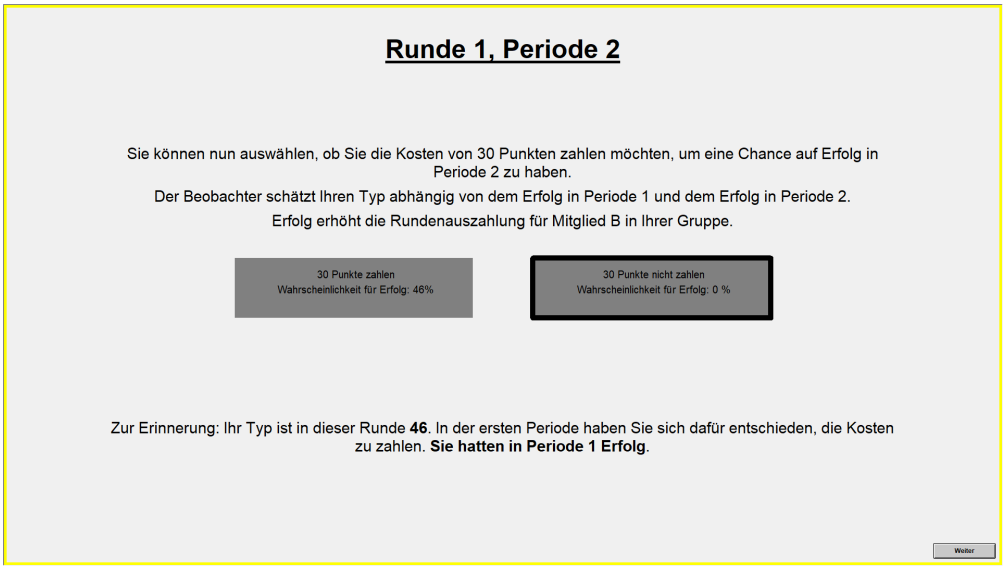

Figure SM3: IPE treatment, Agent.

Notes: "Round 1, Period 2. You can now choose whether you want to pay the costs of 30 points in order to have a chance for Success in Period 2. The observer guesses your type dependent on your Success in Period 1 and your Success in Period 2. Success increases the roundpayoff of Member B of your group. Paying 30 points. Probability for Success: 46%. Not paying 30 points. Probability for Success: 0%. Reminder: In this round your type is 46. You decided to pay the costs in the first Period. You had Success in the first Period."

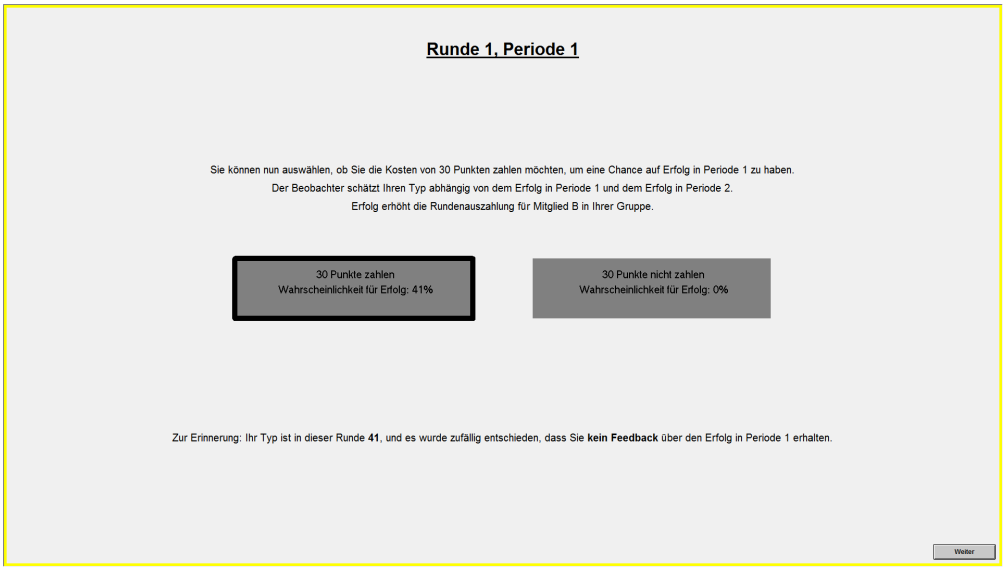

Figure SM4: Exogenous Treatment, Agent.

Notes: "Round 1, Period 1. You can now choose whether you want to pay the costs of 30 points in order to have a chance for Success in Period 1. The observer guesses your type dependent on your Success in Period 1 and your Success in Period 2. Success increases the roundpayoff of Member B of your group. Paying 30 points. Probability for Success: 41%. Not paying 30 points. Probability for Success: 0%. Reminder: In this round your type is 41, and it was randomly determined that you receive no feedback about the Success in Period 1."

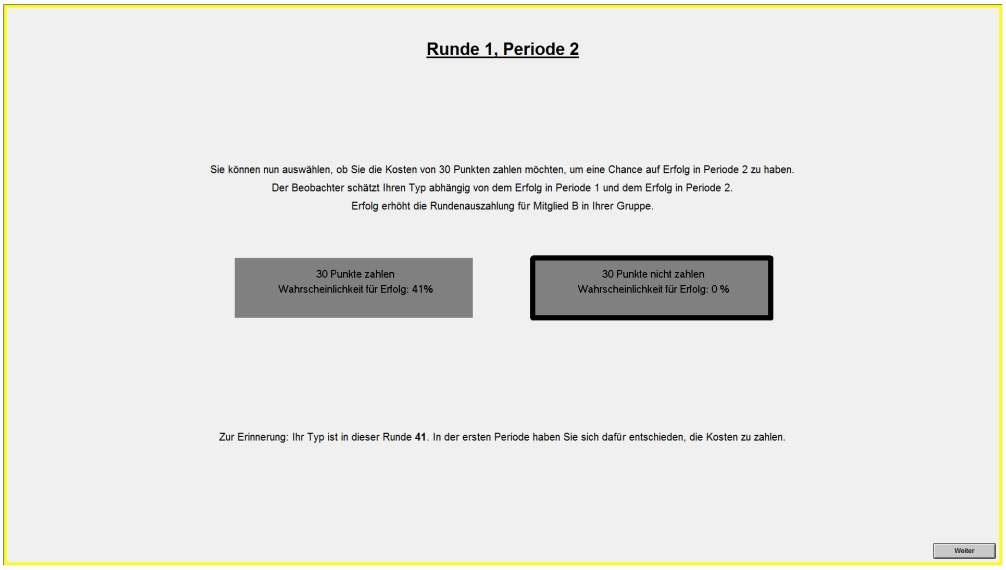

Figure SM5: NoIPE Treatment, Agent.

Notes: "Round 1, Period 2. You can now choose whether you want to pay the costs of 30 points in order to have a chance for Success in Period 2. The observer guesses your type dependent on your Success in Period 1 and your Success in Period 2. Success increases the roundpayoff of Member B of your group. Paying 30 points. Probability for Success: 41%. Not paying 30 points. Probability for Success: 0%. Reminder: In this round your type is 41. You decided to pay the costs in the first Period."

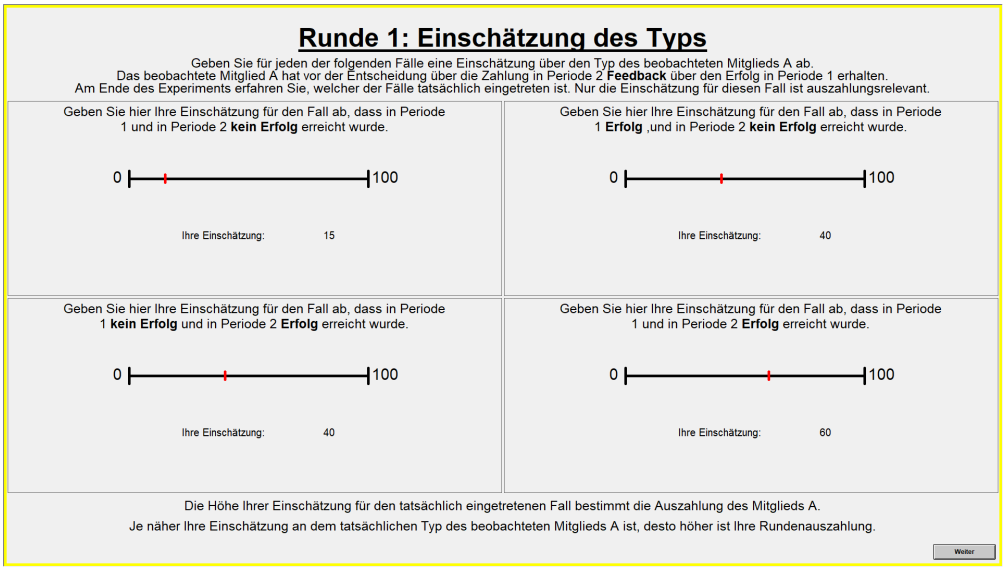

Figure SM6: Observer.

Notes: "Round 1: Type Assessment. Enter for each of the following cases an estimate about Member A's type. The observed Member A received Feedback about the Period 1 Success prior to the decision whether to pay in Period 2. At the end of the experiment you will learn which of the cases indeed occurred. Only the assessment for this case is payoff-relevant. Enter your assessment for the case that no Success was achieved in Period 1 and Period 2. Your Assessment: 15. Enter your assessment for the case that no Success was achieved in Period 1 and Success was achieved in Period 2. Your Assessment: 40. Enter your assessment for the case that Success was achieved in Period 1 and no Success was achieved in Period 2. Your Assessment: 40. Enter your assessment for the case that Success was achieved in Period 1 and Period 2. Your Assessment: 60. The level of your assessment for the relevant case determines the payoff of Member A. The closer your assessment to the true type of the observed Member A, the higher is your round payoff."

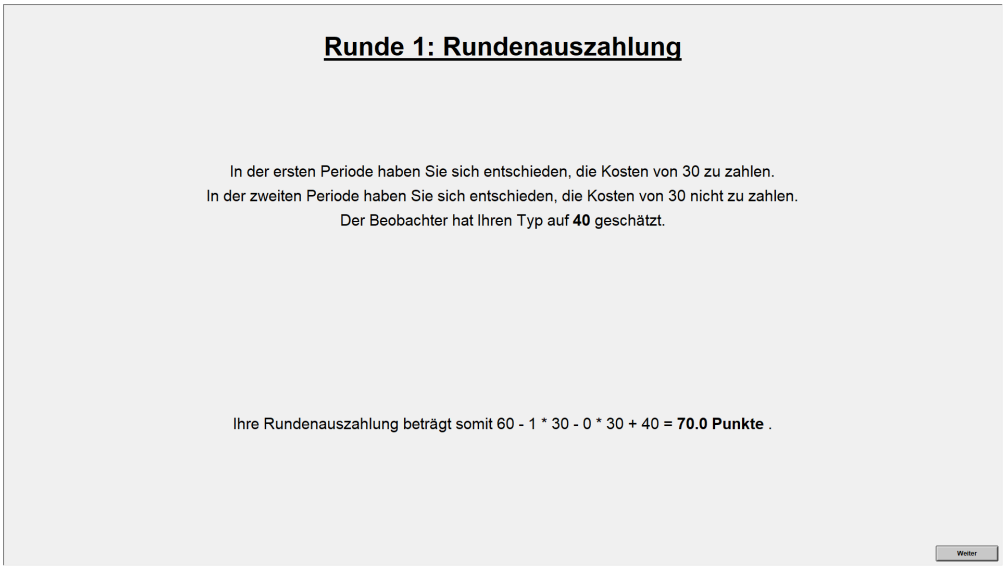

Figure SM7: Agent.

Notes: "Round 1: roundpayoff. In the first period you decided to pay the costs of 30. In the second Period you decided not to pay the costs of 30. The observer estimated your type to be 40. Your roundpayoff is therefore: $60 - 1 * 30 - 0 * 30 + 40 = 70$ Points."

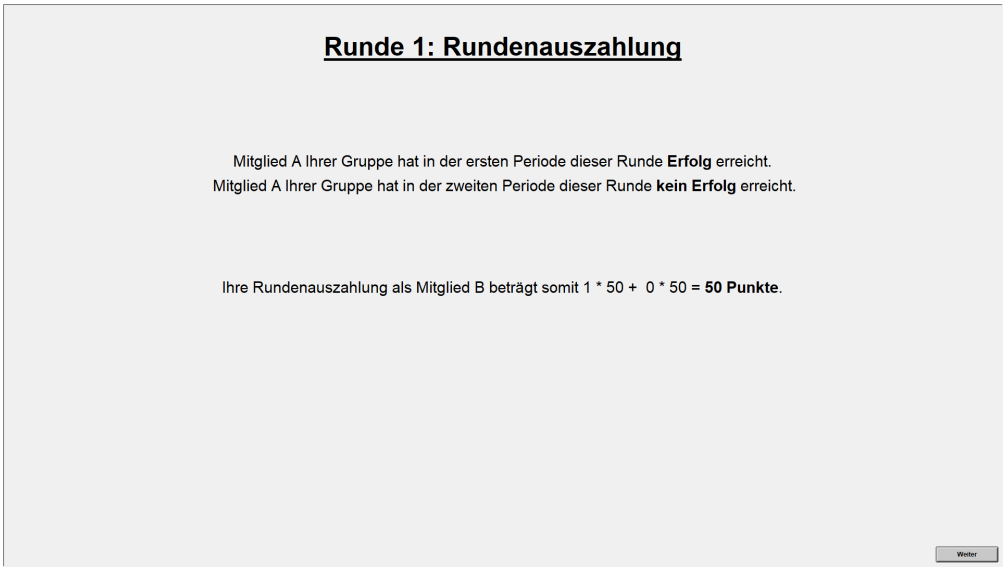

Figure SM8: Principal.

Notes: "Round 1: Roundpayoff. Member A in your Group achieved Success in the first Period of this round. Member A of your group achieved no Success in the second Period of this round. Your roundpayoff as Member B is therefore: $1 * 50 + 0 * 50 = 50$ Points."

SM2 Quiz Treatments

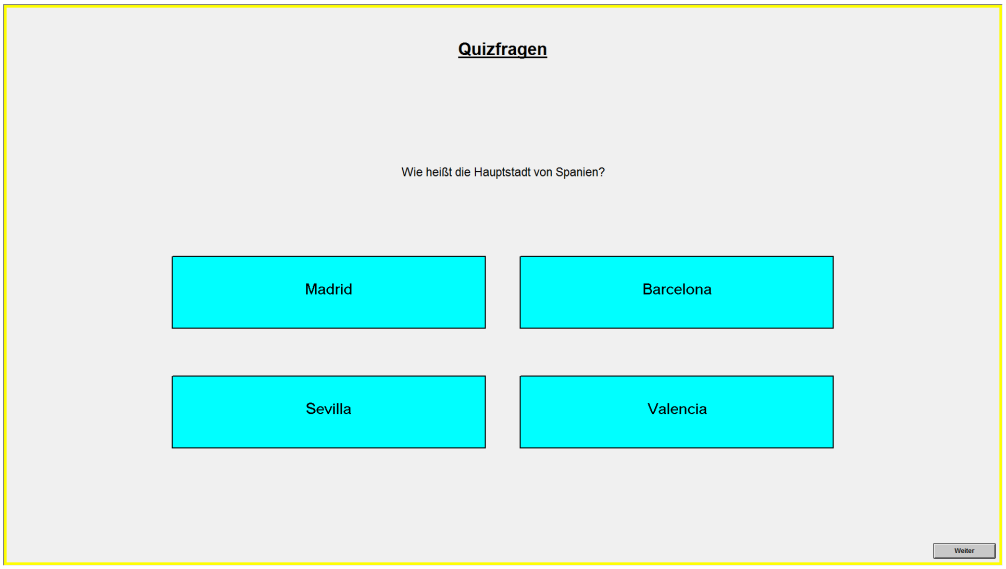

Figure SM9: Quiz questions.

Notes: "What is the capital of Spain? Madrid. Barcelona. Sevilla. Valencia."

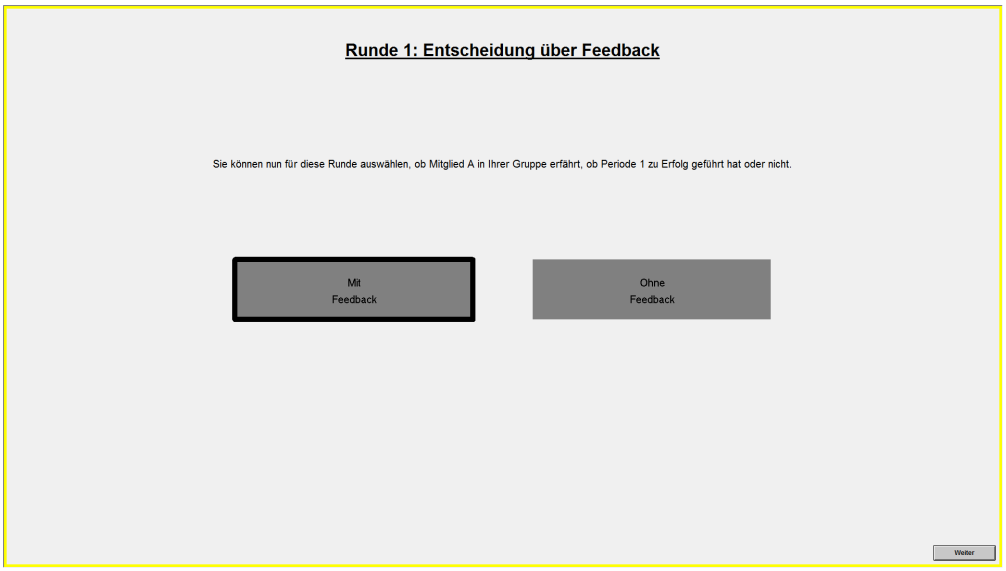

Figure SM10: Endogenous treatments, Principal.

Notes: "Round 1: Decision over feedback. You can now choose for this round whether member A from your group learns whether period 1 resulted in success or not. With feedback. Without feedback."

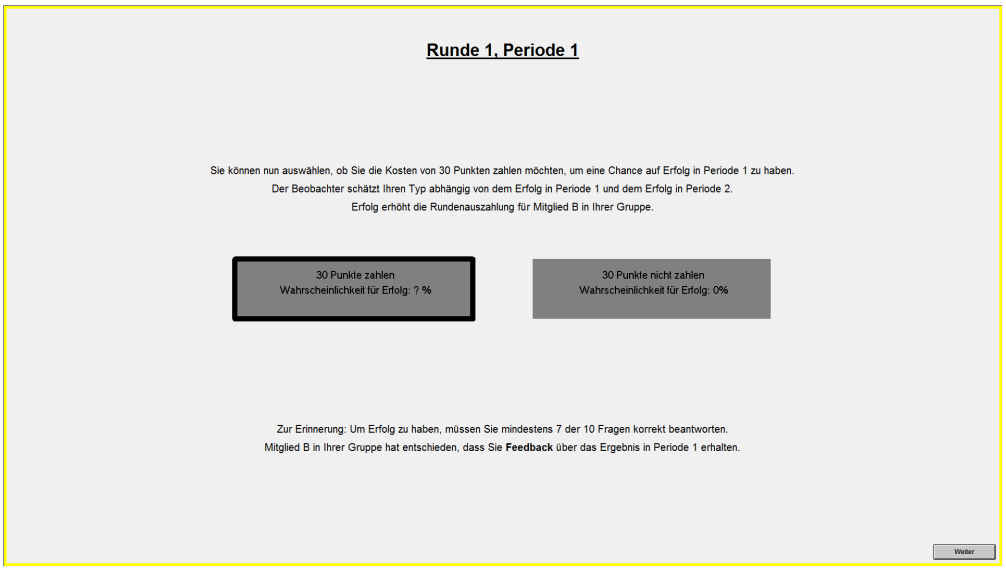

Figure SM11: Endogenous treatments, Agent.

Notes: "Round 1, Period 1. You can now choose whether you want to pay the costs of 30 points in order to have a chance for Success in Period 1. The observer guesses your type dependent on your Success in Period 1 and your Success in Period 2. Success increases the roundpayoff of Member B of your group. Paying 30 points. Probability for Success: ?%. Not paying 30 points. Probability for Success: 0%. Reminder: In order to have success, you need to correctly answer at least seven out of the 10 questions. Member B of your group determined that you receive feedback about the result in Period 1."

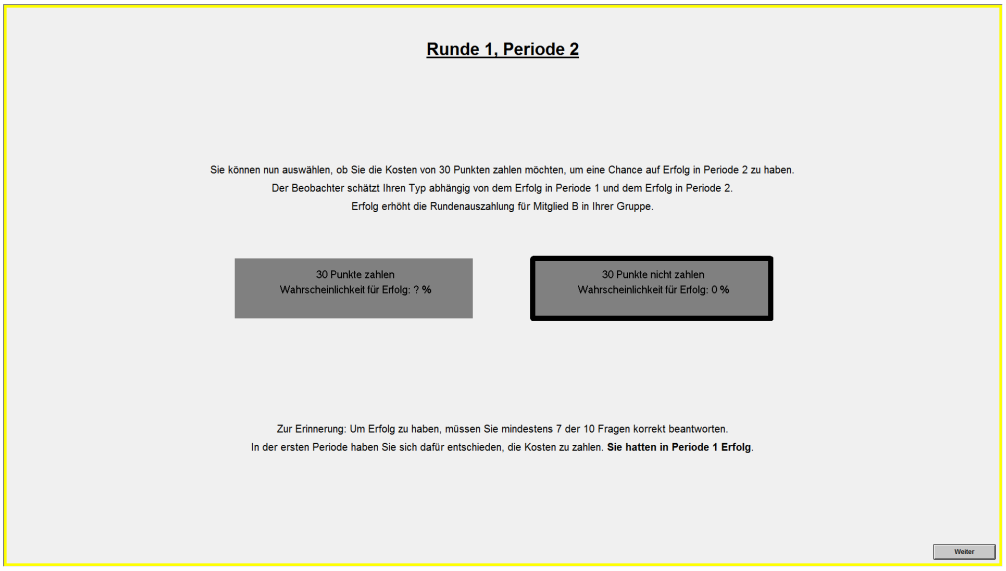

Figure SM12: IPE treatment, agent.

Notes: "Round 1, Period 2. You can now choose whether you want to pay the costs of 30 points in order to have a chance for Success in Period 2. The observer guesses your type dependent on your Success in Period 1 and your Success in Period 2. Success increases the roundpayoff of Member B of your group. Paying 30 points. Probability for Success: ?%. Not paying 30 points. Probability for Success: 0%. Reminder: In order to have success, you need to correctly answer at least seven out of the 10 questions. You decided to pay the costs in the first Period. You had Success in period 1."

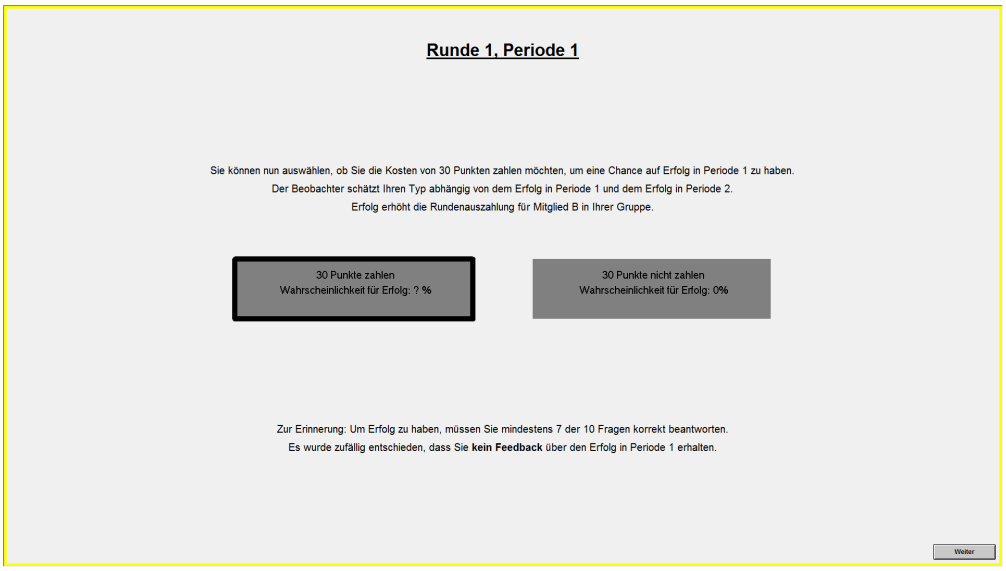

Figure SM13: Exogenous treatments, Agent.

Notes: "Round 1, Period 1. You can now choose whether you want to pay the costs of 30 points in order to have a chance for Success in Period 1. The observer guesses your type dependent on your Success in Period 1 and your Success in Period 2. Success increases the roundpayoff of Member B of your group. Paying 30 points. Probability for Success: ?%. Not paying 30 points. Probability for Success: 0%. Reminder: In order to have success, you need to correctly answer at least seven out of the 10 questions. It was randomly determined that you receive no feedback about the success in Period 1."

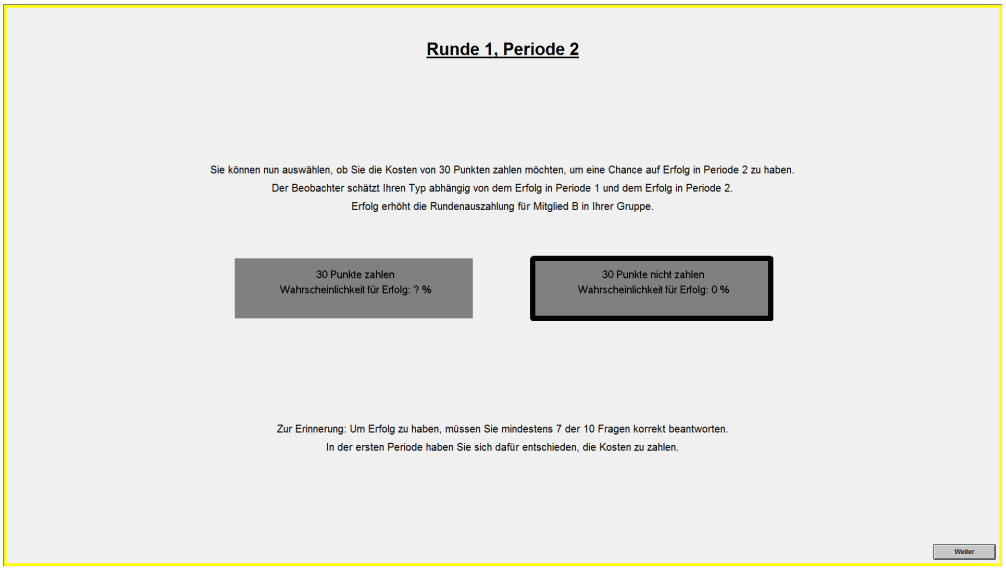

Figure SM14: NoIPE, Agent.

Notes: "Round 1, Period 2. You can now choose whether you want to pay the costs of 30 points in order to have a chance for Success in Period 2. The observer guesses your type dependent on your Success in Period 1 and your Success in Period 2. Success increases the roundpayoff of Member B of your group. Paying 30 points. Probability for Success: ?%. Not paying 30 points. Probability for Success: 0%. Reminder: In order to have success, you need to correctly answer at least seven out of the 10 questions. You decided to pay the costs in the first Period."

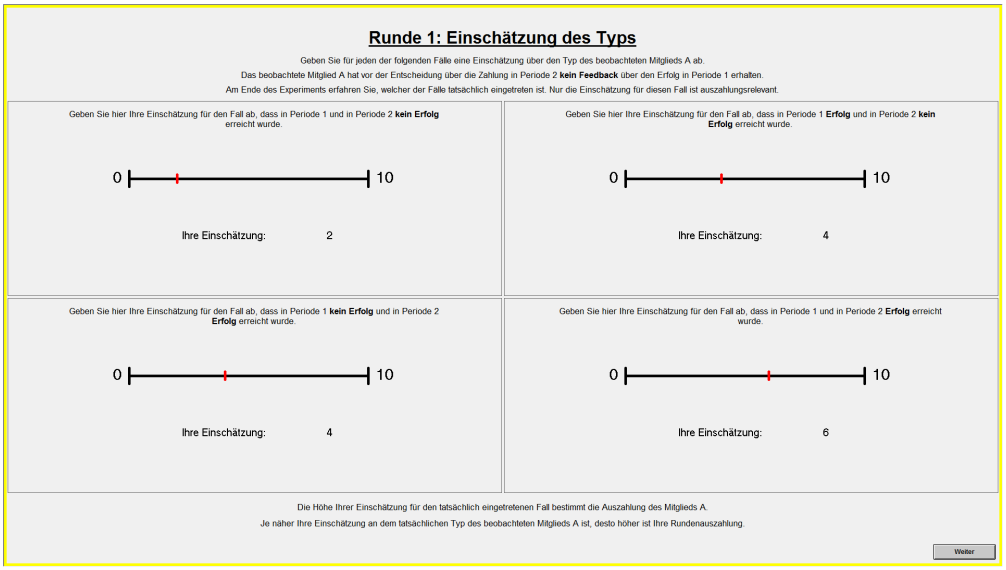

Figure SM15: Observer.

Notes: "Round 1: Type Assessment. Enter for each of the following cases an estimate about Member A's type. The observed Member A received no Feedback about the Period 1 Success prior to the decision whether to pay in Period 2. At the end of the experiment you will learn which of the cases indeed occurred. Only the assessment for this case is payoff-relevant. Enter your assessment for the case that no Success was achieved in Period 1 and Period 2. Your Assessment: 2. Enter your assessment for the case that no Success was achieved in Period 1 and Success was achieved in Period 2. Your Assessment: 4. Enter your assessment for the case that Success was achieved in Period 1 and no Success was achieved in Period 2. Your Assessment: 4. Enter your assessment for the case that Success was achieved in Period 1 and Period 2. Your Assessment: 6. The level of your assessment for the relevant case determines the payoff of Member A. The closer your assessment to the true type of the observed Member A, the higher is your round payoff."



Figure SM16: Agent.

Notes: "Round 1: Roundpayoff. In the first period you decided to pay the costs of 30. In the second Period you decided not to pay the costs of 30. The observer estimated your type to be 4. Your roundpayoff is therefore: $60 - 1 * 30 - 0 * 30 + 10 * 4 = 70$ Points."

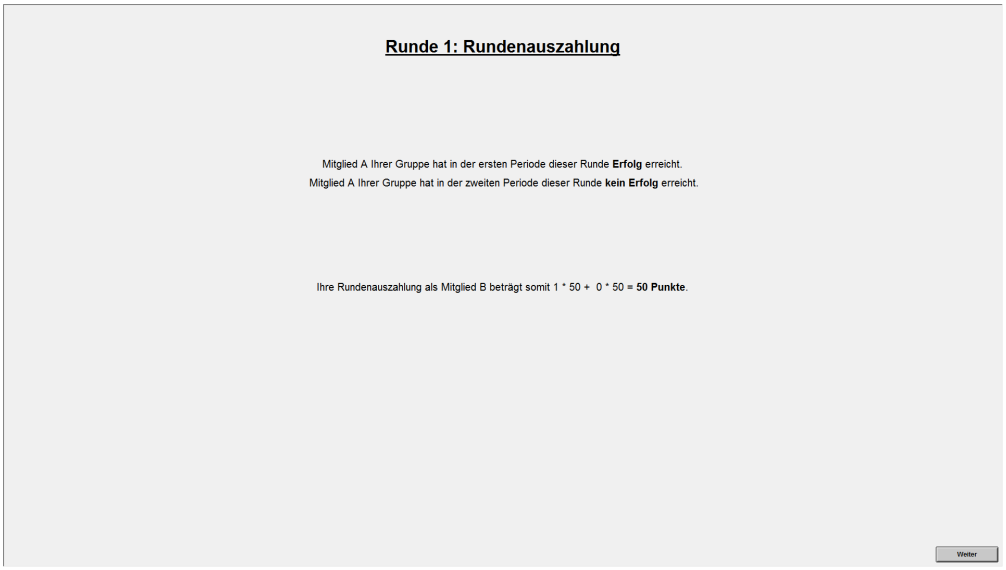

Figure SM17: Principal.

Notes: "Round 1: Roundpayoff. Member A in your Group achieved Success in the first Period of this round. Member A of your group achieved no Success in the second Period of this round. Your roundpayoff as Member B is therefore: $1 * 50 + 0 * 50 = 50$ Points."

SM2 Experimental Instructions

An English translation of the Instructions for the Quiz treatments (endogenous IPE first) as well as for the Abstract treatments (exogenous IPE first) is provided. The instructions of the other treatments (Quiz, exogenous IPE first; Abstract, endogenous first) are very similar and omitted here. The original German instructions are part of the replication material.

SM 2.1 Quiz Treatment

Overview

Welcome to this experiment. We ask you not to speak with other participants during the experiment and to switch off your mobile phones and other mobile electronic devices.

For participating in today's session, the amount you earn will be paid out to you in cash at the end of the experiment. The size of the payment depends partly on your decisions, partly on the decisions of other participants, and partly on chance. It is therefore important that you carefully read and understand the instructions before the experiment begins.

In this experiment, every interaction between participants takes place via the computers in front of which you are seated. You will interact with each other anonymously and your decisions will only be stored together with your random ID number. Neither your name nor the names of other participants will be revealed, neither today nor in future written analyses.

Today's session consists of several rounds. At the end, a total of 2 or 3 rounds will be randomly selected and paid out. The rounds that are not selected will not be paid. Your payment amount results from the points earned in the selected rounds, converted into euros, as well as your show-up fee of 6 euros. The conversion of points into euros is as follows: Each point is worth 4 cents, so that:

$$25 \text{ Points} = 1 \text{ Euro}$$

Each participant will be paid privately so that other participants cannot see how much you have earned.

Experiment

The experiment consists of two parts with a total of 10 rounds. The rounds within Parts 1 (Rounds 1–5) and 2 (Rounds 6–10) are identical in structure. The instructions for the second part will be provided after the completion of the first part.

At the beginning of the experiment, you will be randomly assigned either the role of **Member A** or **Member B / Observer**. Each participant keeps this role for the entire experiment.

The Group

In the first round, you will be assigned to a large group of 8, 10, or 12 members. At the beginning of each round, two participants from your respective large group will be randomly selected to form a group in that round. Each group consists of one Member A and one Member B. The members of a group are therefore randomly reassigned in every round.

General Structure of a Round

Each round consists of two periods. In each period, each Member A answers 10 quiz questions about the capital city of randomly selected countries. Before Member A sees the questions, they decide whether to pay a cost of 30 points in order to have a chance of success in that period. Success occurs if the cost of 30 points is paid and at least 7 out of the 10 questions are answered correctly. Success in a period determines the payoff of Member B in the respective group. At the end of a round, an observer estimates, based on the successes of this Member A in that round, how many questions this Member A answered correctly at the beginning of the experiment. This estimate determines the round payoff of Member A.

Member A

At the beginning of the experiment, each member answers 10 questions about the capital cities of countries (see Figure 1). The countries asked about are randomly selected for each member, with each of the 194 possible countries being equally likely. The number of correctly answered questions determines the type of Member A in this experiment (a number between 0 and 10). Member A keeps this type for the entire experiment. Member B in the respective group does not learn the type of Member A.

Member A starts each round with a budget of 60 points.

Figure 1: Example image of a randomly selected quiz question. Each Member A answers 10 such randomly selected questions in every period and once at the beginning of the experiment..

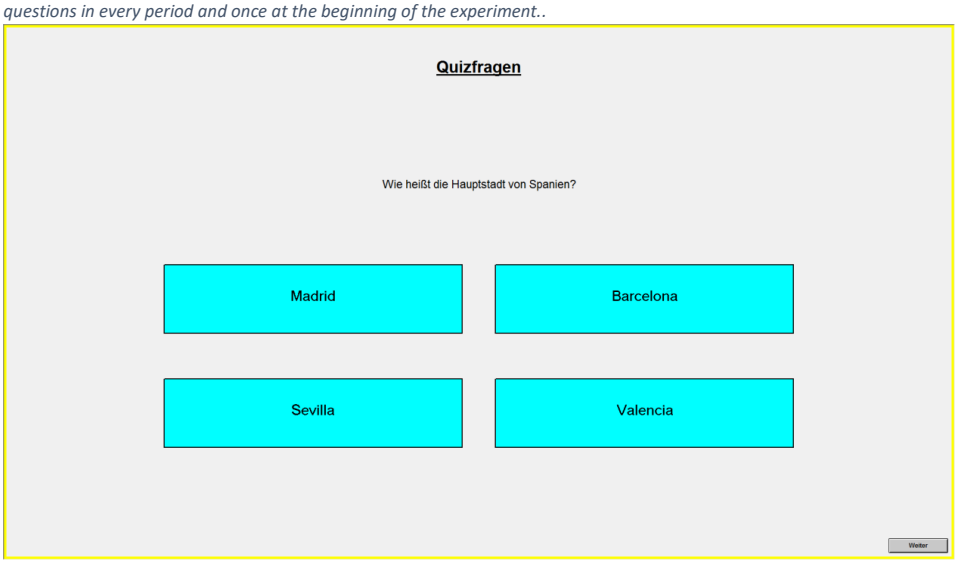

In each round, Member A must make two decisions (see Figures 2 and 3):

1. Whether to pay a cost of 30 points in the first period in order to have a chance of success in the first period.
2. Whether to pay a cost of 30 points in the second period in order to have a chance of success in the second period.

After each choice of whether to pay the 30 points in order to have a chance of success, Member A again answers 10 questions about the capital cities of randomly selected countries. Success means answering at least 7 of the 10 questions correctly. If Member A decides not to pay the cost of 30 points in a period, success in that period cannot occur, regardless of how many questions are answered correctly. If Member A decides to pay the cost of 30 points in a period, the member has success if at least 7 of the 10 following questions are answered correctly. If fewer than 7 questions are answered correctly, there is no success. No other participant learns whether Member A decided to pay the cost.

Success in a period increases the round payoff of Member B in the same group. An observer (a Member B from another group) provides, depending on the success outcomes in a round, an estimate of how many questions this Member A answered correctly at the beginning of the experiment. This estimate determines the round payoff of Member A: the higher the estimate given by Member B, the higher the round payoff of Member A.

Figure 2: Example image of Member A's decision whether to pay 30 points in the first period in order to have a chance of success in the first period. In this example, Member B of the same group has decided that Member A receives feedback about the success of Period 1.

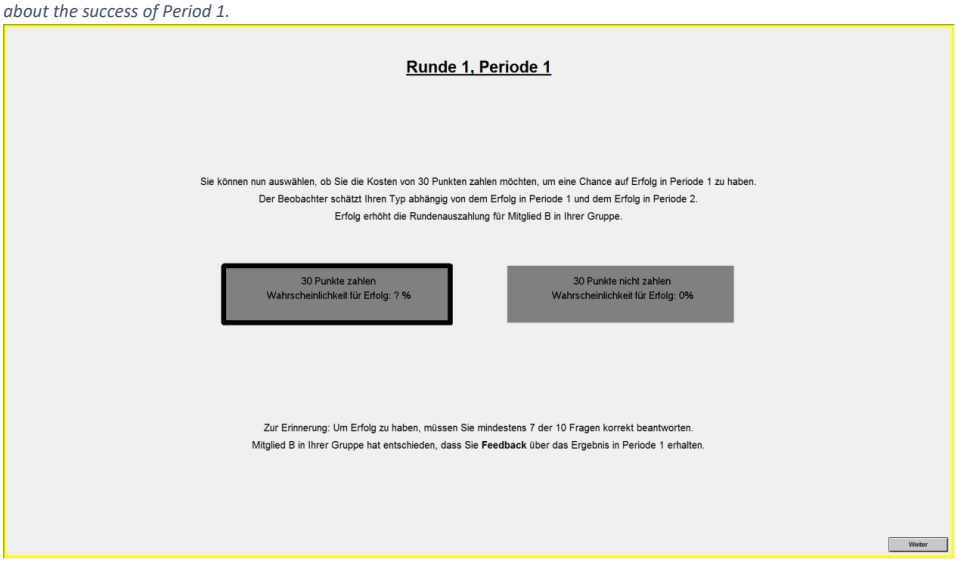

Figure 3: Example image of Member A's decision whether to pay 30 points in the second period in order to have a chance of success. In this example, Member B of the same group has decided that Member A receives feedback about the success of Period 1. In this example, Member A was successful in the first period.

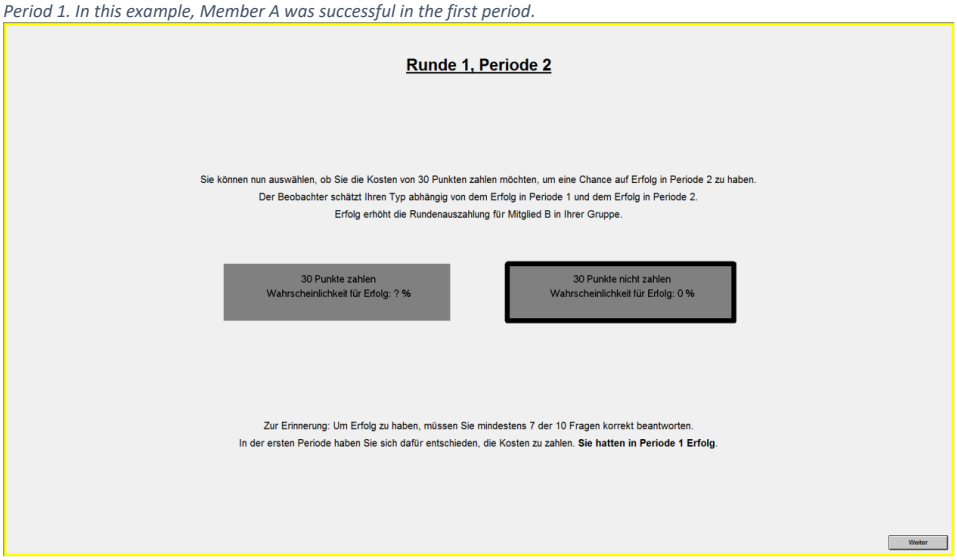

The payoff of Member A in a round is therefore

- $60\ (starting\ endowment) + 10 * estimate\ of\ the\ observer$, if Member A decided not to pay the 30 points for chances of success in either period.
- $60\ (starting\ endowment) - 30\ (payment\ for\ chance\ of\ success\ in\ one\ period) + 10 * estimate\ of\ the\ observer$, if Member A decided to pay the 30 points for chances of success in one of the two periods.
- $60\ (starting\ endowment) - 60\ (payment\ for\ chance\ of\ success\ in\ one\ period) + 10 * estimate\ of\ the\ observer$, if Member A decided to pay the 30 points for chances of success in both periods.

Example: A Member A decides not to pay the 30 points for a chance of success in the first period. In the second period, Member A decides to pay the 30 points for a chance of success. The observer estimates the type of Member A to be 6 if no success occurred in the first period and success occurred in the second period (the observer therefore estimates that this Member A answered 6 out of 10 questions correctly at the beginning of the experiment). Furthermore, the observer estimates the type of Member A to be 3 if no success occurred in either of the two periods of the current round (the observer therefore estimates that this Member A answered 3 out of 10 questions correctly at the beginning of the experiment). If Member A answers at least 7 of the 10 questions correctly in the second period, the payoff is: 60 - 30 + 60 = 90 points, and if fewer than 7 questions are answered correctly, the payoff of Member A is 60 - 30 + 30 = 60 points.

Figure 4: Example image of the round summary for Member A.

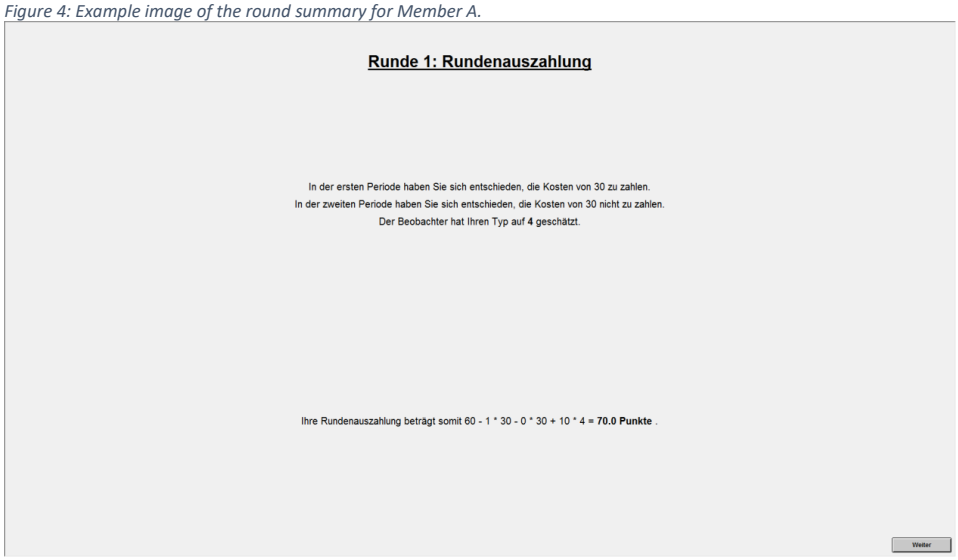

Member B / Observer of another group

Member B also answers 10 randomly selected quiz questions at the beginning of the experiment in order to get an impression of the difficulty of the questions. Afterwards, Member B learns how many of the 10 questions were answered correctly.

At the beginning of each round, Member B decides whether Member A of the same group receives feedback about success in the first period or not (see Figure 5). This determines whether Member A of the same group, after deciding whether to pay 30 points in the first period but before making the decision for Period 2, learns whether success occurred in Period 1 or not.

At the end of each round, Member B of the group receives 50 points for each success of Member A of the same group. If Member A therefore has success in both periods of a round, the round payoff of Member B is 100 points. If Member A has success in only one period, the round payoff of Member B is 50 points. If Member A has no success in both periods, the round payoff of Member B is 0 points. Member B never learns whether Member A of the same group paid 30 points in a period in order to have a chance of success or not.

Figure 5: Decision of Member B whether Member A of the same group is informed about the success in Period 1 in this round or not.

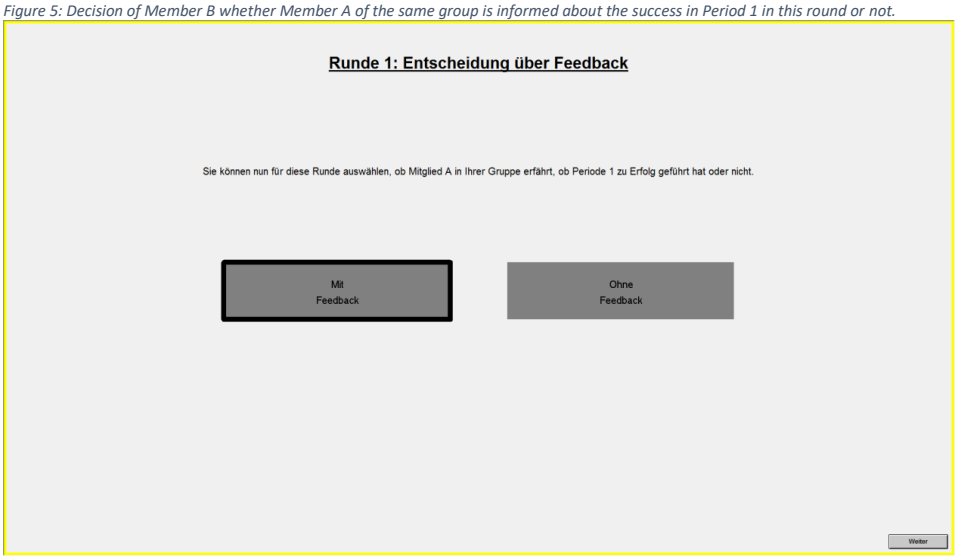

As an **observer**, a Member B observes the Member A of another group. Here, the observer provides an estimate of the type of Member A for each possible success outcome (see Figure 6).

The payoff for the observer is as follows:

$$Round\ payoff\ =\ 100\ -10 * |submitted\ estimate - actual\ type\ of\ Member\ A|$$

The closer the observer's estimate is to the true type of Member A, the higher the observer's round payoff. The estimate submitted for the case that actually occurred is relevant for this. The estimate given for the case that actually occurred determines the round payoff of the observed Member A.

Only at the end of the entire experiment does an observer learn their round payoff, and before that does not learn which case for which they submitted an estimate was actually relevant in a round.

Figure 6: Example image of the observer's estimate of the type of Member A from another group. Only at the end of the experiment does the observer learn which of the four cases actually occurred and is therefore payoff-relevant for them..

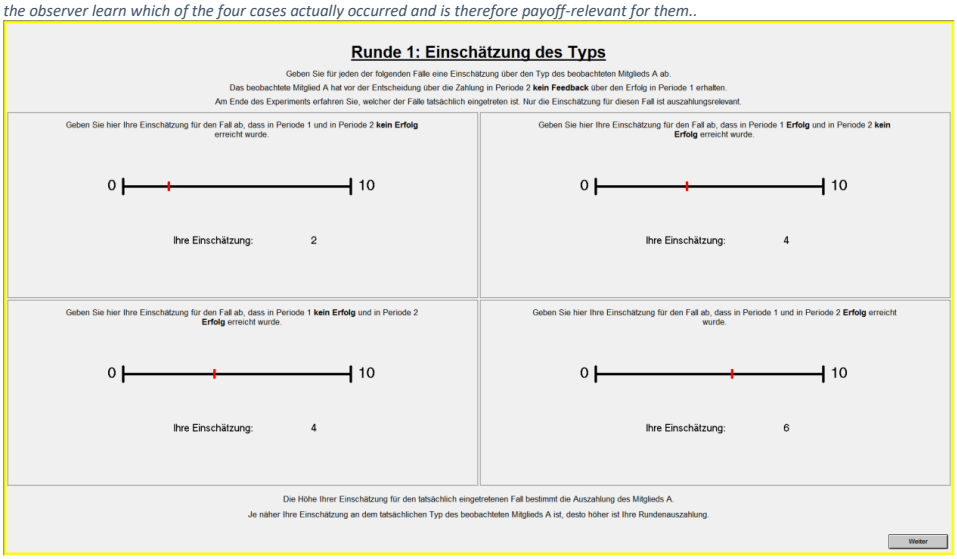

Figure 7: Example image of the round summary for Member B.

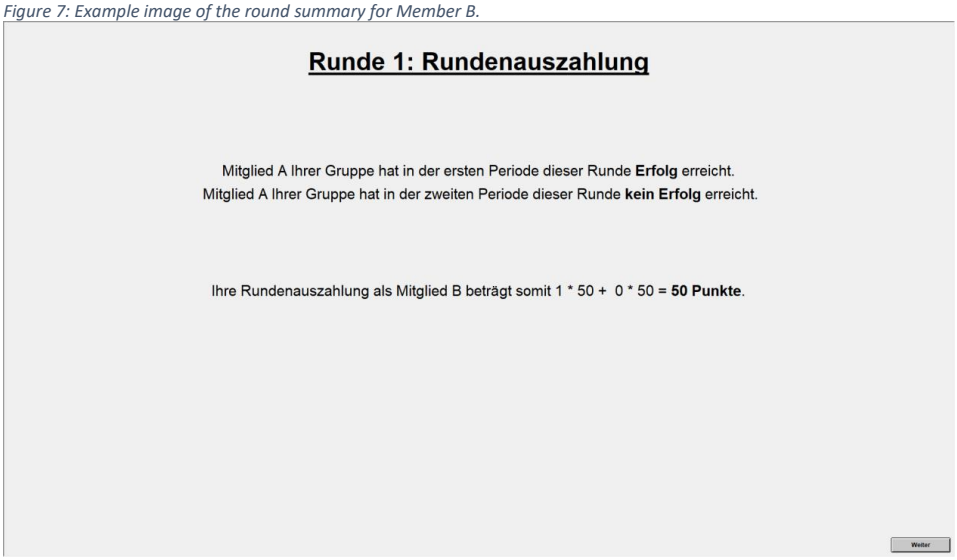

Feedback of a round

Member B learns at the end of each round how often success was achieved by Member A (0, 1, or 2 times). Member B does not learn the type of Member A, and also not whether Member A paid the cost for the chance of success.

An observer does not learn at the end of the round which of the four cases actually occurred. Only at the end of the experiment does the observer learn, for the two randomly selected rounds, the true type of the observed Member A. The observer does not learn whether this Member A paid the cost for the chance of success.

Member A and Member B learn, after the 10 questions at the beginning of the experiment, how many questions were answered correctly. In each round, Member A learns whether success was achieved in the second period. In addition, Member A learns the round payoff, that is, the observer's estimate of the type of Member A. Member B of the same group decides whether Member A learns in a round whether success was achieved in the first period. If Member B of a group chooses a round "with feedback", Member A learns after the first period and before the second period of a round whether success was achieved in the first period. If Member B chooses a round "without feedback", Member A never learns whether Period 1 of that round resulted in success.

Total Payment

The payment at the end of the experiment for Member A consists of:

- Round payoffs from 3 randomly selected rounds.
- 6 euros for arriving on time.

The payment at the end of the experiment for Member B / Observer consists of:

- Round payoffs for Member B from 2 randomly selected rounds.
- Round payoffs for the observer from 2 randomly selected rounds.
- 6 euros for arriving on time.

Questions?

Please take your time to go through the instructions once again. If you have a question, please raise your hand. An experimenter will then come to your seat and answer the question personally. If you believe that you have understood everything, please click "Start" on the screen. Next, control questions will appear on the screen. The control questions are intended to ensure that every participant has understood the instructions. The answers do not affect the payment. After answering the control questions, you can click "continue" to see the solutions and compare your answers with the correct ones. Once all participants have answered the control questions and all uncertainties have been clarified, the experiment will begin.

SM 2.2 Abstract Treatment

Overview

Welcome to this experiment. We ask you not to speak with other participants during the experiment and to switch off your mobile phones and other mobile electronic devices.

For participating in today's session, the amount you earn will be paid out to you in cash at the end of the experiment. The size of the payment depends partly on your decisions, partly on the decisions of other participants, and partly on chance. It is therefore important that you carefully read and understand the instructions before the experiment begins.

In this experiment, every interaction between participants takes place via the computers in front of which you are seated. You will interact with each other anonymously and your decisions will only be stored together with your random ID number. Neither your name nor the names of other participants will be revealed, neither today nor in future written analyses.

Today's session consists of several rounds. At the end, a total of 2 or 3 rounds will be randomly selected and paid out. The rounds that are not selected will not be paid. Your payment amount results from the points earned in the selected rounds, converted into euros, as well as your show-up fee of 4 euros. The conversion of points into euros is as follows: Each point is worth 4 cents, so that:

$$25 \text{ Points} = 1 \text{ Euro}$$

Each participant will be paid privately so that other participants cannot see how much you have earned.

Experiment

The experiment consists of two parts with a total of 10 rounds. The rounds within Parts 1 (Rounds 1 – 5) and 2 (6 – 10) are identical in structure. The instructions for the second part will be provided after the completion of the first part.

At the beginning of the experiment, you will be randomly assigned either the role of **Member A** or **Member B / Observer**. Each participant keeps this role for the entire experiment.

The Group

In the first round, you will be assigned to a large group of 8, 10, or 12 members. At the beginning of each round, two participants from your respective large group will be randomly selected to form a group in that round. Each group consists of one Member A and one Member B. The members of a group are therefore randomly reassigned in every round.

General Structure of a Round

Each round consists of two periods. In each period, Member A can decide whether to pay a cost of 30 points in order to have a chance of success in that period. The probability of success, if the cost of 30 points is paid, is determined by the type of Member A (between 0 and 100%). Success in a period determines the payoff of Member B in the respective group. At the end of a round, an observer estimates the type of Member A based on the successes of this Member A in that round. This estimate determines the round payoff of Member A.

Member A

Member A starts each round with a budget of 60 points. At the beginning of each round, each Member A is randomly assigned a type. This type is a number between 0 and 100, with each number being equally likely. This type remains fixed for the entire round. Member B of the same group does not learn the randomly assigned type of Member A. In each round, Member A must make two decisions (see Figures 1 and 2):

1. Whether to pay a cost of 30 points in the first period in order to have a chance of success in the first period.
2. Whether to pay a cost of 30 points in the second period in order to have a chance of success in the second period.

If Member A decides not to pay the cost of 30 points in a period, success in that period cannot occur. If Member A decides to pay the cost of 30 points in a period, the chances of success in that period are determined by the type: a Member A of type 60, for example, has a probability of success of 60% in every period in which they pay the 30 points. No other participant learns whether Member A decided to pay the cost.

Success in a period increases the round payoff of Member B in the same group. An observer (a Member B from another group) provides an estimate of the type of Member A, depending on the success outcomes in a round. This estimate determines the round payoff of Member A: the higher the estimate of the type given by Member B, the higher the round payoff of Member A.

Figure 1: Example image of Member A's decision whether to pay 30 points in the first period in order to have a chance of success in the first period. In this example, the randomly assigned type is 41, and it was randomly determined that Member A does not receive feedback about the success of Period 1.

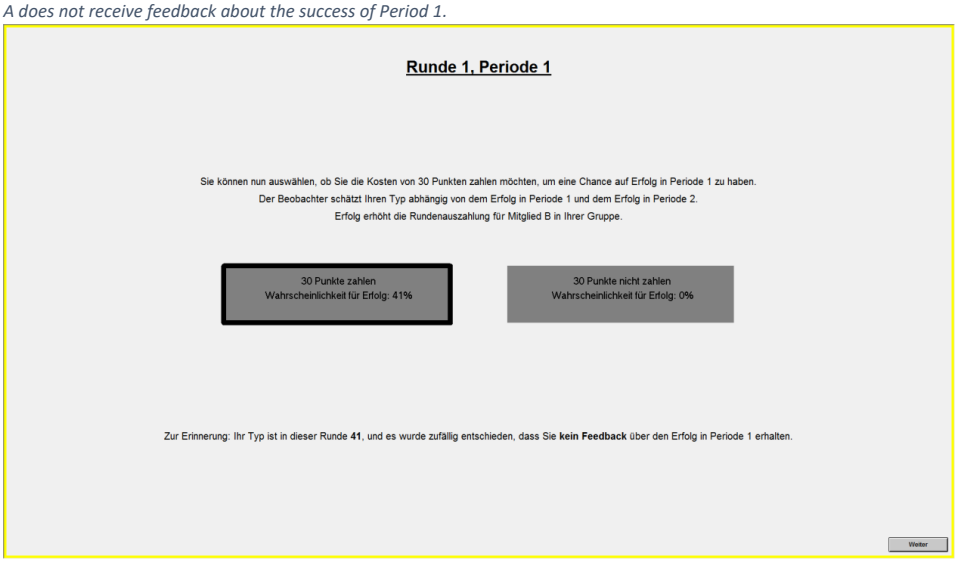

Figure 2: Example image of Member A's decision whether to pay 30 points in the second period in order to have a chance of success. In this example, the randomly assigned type is still 41, and it was randomly determined that Member A does not receive feedback about the success of Period 1.

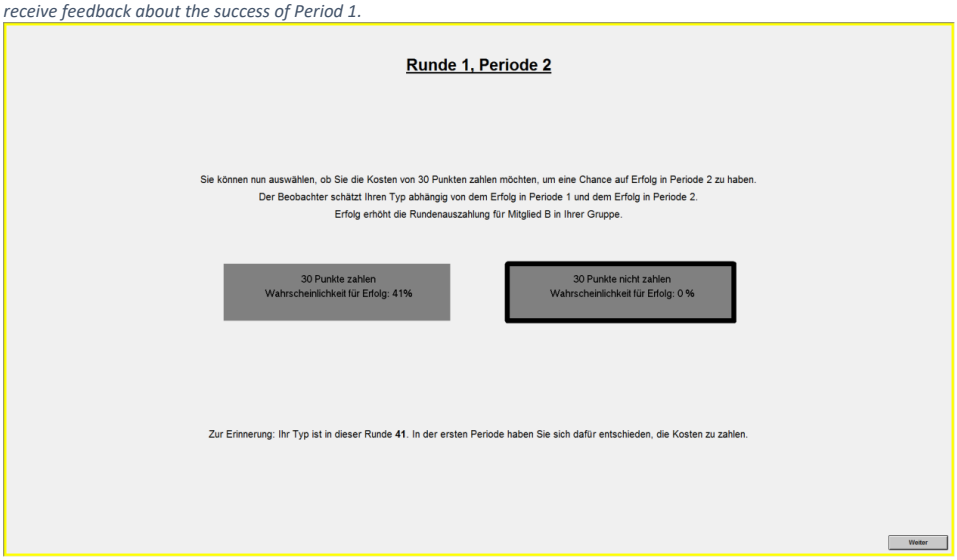

The payoff of Member A in a round is therefore:

- $60\ (starting\ endowment) +\ estimate\ of\ the\ observer$, if Member A decided not to pay the 30 points for chances of success in either period.
- $60\ (starting\ endowment) - 30\ (payment\ for\ chance\ of\ success\ in\ one\ period) +\ estimate\ of\ the\ observer$, if Member A decided to pay the 30 points for chances of success in one of the two periods.
- $60\ (starting\ endowment) - 60\ (payment\ for\ chance\ of\ success\ in\ both\ periods) +\ estimate\ of\ the\ observer$, if Member A decided to pay the 30 points for chances of success in both periods.

Example: A Member A of type 20 decides not to pay the 30 points for a chance of success in the first period. In the second period, Member A decides to pay the 30 points for a chance of success. The observer estimates the type of Member A to be 65 if no success occurred in the first period and success occurred in the second period. In addition, the observer estimates the type of Member A to be 35 if no success occurred in either of the two periods. With 20% probability (the type of Member A), Member A has success in the second period and the payoff is therefore: 60 - 30 + 65 = 95 points. With 80% probability, Member A has no success in the second period and the payoff of Member A is 60 - 30 + 35 = 65 points.

Figure 3: Example image of the round summary for Member A.

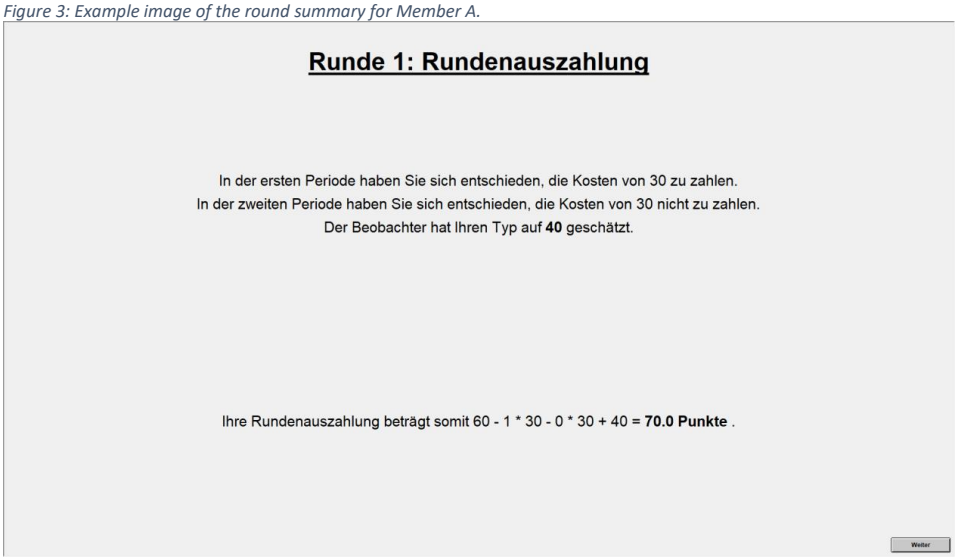

Member B / Observer of another group

At the end of each round, Member B of the group receives 50 points for each success of Member A of the same group. If Member A therefore has success in both periods of a round, the round payoff of Member B is 100 points. If Member A has success in only one period, the round payoff of Member B is 50 points. If Member A has no success in both periods, the round payoff of Member B is 0 points. Member B never learns whether Member A of the same group paid 30 points in a period in order to have a chance of success or not.

As an **observer**, a Member B observes the Member A of another group. Here, the observer provides an estimate of the type of Member A for each possible success outcome (see Figure 4).

The payoff for the observer is as follows:

$$Round\ payoff\ =\ 100\ -|submitted\ estimate - actual\ type\ of\ Member\ A|$$

The closer the observer's estimate is to the true type of Member A, the higher the observer's round payoff. The estimate submitted for the case that actually occurred is relevant for this. The estimate given for the case that actually occurred determines the round payoff of the observed Member A.

Only at the end of the entire experiment does an observer learn their round payoff, and before that does not learn which case for which they submitted an estimate was actually relevant in a round.

Figure 4: Example image of the observer's estimate of the type of Member A from another group. Only at the end of the experiment does the observer learn which of the four cases actually occurred and is therefore payoff-relevant for them.

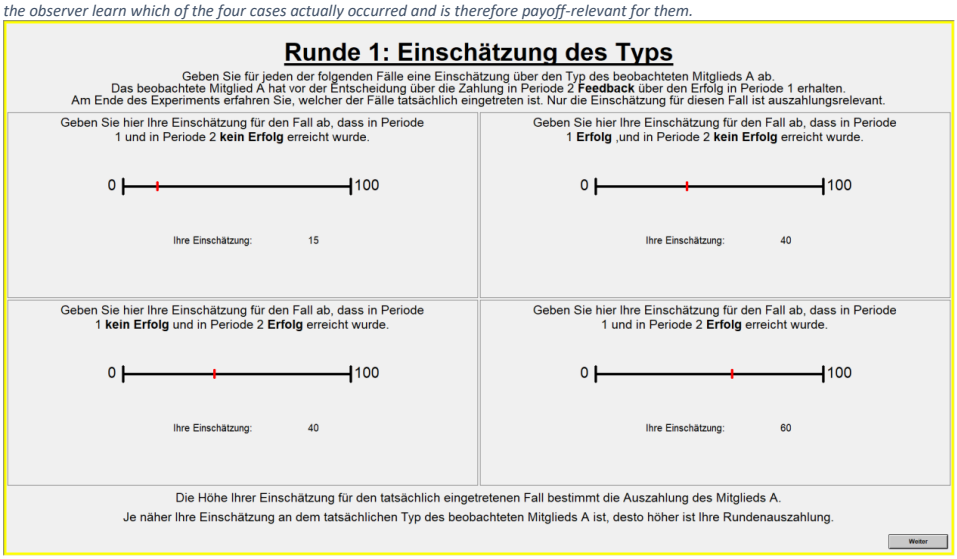

Figure 5: Example image of the round summary for Member B.

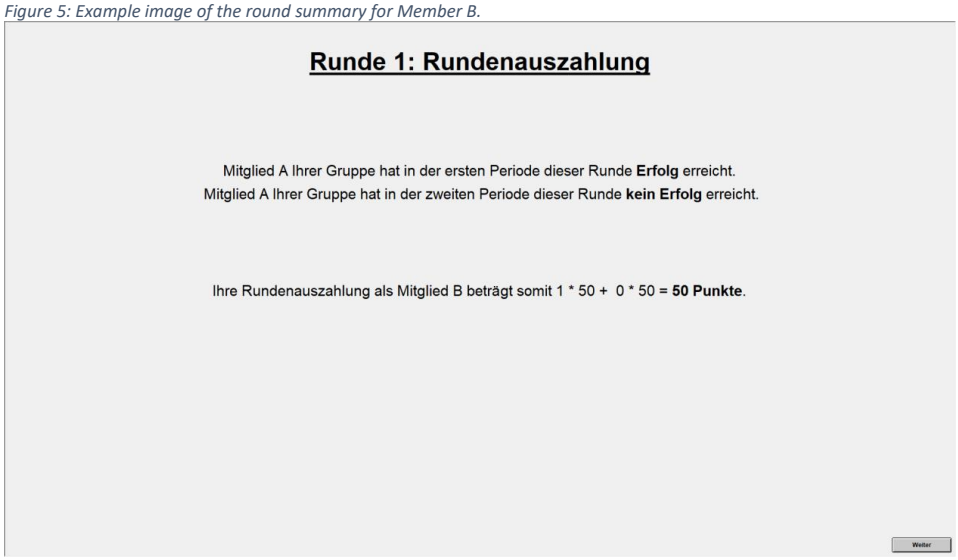

Feedback of a round

Member B learns at the end of each round how often success was achieved by Member A (0, 1, or 2 times). Member B does not learn the type of Member A, and also not whether Member A paid the cost for the chance of success.

An observer does not learn at the end of the round which of the four cases actually occurred. Only at the end of the experiment does the observer learn, for the two randomly selected rounds, the true type of the observed Member A. The observer does not learn whether this Member A paid the cost for the chance of success.

In each round, Member A learns their own type and whether they achieved success in the second period. In addition, Member A learns the round payoff, that is, the observer's estimate of the type of Member A. At the beginning of a round, it is randomly determined whether Member A learns in that round whether success was achieved in the first period. In a round "with feedback", Member A learns after the first period and before the second period of a round whether success was achieved in the first period. In a round "without feedback", Member A never learns whether Period 1 of that round resulted in success.

Total Payment

The payment at the end of the experiment for Member A consists of:

- Round payoffs from 3 randomly selected rounds.
- 4 euros for arriving on time.

The payment at the end of the experiment for Member B / Observer consists of:

- Round payoffs for Member B from 2 randomly selected rounds.
- Round payoffs for the observer from 2 randomly selected rounds.
- 4 euros for arriving on time.

Questions?

Please take your time to go through the instructions once again. If you have a question, please raise your hand. An experimenter will then come to your seat and answer the question personally. If you believe that you have understood everything, please click “Start” on the screen. Next, control questions will appear on the screen. The control questions are intended to ensure that every participant has understood the instructions. The answers do not affect the payment. After answering the control questions, you can click “continue” to see the solutions and compare your answers with the correct ones. If all participants have answered the control questions and all uncertainties have been clarified, the experiment will begin.