

The Silent Treatment*

Geoffroy de Clippel
Brown University

Kfir Eliaz
Michigan and TAU

Kareen Rozen[†]
Brown University

August 2016

Abstract

A principal's attention is repeatedly sought by multiple agents, each eager for his idea to be implemented. An idea's quality stochastically affects the principal's profit, and agents' abilities to generate good ideas may be private information. The principal is unable to review proposals before choosing one each period. She can provide incentives only through her selection rule among proposals, but cannot commit to this rule in advance. We show how she may discipline agents to exercise restraint, achieving her first-best in an intuitive belief-free equilibrium. Whether first best is achievable hinges on the worst possible agent, the organization's 'weakest link.'

Keywords: limited attention, organizations, belief-free equilibrium, mechanism design without commitment, multi-armed bandit

*This paper benefitted from seminar audiences at Brown University and the University of Wisconsin-Madison. Special thanks to Nageeb Ali, Heski Bar-Isaac, Johannes Hörner, Phil Reny, Larry Samuelson and Lones Smith for their helpful suggestions.

[†]Corresponding author. Address: 64 Waterman Street, Department of Economics, Brown University, Providence, RI 02912. Email: kareen.rozen@brown.edu

1 Introduction

“Blessed is the man who, having nothing to say, abstains from giving us wordy evidence of the fact.” – George Eliot

Most people repeatedly find themselves engaged in trying to meet the demands for their own attention by others, or demanding others’ attention themselves. We send emails and texts expecting a prompt response. At the same time, we are flooded with incoming communications demanding our immediate attention. Meetings drag on, with too many participants inclined to make themselves heard. Open source platforms become cluttered with insignificant updates posted by contributors vying for credit.

As illustrated by these examples, communicating without restraint imposes an externality on decision makers: good ideas risk getting lost in the clamor for attention. This externality arises because it is costly, or even impossible, for a decision maker to consider each and every idea. Consequently, a decision maker is forced to choose which ideas to consider among the many proposed, and this may be challenging if little information is known a priori about each idea. However, since the set of available ideas, or the set of items that demand the decision maker’s attention, is generated by *strategic* agents, a decision maker may be able to influence the quality of this set by incentivizing agents to be *discerning*; that is, to communicate information only when it is important. How should agents then be incentivized to help the decision maker make optimal choices?

We address this question within a simple and stylized dynamic setting. In every period, a principal has a problem to solve and seeks proposals from multiple agents, who may be her subordinates, consultants or independent contributors. Each agent comes up with a new idea in each period, which is of good quality with probability θ and of bad quality otherwise. The principal and agents have a conflict of interest. While implementing any idea is better for the principal than implementing none at all, she prefers to implement the highest quality idea available. An agent, on the other hand, benefits whenever his idea is the one selected. Instead of taking the influx of proposals in each period as exogenous, we model those proposals as originating from *strategic* agents who can choose whether to propose an idea at all. An agent knows the quality of his own idea when deciding whether to propose it. However, quality is only a noisy indicator of the profit the principal would get if she implements the idea. Profit may be high or low, with good ideas yielding high profit with probability γ and bad ideas yielding high profit with a smaller probability β . The principal is unable to spend the time to evaluate whether the proposed ideas are good or bad before choosing which one of them to implement that period. An idea’s profit is realized only once it is implemented; the principal cannot know what her profit would have been from unchosen

ideas. The principal seeks to maximize her discounted sum of profits. The only tool at the principal's disposal for providing punishments or rewards to agents is the procedure by which she selects among proposals in each period. The principal cannot commit to a selection rule in advance; it must be rational for her to follow her selection rule in equilibrium.

We say that the principal achieves her first best when there is a strategy profile and threshold patience level such that (i) the strategy profile is a perfect Bayesian Nash equilibrium (PBNE) if agents' discount factors exceed the threshold; and (ii) the strategy profile leads to the selection of the highest quality idea in every period. Our first main result establishes the existence of a unique threshold probability $\theta^* \leq 1/2$ that characterizes when the principal can achieve her first best. As it turns out, her ability to do so hinges on the talent of agents in her organization. If the probability that an agent has a good idea is below this threshold, then the principal's first best *cannot* be achieved. If, however, the probability of a good quality idea is above θ^* , then the principal can achieve her first best if profits are sufficiently informative of an idea's quality.¹ In this case, the principal's first best is achieved through a simple and intuitive strategy profile that we call the *Silent Treatment*.

The Silent Treatment strategy profile is defined as follows. In any period, one agent is designated as the *agent of last resort*, and all other agents are designated as *discerning*. The agent of last resort proposes his idea regardless of its quality. Each discerning agent proposes his idea if it is good, and remains silent otherwise. The principal selects the idea proposed by the agent of last resort if it is the only one available. Otherwise, the principal ignores the proposal of the last resort agent and selects among the discerning agents' proposals by randomizing uniformly. The initial agent of last resort is chosen arbitrarily, and remains in that role until the principal realizes a low profit for the first time. Going forward, the agent of last resort is the most recent agent whose idea yielded low profit for the principal.

The Silent Treatment strategy profile has a number of desirable properties. First, it requires players to keep track of very little information: they need only know who was the last agent whose idea yielded low profit. Second, it does not require the agents to punish the principal (the mechanism designer) to ensure that she follows the strategy. This can be seen from the fact that whenever the silent treatment strategy profile is a PBNE, then it remains a PBNE even when the principal's discount factor is zero. Third, it is independent of the probability of a good idea (θ), and is robust to having privately observed heterogeneity in the ability of agents to generate good ideas. Consequently, the principal need not engage in complicated inferences about abilities, and the agents need not be concerned with signalling when deciding whether to propose an idea.

To demonstrate the last point, we enrich our benchmark model by assuming that each

¹As captured by the likelihood ratio $\frac{1-\beta}{1-\gamma}$ that low profit arises from a bad idea versus a good idea

agent i is characterized by an ability $\theta_i \in [\underline{\theta}, \bar{\theta}]$, which is the probability that he has a good idea in a period. Agents' abilities are not observed by the principal and may or may not be observed by other agents. Our second main result establishes that if $\underline{\theta} > \theta^*$, then so long as profits are sufficiently informative of an idea's quality, the silent treatment strategy profile attains the principal's first-best in an *ex-post* PBNE for *any* realized vector of abilities. That is, the silent treatment profile constitutes a *belief-free* equilibrium.² If, however, $\underline{\theta} < \theta^*$, then *no* belief-free equilibrium can attain the principal's first best. Thus the organization's worst possible agent, its 'weakest link,' determines what is achievable. Establishing our result is complicated by the heterogeneity in agents' abilities: an agent's continuation payoff depends on the ability of the last resort agent, which varies as different agents are relegated to this role. Our method of proof overcomes this challenge by taking advantage of the special properties of the matrices emerging from the agents' value functions and incentive conditions. Belief-freeness is achieved because there is sufficient slack in the agents' incentive conditions (not because of indifference). Discerning agents simply prefer to propose only good ideas in equilibrium, and this is all the information the principal needs to know to select a proposal.

Our analysis also applies in the limiting case $\gamma = 1$, in which a good idea generates high profit with probability one.³ Indeed, the silent treatment strategy profile remains a belief-free equilibrium if $\underline{\theta} > \theta^*$ and the agents are sufficiently patient. In this limit case, the principal will be able to infer that an agent proposed a bad idea upon receiving a low profit. As a result, the principal will be able to attain her first best in equilibrium for a wider range of parameters by resorting to harsher punishments than those in the silent treatment strategy profile. Such punishments, however, would affect the size and/or quality of the principal's pool of proposals, and therefore reduce her profits, should cheating actually occur. To put the matter in context, observe that the literature on mechanism design generally interprets the principal as having the ability to induce participants to coordinate on her most preferred equilibrium. In mechanism design without commitment, we suggest that the principal may be interpreted analogously, as a special player who is able to induce the participants to coordinate on her most preferred equilibrium *at any history*, whether on or off the equilibrium

²In the context of repeated games with imperfect monitoring, see, for instance, Piccione (2002) and Ely and Välimäki (2002). They look to belief-freeness as a way to generate robustness to the particular structure of private monitoring in a prisoner's dilemma, and carefully construct mixed strategies where each player is indifferent between all his actions, no matter his opponent's private history. Hörner, Lovo and Tomala (2011) consider the notion of belief-free equilibrium in games of incomplete information. They examine the existence of belief-free equilibrium payoffs under certain types of information structures and reward functions, and also use sophisticated strategies involving randomization to delimit what is achievable. In contrast to these works, in our simple model, the principal's first best is attained with agents playing pure strategies; and we analytically characterize the threshold ability level below which the first best is impossible to attain in a belief-free way.

³This is true regardless of the probability β that a bad idea generates high profit, so long as $0 \leq \beta < 1$.

path. With this new notion of credibility in mind, our earlier result can be strengthened: a belief-free PBNE achieves the principal’s first best in the stage game played after each history *if and only if* her first best is achievable via the Silent Treatment strategy profile.

The paper is organized as follows. Section 2 discusses connections to the literature. Section 3 presents our benchmark model, in which agents all have the same commonly known ability. Section 4 introduces the silent treatment strategy profile and characterizes when the principal can achieve her first best. Section 5 studies the case of heterogenous and privately known agents’ abilities. Section 6 considers the limiting case in which good ideas surely yield high profit. Section 7 concludes.

2 Related Literature

Our paper relates to several strands of literature. When the qualities of agents’ ideas are privately drawn, the principal’s problem is reminiscent of a multi-armed bandit problem (Gittins and Jones, 1974), with the twist that the bandits’ arms respond to incentives and strategically decide whether to make themselves available in each period.⁴ In the classic multi-armed bandit problem, a decision maker faces multiple arms, each of which has an unknown probability of success. The decision maker wants to maximize her discounted sum of payoffs from pulling the arms, but faces a tradeoff between learning which arms are best and exploiting those that have succeeded so far. The solution to this classic problem, based on the Gittins index, does not achieve the decision maker’s first best, and uses a sophisticated learning strategy requiring commitment and patience. By contrast, in our setting, the principal achieves her first best by properly incentivizing the arms (the agents) without having to infer quality levels, without having to commit to a strategy, and for any level of her patience.⁵

The problem we study may be thought of as dynamic mechanism design without transfers when the planner is a player (and therefore, cannot commit). In our model, there is no institutional device that enables the principal to credibly commit to a policy, and the agents’ payoffs cannot be made contingent on the payoff to the principal. This could be due to the fact that the principal’s payoff cannot be verified by an outside party (e.g., it may include intangible elements such as perceived reputation), or because of institutional constraints that

⁴Bergemann and Välimäki (2008) offer a nice survey of applications of multi-armed bandit problems to economics. For the case of a one-armed bandit, Bar-Isaac (2003) endogenizes the arm’s availability by allowing a monopolist who sells a good of unknown quality to choose if to sell on the market each period.

⁵Some recent work in computer science considers a different generalization allowing for strategic bandits, whereby the bandits make a one-time decision of whether to be available and with what probability of success, in response to the algorithm determining how arms will get pulled in the future. Algorithms are then compared based on the criterion of minimal regret. See Ghosh and Hummel (2012).

preclude such contracts, as in most public organizations where subordinates may suggest ideas and improvements to an executive decision-maker. A recent paper in this literature on dynamic mechanism design with neither transfers nor commitment is Li, Matouschek and Powell (2016) who study the dynamics of power in an infinitely repeated game between a principal and one agent. In each period, the principal decides whether to entrust the choice of a project to the agent, who privately observes which projects are available and whose preferred project differs from that of the principal. The authors characterize the public perfect equilibrium that maximizes the principal's expected payoff, and show that eventually, the principal will not be able to make effective use of the agent's information: either the principal will end up relying on the agent to make all choices, or ignore the agent's recommendations entirely. Which extreme occurs depends on random outcomes early in the game. An opposite dynamic is shown by Lipnowski and Ramos (2015), who analyze an infinitely repeated game between a principal and one agent. In their model, the principal decides each period whether to delegate the choice of a project, the quality of which is observed only by the agent. While the agent cares only about the project's quality, the principal also cares whether the quality exceeds a fixed cost. They show delegation occurs often at the start of the relationship, but that eventually, it will rarely occur. Among other modeling differences with these papers, we consider a multi-agent setting. The competition between agents in our model is a driving factor in the results: if there were only one agent, the principal could achieve no better than having him propose all his ideas, both good and bad. The principal's best equilibrium in our model achieves her unconstrained first best, and does not exhibit nonstationary dynamics (indeed, the equilibrium is Markovian with respect to the identity of the agent of last resort).

Our paper contributes to the emerging literature on allocation dynamics in repeated games. Two recent papers, Board (2011) and Andrews and Barron (2016), study how a principal (firm) chooses each period among multiple agents (contractors or suppliers) whose characteristics are perfectly observed by the principal, but whose post-selection action is subject to moral hazard. Both papers consider relational contracts and thus allow for players to make transfers in the repeated game. Board (2011) considers a hold-up problem, where the chosen contractor each period decides how much to repay the principal for her investment. Assuming that the principal can commit to the selection rule, Board shows that it is optimal to be loyal to a subset of 'insider' contractors, because the rents the principal must promise to entice the contractor to repay act as an endogenous switching cost. He shows that this bias towards loyalty extends when the principal cannot commit, so long as she is sufficiently patient. Relaxing the assumption of commitment and introducing imperfect monitoring in the moral hazard problem, Andrews and Barron (2016) consider a firm who repeatedly faces

multiple possible, ex-ante symmetric suppliers. A supplier’s productivity level is redrawn each period but is observable to the principal. The principal approaches a supplier and, upon agreeing to the relationship, the supplier makes a hidden, binary effort choice yielding a stochastic profit for the principal. Under the assumption of private monitoring (that each agent observes only his own history with the principal), they show that the principal’s first best can be achieved for the widest possible range of discount factors by a ‘favored supplier’ allocation rule. Each period, the principal must choose a supplier from among those with the highest observed productivity level, but breaks ties in favor of the agent who most recently yielded high profit. There are several interesting differences with our paper. First, we study a problem without transfers. Furthermore, we study a problem of adverse selection: the principal’s problem is precisely that she cannot observe the distinguishing characteristic – the quality – of the agents’ ideas. In our model, an aim of the principal’s selection rule is to influence her set of proposers; thus the set of possible agents in each period is endogenous to the problem. Another interesting difference with Andrews and Barron (2016) is that our results rely on the history being at least partially public: the identity of the current agent of last resort must be known to all players. By contrast, Andrews and Barron point out that if they were to relax private monitoring, then the agents could collectively punish the principal and the optimal allocation rule would become stationary (independent of past performance). As discussed earlier, the Silent Treatment strategy profile does not rely on punishing the principal. Whenever it is an equilibrium, it remains so for any discount factor of the principal, even if she is fully myopic.

Finally, our paper joins a growing literature on attention and other cognitive constraints in organizations (see Garicano and Prat (2013) for a thorough survey of earlier works and a detailed history of the field; see Halac and Prat (2015) and Dessein, Galeotti and Santos (2015) for some more recent work on the subject). Our work brings a new aspect to this literature, pointing out that because interactions in organizations are often repeated, it may be possible to compensate for a lack of attention by using intertemporal incentives to minimize superfluous communication. In our model, the principal is able to filter agents’ ideas without paying attention to their quality, without considering what agents’ abilities may be, and while remembering very little information from the past.

3 Benchmark Model

There is one principal and a set $\mathcal{A} = \{1, \dots, n\}$ of $n \geq 2$ agents who individually and independently come up with a new *idea* for the principal at each period $t = 1, 2, 3, \dots$. An idea’s profit to the principal is either high (H) or low (L), where $H > L \geq 0$. Her profit

depends stochastically on the quality of the idea that she implements. An idea's quality is either *good* or *bad*. A good idea has probability $\gamma \in (0, 1)$ of generating high profit for the principal; while a bad idea generates high profit with a strictly smaller probability β .⁶ There is a commonly known probability θ that an agent's idea in any given period is good.

In every period, the stage game unfolds as follows. Knowing the quality of his idea, each agent decides whether to propose it to the principal. The principal then decides which idea, if any, to implement among these proposals. Figuring out an idea's quality prior to its implementation requires the principal's attention. This is costly, which can be modeled, for instance, via an explicit cost of reviewing ideas, or by introducing a capacity constraint which induces an implicit cost. The next section shows how the principal may take advantage of the repeated nature of her interactions with the agents to reach her first best, even when her attention is fully constrained and she cannot review any ideas at all. There is thus no need to explicitly model costs to make this point.

Agent i gets a positive payoff u_i in period t if the principal picks his idea at t . Agent i 's objective is to maximize the expectation of the discounted sum $\sum_{t=0}^{\infty} \delta_i^t u_i 1\{x_t = i\}$, where δ_i is agent i 's discount factor, $1\{\cdot\}$ is the indicator function and $x_t \in \mathcal{A} \cup \{\emptyset\}$ is the identity of the agent whose idea the principal picks in period t , if any. The principal's profit in a period is zero if she does not implement any idea, and is otherwise equal to the realized profit of the idea that she implements. Her objective is to maximize the expectation of the discounted sum $\sum_{t=0}^{\infty} \delta_0^t y_t$, where δ_0 is the principal's discount factor and $y_t \in \{0, L, H\}$ is her period- t profit.

The players observe which agent's idea is chosen by the principal and the realized value of that idea.⁷ We define a *history* at any period t as the sequence

$$h^t = ((x_0, y_0, S_0), \dots, (x_{t-1}, y_{t-1}, S_{t-1})),$$

where $S_\tau \subseteq \mathcal{A} \cup \{\emptyset\}$ is the set of agents who proposed their ideas in each period $\tau < t$ and, as defined above, x_τ and y_τ denote the implemented idea's proposer and its realized profit, if any.

A *strategy for agent i* determines, for each period t , the probability with which he reports his idea to the principal as a function of his current idea's quality and the history of the game. A *strategy for the principal* determines, for each period t , a lottery over whose idea to select (if any) from among the set of agents currently proposing an idea, given that set of proposers and the history of the game. We apply the notion of *perfect Bayesian Nash*

⁶Note that β may be zero. We discuss the case $\gamma = 1$ in Section 6.

⁷Our results would not change if players could observe more, nor would they change if they only observed the identity of the last agent whose idea yielded low profit for the principal.

equilibrium. We view an equilibrium as a mechanism selected by a principal who is unable to commit. The principal cannot influence nature (the probability of good ideas, and the stochasticity of profit), but would ideally like to overcome the incentive problem of agents. The *first-best outcome* from the principal's point of view is to be able to implement, in every period, a good idea whenever one exists and a bad idea otherwise.

4 Analysis of the Benchmark Model

We think of this game as a mechanism design problem without commitment. The principal wants to design a selection rule to maximize her payoff, but cannot commit to any set of rule. Instead, her rule must be justified endogenously, as an optimal response to that of the agents in equilibrium. Can the principal reach her first best in this circumstance? It turns out that the answer to this question hinges on the ability of the agents in her organization.

A strategy profile achieves the principal's first best if a good idea is implemented in all rounds where at least one agent has a good idea, and a bad idea is implemented in all other rounds.⁸ We say that the principal's first best is *achievable at equilibrium* if there exists $\delta < 1$ and a strategy profile that achieves the principal's first best and that forms a PBNE whenever $\delta_i \geq \delta$, for all $i \in A$.

Our first result provides a characterization of the range of θ 's for which the principal can achieve her first-best when agents are patient enough, provided that profits are sufficiently informative of quality.

Proposition 1. *Define the threshold ability level $\theta^* = 1 - \sqrt[n-1]{\frac{1}{n}}$. Then:*

- (i) *If $\theta < \theta^*$, then the principal's first best cannot be achieved at equilibrium.*
- (ii) *If $\theta > \theta^*$, then the principal's first best is achievable at equilibrium provided that profits are sufficiently informative of quality, meaning that $\frac{1-\beta}{1-\gamma}$ is large enough.*

Proof. We start with the negative result when $\theta < \theta^*$. Suppose that there is a strategy profile that forms a PBNE for $\delta_i = \delta$, for all $i \in \mathcal{A}$, and that achieves the principal's first best. Achieving the principal's first best implies that, at each history h , there is an agent $i(h) \in \mathcal{A}$ such that agents other than $i(h)$ propose good ideas only, $i(h)$ proposes his idea whatever its quality, the principal picks $i(h)$ only when he is the sole proposer, and otherwise picks an agent other than $i(h)$.

⁸Of course, the principal would prefer picking only high-profit ideas when possible, but no one knows at the selection stage which ideas will turn out successful.

An agent j could follow the strategy of proposing his idea in each round, whatever its quality. By doing this, the agent gets picked with probability $(1 - \theta)^{n-1}$ at any history h with $j = i(h)$, and he gets picked with probability at least $(1 - \theta)^{n-2}$ at any history h with $j \neq i(h)$. Each agent can thus secure himself a discounted likelihood of being picked which is larger than or equal to $(1 - \theta)^{n-1}/(1 - \delta)$.

To achieve her first best at equilibrium, the principal picks exactly one agent in each round. So, in total, the aggregate discounted likelihood of being picked is $1/(1 - \delta)$. The equilibrium could not exist if $1/(1 - \delta)$ were strictly smaller than the aggregate discounted likelihood of being picked that agents can secure, that is, n times $(1 - \theta)^{n-1}/(1 - \delta)$. That relationship holds if and only if $\theta < \theta^*$, thereby proving the first part of the result.

The proof of the positive result for $\theta > \theta^*$ is constructive, and will follow from Proposition 2. Indeed, we will define shortly a strategy profile that achieves the principal's first best, and characterize under which conditions on β , γ , and δ it forms a PBNE. The positive result for $\theta > \theta^*$ will follow at once. ■

The threshold ability θ^* in Proposition 1 depends on the number of agents, n . Viewed as a function of n , θ^* decreases in n , and tends to 0 as n tends to infinity. For instance, θ^* is 0.5 when n equals two, and approximately 0.42 and 0.37 for n equal to three and four, respectively. We next show that, regardless of the number of agents, the principal's first best can be achieved through the following strategy profile when $\theta > \theta^*$.

Definition 1 (The *Silent Treatment* Strategy Profile). At each history, one agent is designated as the *agent of last resort*, and all other agents are designated as *discerning*. The agent of last resort proposes his idea independently of its quality, while each discerning agent proposes his idea if and only if it is good. The principal selects the idea proposed by the agent of last resort if it is the only one available. Otherwise, the principal ignores the proposal of the last resort agent and selects among the discerning agents' proposals by randomizing uniformly. The initial agent of last resort is chosen arbitrarily, and remains in that role so long as all the principal's past profits were high. Otherwise, the agent of last resort is the most recent agent whose idea yielded low profit for the principal.

Under the Silent Treatment strategy profile, the principal is sure to implement an idea each period, and will select a good idea whenever one exists. Indeed, if none of the discerning agents have a good idea, then there is always a proposal available from the agent of last resort. This begs the question, when does the silent treatment strategy profile constitute an equilibrium? To this end, it will be important to define the following quantities. At the very beginning of a period – before ideas' qualities are realized – we have:

- the *ex ante* probability that the last resort agent is chosen is $\rho = (1 - \theta)^{n-1}$;

- the *ex ante* probability of being selected as a discerning agent is $\frac{1-\rho}{n-1}$,
- the premium (in terms of the increased *ex ante* probability of selection) from being a discerning agent, instead of the agent of last resort, is $\pi = \frac{1-\rho}{n-1} - \rho = \frac{1-n\rho}{n-1}$.

We are now ready to characterize when the Silent Treatment strategy profile forms a PBNE.

Proposition 2. *The Silent Treatment strategy profile forms a PBNE if and only if for every agent i ,*

$$\delta_i \geq \frac{1}{\gamma + (\gamma - \beta)\pi}.$$

Proof. Assume that players follow the Silent Treatment strategy profile. It is easy to see that neither the agent of last resort, nor the principal, have profitable unilateral deviations. We need to check that a discerning agent wants to propose good ideas, and refrain from proposing bad ideas.

Let σ be the probability that a discerning agent is picked conditional on him proposing his idea, that is, $\sigma = \frac{1-\rho}{(n-1)\theta}$. A discerning agent i will refrain from proposing a bad idea if

$$\delta_i V_i^D \geq \underbrace{\sigma}_{i \text{ selected}} \left((1 - \delta_i)u_i + \underbrace{\beta\delta_i V_i^D}_{\text{high profit}} + \underbrace{(1 - \beta)\delta_i V_i^{LR}}_{\text{low profit, } i \text{ becomes last resort agent}} \right) + \underbrace{(1 - \sigma)\delta_i V_i^D}_{i \text{ not selected}}, \quad (1)$$

where V_i^D and V_i^{LR} , represent i 's average discounted payoff (before learning his idea's quality) under the Silent Treatment strategy profile when he is discerning and when he is the agent of last resort, respectively. Similarly, a discerning agent i will propose a good idea if

$$\underbrace{\sigma}_{i \text{ selected}} \left((1 - \delta_i)u_i + \underbrace{\gamma\delta_i V_i^D}_{\text{high profit}} + \underbrace{(1 - \gamma)\delta_i V_i^{LR}}_{\text{low profit, } i \text{ becomes last resort agent}} \right) + \underbrace{(1 - \sigma)\delta_i V_i^D}_{i \text{ not selected}} \geq \delta_i V_i^D. \quad (2)$$

Let us first examine incentive condition (1). We subtract $\delta_i V_i^{LR}$ from both sides of the inequality (1), and let Δ_i represent $V_i^D - V_i^{LR}$. Then incentive condition (1) is equivalent to

$$\delta_i \Delta_i \geq \sigma(1 - \delta_i)u_i + \sigma\beta\delta_i \Delta_i + (1 - \sigma)\delta_i \Delta_i,$$

which can be rearranged to obtain the inequality

$$\Delta_i \geq \frac{(1 - \delta_i)u_i}{(1 - \beta)\delta_i}. \quad (3)$$

Similar computations show that inequality (2) is equivalent to

$$\Delta_i \leq \frac{(1 - \delta_i)u_i}{(1 - \gamma)\delta_i}. \quad (4)$$

The payoff difference Δ_i from being a discerning agent instead of the last resort agent can be computed through the recursive equations defining V_i^D and V_i^{LR} . Since a discerning agent proposes only good ideas under the Silent Treatment strategy profile, while a last resort agent proposes all ideas but is chosen only when his is the only one available, these equations are:

$$\begin{aligned} V_i^D &= \overbrace{\theta\sigma}^{\text{good idea, selected}} \left((1 - \delta_i)u_i + \gamma\delta_i V_i^D + (1 - \gamma)\delta_i V_i^{LR} \right) + \overbrace{(1 - \theta\sigma)}^{\text{bad idea, or not selected}} \delta_i V_i^D, \\ V_i^{LR} &= \underbrace{\rho}_{\text{selected}} \left((1 - \delta_i)u_i + \delta_i V_i^{LR} \right) + \underbrace{(1 - \rho)}_{\text{not selected}} \left(\underbrace{\gamma\delta_i V_i^{LR} + (1 - \gamma)\delta_i V_i^D}_{\text{low profit, switch to discerning}} \right). \end{aligned} \quad (5)$$

Replacing V_i^D by $V_i^{LR} + \Delta_i$, notice that the expression for V_i^{LR} can be rewritten as

$$V_i^{LR} = \rho(1 - \delta_i)u_i + \delta_i V_i^{LR} + (1 - \rho)(1 - \gamma)\delta_i \Delta_i.$$

Subtracting this new expression for V_i^{LR} from that for V_i^D in (5), we get:

$$\Delta_i = \pi(1 - \delta_i)u_i + \theta\sigma\gamma\delta_i\Delta_i + (1 - \theta\sigma)\delta_i\Delta_i - (1 - \rho)(1 - \gamma)\delta_i\Delta_i,$$

or

$$\Delta_i = \frac{\pi(1 - \delta_i)u_i}{1 - \delta_i + \delta_i(1 - \gamma)(1 + \pi)}.$$

Using this expression for Δ_i , we conclude that the incentive condition (4) for proposing good ideas is always satisfied, and that the incentive condition (3) for withholding bad ideas is satisfied if and only if $\delta_i \geq 1/(\gamma + (\gamma - \beta)\pi)$, as claimed. ■

From Proposition 2, we see that the Silent Treatment strategy profile forms a PBNE when agents are patient enough if and only if $\gamma + (\gamma - \beta)\pi > 1$, or

$$\pi > \frac{1 - \gamma}{\gamma - \beta} = \frac{1}{\tau - 1},$$

where τ is the likelihood ratio $\frac{1 - \beta}{1 - \gamma}$ that a low profit realization arises from a bad idea versus a good idea. Thus, as soon as π is strictly positive, the principal's first best is achievable at equilibrium provided that profits are sufficiently informative of quality. Moreover, π is strictly positive if and only if $\theta > \theta^*$ (as can be seen using $\pi = \frac{1 - n\rho}{n - 1}$ and $\rho = (1 - \theta)^{n-1}$).

This proves the positive result in Proposition 1.

The Silent Treatment strategy profile has several desirable properties. First, the principal and the agents need not observe, nor remember, much information about past behavior. It suffices for them to know at all histories the identity of the current agent of last resort. Second, the principal’s selection rule is optimal for her (thereby providing endogenous commitment) without relying on the agents to punish her if she deviates from it. While efficient equilibria in repeated games oftentimes rely on any deviator to be punished by others, we would find it unnatural if the principal were to follow her part of the equilibrium that achieves her first best only because of the fear of having the agents punish her otherwise. It is difficult to provide a simple definition of what it means for a strategy profile not to rely on the agents to punish the principal. Even so, we can be certain that the Silent Treatment strategy profile does have this feature, since the profile remains a PBNE even when the principal’s discount factor is set to zero. Indeed, notice that the principal’s discount factor does not enter Proposition 2; only the discount factors of the agents matter. Third, as we will now argue in the second main part of the paper, the Silent Treatment strategy profile achieves the principal’s first best in a belief-free way when there is uncertainty about the ability of different agents to have good ideas.

5 Uncertain Abilities

Remember that θ represents the probability of having a good idea in any period. Thus, it measures an agent’s *ability*. So far, agents’ abilities were commonly known and identical. More realistically, suppose that agents may differ in their ability. Each agent i knows his own ability θ_i , but the principal cannot observe it. Agents may or may not know each others’ abilities either. It is only common knowledge that every agent’s ability belongs to an interval $[\underline{\theta}, \bar{\theta}] \subseteq [0, 1]$.⁹ What can the principal do in this case?

This scenario is reminiscent of a multi-armed bandit problem, where pulling an arm in a period is a metaphor for picking an agent’s idea. The new feature, however, is that arms are strategic: they can choose whether to be available in a period. Following the lessons from the multi-armed bandit literature, the first thought might be to study the principal’s optimal tradeoff between ‘experimentation’ to learn about agents’ abilities and ‘exploitation’ by giving priority to the most promising agents. In the classic bandit problem, the Gittins index offers an elegant (but typically not closed-form) solution for which arm to choose each period.

⁹This strictly generalizes our benchmark model, which can be seen as the special case $\underline{\theta} = \bar{\theta} = \theta$.

Applied to our setting, the classic solution falls short of the principal’s first best: experimentation necessarily implies efficiency losses. In this section, we show that the principal can still achieve her first best under incomplete information. As before, using the Silent Treatment strategy profile, she has a simple way to use the repeated nature of her interactions to incentivize the agents. The equilibrium is robust, in the sense that it forms an ex-post PBNE for any realized vector $\vec{\theta} = (\theta_1, \dots, \theta_n)$ of agents’ abilities; that is, it constitutes a belief-free equilibrium. To show this, we must first consider the scenario in which abilities are heterogenous but commonly known.

5.1 Commonly Known Heterogenous Abilities

Consider the ex-post game in which the vector of agents’ abilities is commonly known to be $\vec{\theta}$. Is the Silent Treatment strategy profile still an equilibrium? The behavior prescribed for the principal and agent of last resort are clearly best responses to others’ strategies. It remains to check that a discerning agent is willing to propose good ideas and refrain from proposing bad ideas.

To do this, we must consider an agent’s payoffs and incentives when he is a discerning agent and when he is the agent of last resort, conditional on all players following the Silent Treatment strategy profile. The difficulty here is that unlike in Section 4, an agent’s average discounted payoff depends not only on the different ability levels of agents, but also on the *identity* of the agent of last resort. Indeed, a discerning agent’s payoff depends on how often other discerning agents propose their ideas, which in turn depends on their ability. A discerning agent’s payoff is thus impacted by which of the $n - 1$ other agents is removed from the discerning pool in order to serve as the agent of last resort. Moreover, the agent of last resort will vary over time as low profits are realized.

We will use $V_i^{LR}(\vec{\theta})$ to denote i ’s average discounted payoff under the Silent Treatment strategy profile when he is the current agent of last resort; and use $V_i^D(\vec{\theta}, \ell)$ to denote i ’s average discounted payoff under the Silent Treatment strategy profile when he is discerning and agent $\ell \in \mathcal{A} \setminus \{i\}$ is the current agent of last resort.

Important probabilities. To understand agents’ payoffs and incentives, we must understand the probability with which an agent’s idea is selected by the principal, assuming that the other agents and the principal all follow the Silent Treatment strategy profile. There are different possible circumstances to consider. We let the probability that i is picked when he is the agent of last resort be denoted by $\rho_i(\vec{\theta})$. When ℓ is the agent of last resort, we let the probability that a discerning agent i is picked, *conditional on his proposing an idea*, be denoted by $\sigma_i(\vec{\theta}, \ell)$. When ℓ is the agent of last resort, we let the probability that a discerning

agent j is picked, *conditional on another discerning agent i proposing but not being picked*, be denoted by $p_j(\vec{\theta}, i, \ell)$. Finally, when ℓ is the agent of last resort, we let the probability that a discerning agent j is picked, *conditional on another discerning agent i not proposing*, be denoted by $q_j(\vec{\theta}, i, \ell)$. These probabilities are given as follows:

$$\begin{aligned}\rho_i(\vec{\theta}) &= \prod_{k \neq i} (1 - \theta_k), \\ \sigma_i(\vec{\theta}, \ell) &= \frac{\sum_{S \subseteq \mathcal{A} \setminus \{\ell\}: i \in S} \frac{1}{|S|} \prod_{k \in S} \theta_k \prod_{k \notin S, k \neq \ell} (1 - \theta_k)}{\theta_i}, \\ p_j(\vec{\theta}, i, \ell) &= \frac{\sum_{S \subseteq \mathcal{A} \setminus \{\ell\}: i, j \in S} \frac{1}{|S|} \prod_{k \in S} \theta_k \prod_{k \notin S, k \neq \ell} (1 - \theta_k)}{\theta_i (1 - \sigma_i(\vec{\theta}, \ell))}, \\ q_j(\vec{\theta}, i, \ell) &= \frac{\sum_{S \subseteq \mathcal{A} \setminus \{i, \ell\}: j \in S} \frac{1}{|S|} \prod_{k \in S} \theta_k \prod_{k \notin S, k \neq \ell} (1 - \theta_k)}{1 - \theta_i}.\end{aligned}$$

The expression for $\rho_i(\vec{\theta})$ follows because a last resort agent is selected under the Silent Treatment strategy profile if and only if his is the only proposal, which occurs if and only all discerning agents have bad ideas. To understand the expression for $\sigma_i(\vec{\theta}, \ell)$, observe that while agent i 's proposal is selected uniformly from among any set of discerning agents' proposals, we must consider all different possible sets of proposers and their probabilities. The probabilities $\rho_i(\vec{\theta})$ and $\sigma_i(\vec{\theta}, \ell)$ are needed to characterize the equilibrium value functions of agents. The final two probabilities $p_j(\vec{\theta}, i, \ell)$ and $q_j(\vec{\theta}, i, \ell)$, whose expressions follow from similar reasoning, will be needed to capture incentive conditions. We begin by studying the latter.

Incentive conditions in terms of equilibrium payoffs. With these probabilities in mind, if ℓ is the agent of last resort, then the incentive condition for a discerning agent i not to propose a bad idea is given by:

$$\begin{aligned}& \overbrace{\frac{\rho_\ell(\vec{\theta})}{1 - \theta_i}}^{\text{last resort agent chosen}} \delta_i V_i^D(\vec{\theta}, \ell) + \sum_{j \neq i, \ell} \overbrace{q_j(\vec{\theta}, i, \ell)}^{\text{discerning } j \text{ chosen}} \left(\gamma \delta_i V_i^D(\vec{\theta}, \ell) + \overbrace{(1 - \gamma) \delta_i V_i^D(\vec{\theta}, j)}^{\text{low profit, } j \text{ becomes last resort}} \right) \\ & \geq \overbrace{\sigma_i(\vec{\theta}, \ell)}^{i \text{ chosen}} \left((1 - \delta_i) u_i + \beta \delta_i V_i^D(\vec{\theta}, \ell) + \overbrace{(1 - \beta) \delta_i V_i^{LR}(\vec{\theta})}^{\text{low profit, } i \text{ becomes last resort}} \right) \\ & \quad + (1 - \sigma_i(\vec{\theta}, \ell)) \sum_{j \neq i, \ell} \overbrace{p_j(\vec{\theta}, i, \ell)}^{\text{discerning } j \text{ chosen instead}} \left(\gamma \delta_i V_i^D(\vec{\theta}, \ell) + \overbrace{(1 - \gamma) \delta_i V_i^D(\vec{\theta}, j)}^{\text{low profit, } j \text{ becomes last resort}} \right).\end{aligned} \tag{IC_b}$$

Similarly, the incentive condition for i to report a good idea when he is discerning and ℓ is the agent of last resort, is:

$$\begin{aligned}
& \sigma_i(\vec{\theta}, \ell) \left((1 - \delta_i)u_i + \gamma\delta_i V_i^D(\vec{\theta}, \ell) + (1 - \gamma)\delta_i V_i^{LR}(\vec{\theta}) \right) \\
& + (1 - \sigma_i(\vec{\theta}, \ell)) \sum_{j \neq i, \ell} p_j(\vec{\theta}, i, \ell) \left(\gamma\delta_i V_i^D(\vec{\theta}, \ell) + (1 - \gamma)\delta_i V_i^D(\vec{\theta}, j) \right) \\
& \geq \frac{\rho_\ell(\vec{\theta})}{1 - \theta_i} \delta_i V_i^D(\vec{\theta}, \ell) + \sum_{j \neq i, \ell} q_j(\vec{\theta}, i, \ell) \left(\gamma\delta_i V_i^D(\vec{\theta}, \ell) + (1 - \gamma)\delta_i V_i^D(\vec{\theta}, j) \right),
\end{aligned} \tag{IC_g}$$

which differs from Condition IC_b both in the direction of the inequality and because the probability that agent i 's idea generates low profit is γ instead of β .

Incentive conditions IC_b and IC_g are linear in the equilibrium payoffs. It turns out that they depend on these payoffs *only* through the difference in average discounted payoffs from being discerning instead of being the agent of last resort, as the preliminary result below highlights. Because agents are heterogenous, the payoff difference depends on the identity of the agent of last resort. For each agent $i \in \mathcal{A}$ and each possible agent of last resort $\ell \neq i$, we define the payoff difference

$$\Delta V_i(\vec{\theta}, \ell) = V_i^D(\vec{\theta}, \ell) - V_i^{LR}(\vec{\theta}).$$

Let $\Delta \vec{V}_i(\vec{\theta})$ denote the $(n - 1)$ -column vector obtained by varying the agent of last resort in $\mathcal{A} \setminus \{i\}$. We next define two matrices to help state the result. For each i and $\vec{\theta}$, let $M_i^g(\vec{\theta})$ be the $(n - 1)$ -square matrix whose $\ell\ell'$ entry, for all $\ell, \ell' \in \mathcal{A} \setminus \{i\}$, is given by

$$[M_i^g(\vec{\theta})]_{\ell\ell'} = \begin{cases} q_{\ell'}(\vec{\theta}, i, \ell) - p_{\ell'}(\vec{\theta}, i, \ell)(1 - \sigma_i(\vec{\theta}, \ell)) & \text{if } \ell \neq \ell', \\ \rho_\ell(\vec{\theta})/(1 - \theta_i) & \text{if } \ell = \ell'. \end{cases}$$

The diagonal entries of $M_i^g(\vec{\theta})$ capture the probability that all agents other than i and ℓ have bad ideas; while the off-diagonal entries capture the increased probability with which another discerning agent is selected when agent i does not propose his idea, as compared to when i does propose. Next, define $M_i^b(\vec{\theta})$ to be the $(n - 1)$ -square matrix constructed from $M_i^g(\vec{\theta})$ by adding to it the diagonal matrix whose $\ell\ell$ -entry is $\frac{\gamma - \beta}{1 - \gamma} \sigma_i(\vec{\theta}, \ell)$. Finally, $\vec{\sigma}_i(\vec{\theta})$ is the $(n - 1)$ -column vector whose ℓ -th entry, for all $\ell \neq i$, is $\sigma_i(\vec{\theta}, \ell)$. The following lemma, proved in the appendix, characterizes the equilibrium conditions in terms of payoff differences.

Lemma 1. *The Silent Treatment strategy profile constitutes a PBNE of the ex-post game*

with abilities $\vec{\theta}$ if and only if

$$\frac{\delta_i(1-\gamma)}{u_i(1-\delta_i)} M_i^g(\vec{\theta}) \Delta \vec{V}_i(\vec{\theta}) \leq \vec{\sigma}_i(\vec{\theta}) \leq \frac{\delta_i(1-\gamma)}{u_i(1-\delta_i)} M_i^b \Delta \vec{V}_i(\vec{\theta})(\vec{\theta}).$$

Equilibrium Payoffs and Payoff Differences. The above result provides a preliminary characterization of the equilibrium conditions as a function of the average discounted payoff differences between being discerning and being the agent of last resort. We now delve further into agents' payoffs, with the goal of characterizing the payoff differences in terms of exogenous variables only. An agent i 's average discounted payoff $V_i^D(\vec{\theta}, \ell)$ when he is discerning and agent ℓ is the agent of last resort, and his average discounted payoff $V_i^{LR}(\vec{\theta})$ when he is the agent of last resort himself, are jointly determined by the following, recursive system of equations for all possible agents $\ell \neq i$:

$$\begin{aligned} V_i^{LR}(\vec{\theta}) &= \overbrace{\rho_i(\vec{\theta})}^{i \text{ is chosen}} \left((1-\delta_i)u_i + \delta_i V_i^{LR}(\vec{\theta}) \right) + \sum_{j \neq i} \overbrace{\theta_j \sigma_j(\vec{\theta}, i)}^{j \text{ chosen when } i \text{ is last resort}} \left(\gamma \delta_i V_i^{LR}(\vec{\theta}) + \overbrace{(1-\gamma) \delta_i V_i^D(\vec{\theta}, j)}^{\text{low profit, } j \text{ becomes last resort}} \right), \\ V_i^D(\vec{\theta}, \ell) &= \overbrace{\theta_i \sigma_i(\vec{\theta}, \ell)}^{i \text{ chosen when } \ell \text{ is last resort}} \left((1-\delta_i)u_i + \gamma \delta_i V_i^D(\vec{\theta}, \ell) + \overbrace{(1-\gamma) \delta_i V_i^{LR}(\vec{\theta})}^{\text{low profit, } i \text{ becomes last resort}} \right) \\ &\quad + \sum_{j \neq i, \ell} \overbrace{\theta_j \sigma_j(\vec{\theta}, \ell)}^{j \text{ chosen when } \ell \text{ is last resort}} \left(\gamma \delta_i V_i^D(\vec{\theta}, \ell) + \overbrace{(1-\gamma) \delta_i V_i^D(\vec{\theta}, j)}^{\text{low profit, } j \text{ becomes last resort}} \right) + \overbrace{\rho_\ell(\vec{\theta})}^{\ell \text{ is chosen}} \delta_i V_i^D(\vec{\theta}, \ell). \end{aligned} \tag{6}$$

In the appendix, we manipulate the system of equations (6) to derive the average discounted payoff differences. It turns out that the payoff differences depend on the vector of abilities $\vec{\theta}$ only through the likelihood premiums of being picked when discerning versus when the agent of last resort. By contrast to Section 4, under heterogenous abilities there are many such premiums to consider, as the probability of being picked when discerning depends on the vector of abilities $\vec{\theta}$ as well as the identity of the agent of last resort. Formally, for each ℓ and i in \mathcal{A} , let

$$\pi_{\ell i}(\vec{\theta}) = \theta_i \sigma_i(\vec{\theta}, \ell) - \rho_i(\vec{\theta})$$

be agent i 's likelihood premium when ℓ would be the agent of last resort. For each i and each $\vec{\theta}$, let $\vec{\pi}_i(\vec{\theta})$ be the $(n-1)$ -column vector whose ℓ -component is $\pi_{\ell i}(\vec{\theta})$. This vector thus lists the likelihood premiums that are relevant for i , as a function of the agent of last resort. Using the likelihood premiums, define $B_i(\vec{\theta})$ to be the $(n-1)$ -square matrix whose $\ell \ell'$ -entry,

for any ℓ, ℓ' in $\mathcal{A} \setminus \{i\}$, is given by

$$[B_i(\vec{\theta})]_{\ell\ell'} = \begin{cases} \pi_{i\ell'}(\vec{\theta}) - \pi_{i\ell}(\vec{\theta}) & \text{if } \ell \neq \ell', \\ 1 + \pi_{i\ell}(\vec{\theta}) + (1 - \delta_i)/(\delta_i(1 - \gamma)) & \text{if } \ell = \ell'. \end{cases}$$

The next lemma, which is proved in the appendix, shows how the likelihood premiums characterize the payoff differences through the matrix $B_i(\vec{\theta})$.

Lemma 2. *For all i and $\vec{\theta}$, the average discounted payoff differences $\Delta\vec{V}_i(\vec{\theta})$ satisfy the following equation:*

$$B_i(\vec{\theta})\Delta\vec{V}_i(\vec{\theta}) = \frac{u_i(1 - \delta_i)}{\delta_i(1 - \gamma)}\vec{\pi}_i(\vec{\theta}).$$

Equilibrium conditions in terms of exogenous variables. If we could solve for $\Delta\vec{V}_i(\vec{\theta})$ above, then we could combine the previous two lemmas to characterize under which circumstances our strategy profile forms an equilibrium of the ex-post game with abilities $\vec{\theta}$. We show in the appendix that the matrix $B_i(\vec{\theta})$ is strictly diagonally dominant, and therefore invertible.

Proposition 3. *The Silent Treatment strategy profile constitutes a PBNE of the ex-post game with abilities $\vec{\theta}$ if and only if*

$$M_i^g(\vec{\theta})B_i(\vec{\theta})^{-1}\vec{\pi}_i(\vec{\theta}) \leq \vec{\sigma}_i(\vec{\theta}) \leq M_i^b(\vec{\theta})B_i(\vec{\theta})^{-1}\vec{\pi}_i(\vec{\theta}).$$

As can be seen from our analysis, the equilibrium conditions are independent of the principal's discount factor δ_0 , which means that they would hold even if the principal were fully myopic. The equilibrium thus doesn't require that the principal's behavior be enforced by the threat of punishments from agents, which we consider a natural property in a mechanism design context where the principal is the authority. Note, in addition, that the equilibrium conditions are also independent of the payoff u_i each agent i gets when selected. We next turn to the question of whether the Silent Treatment strategy profile forms an equilibrium for all possible ability levels.

5.2 The Silent Treatment as a Belief-Free Equilibrium

The principal may have little information about agents' abilities and would like to guarantee her first-best outcome in all cases. The notion of *belief-free equilibrium* directly addresses the question of equilibrium robustness. The Silent Treatment strategy profile is a belief-free equilibrium if it forms a PBNE for any realized vector of abilities $\vec{\theta}$ in the set $[\underline{\theta}, \bar{\theta}]^A$ of all

possible abilities. The principal's first best is achievable by a belief-free equilibrium if there exists $\delta < 1$ and a strategy profile that achieves the principal's first best and that forms a belief-free equilibrium whenever $\delta_i \geq \delta$, for all $i \in \mathcal{A}$. Proposition 3 can then be used to prove a belief-free extension of our first main result.

Proposition 4. *Consider the ability threshold θ^* defined in Proposition 1. We have:*

- (i) *If $\underline{\theta} < \theta^*$, then the principal's first-best cannot be achieved in any belief-free equilibrium.*
- (ii) *If $\underline{\theta} > \theta^*$, then there exists an informativeness threshold $\tau > 0$ such that, for all (β, γ) with $\frac{1-\beta}{1-\gamma} \geq \tau$, the principal's first best is achievable by a belief-free equilibrium, namely, the Silent Treatment strategy profile.*

The principal's ability to achieve her first best in this setting thus hinges on her worst possible agent, the organization's 'weakest link.' Only when she is certain that the agents all have abilities greater than θ^* can she incentivize them to be discerning. A principal may or may not be able to screen agents to ensure a minimal standard for entry to the organization. The threshold θ^* decreases in the number of agents n , and is always smaller than $1/2$, so it would suffice that agents are simply more likely to have good ideas than bad ones.

Part (i) of Proposition 4, the negative result, follows from the corresponding part of Proposition 1. Indeed, it is more difficult to form an equilibrium for every possible vector $\vec{\theta}$ of ability levels than it is to form an equilibrium for some common ability level θ . By contrast, heterogeneity in abilities makes part (ii) of the result quite challenging to prove. Even though Proposition 3 provides a helpful characterization of equilibrium conditions for the Silent Treatment strategy profile for a given $\vec{\theta}$, it is not obvious *a priori* whether these inequalities can be satisfied when abilities are heterogenous. As an illustration of the difficulties involved, it is not straightforward to check whether the payoff differences $\Delta V_i(\vec{\theta})$ are positive, because the solution to a linear system may be nonpositive even when all the parameters in the system are positive themselves. In the appendix, we overcome this and other challenges by using special properties of the matrices $M_i^g(\vec{\theta})$, $M_i^b(\vec{\theta})$ and $B_i(\vec{\theta})$. Importantly, the likelihood premiums must be strictly positive to make the equilibrium possible, which we show holds if and only if $\underline{\theta} > \theta^*$. Belief-freeness is achieved because the conditions in Proposition 3 hold strictly, and uniformly across all $\vec{\theta}$, when profits are sufficiently informative and agents are sufficiently patient. That is, the agents have strict incentives to follow the Silent Treatment strategy profile when it constitutes an equilibrium.

6 When Good Ideas Give High Profit for Sure

Our analysis so far has restricted attention to $\gamma < 1$, that is, settings in which the principal cannot conclude an idea was bad when low profit is realized. We now consider the special case $\gamma = 1$ in which good ideas surely deliver high profit when implemented. As before, we permit any $\beta \in [0, \gamma)$. Indeed, so long as we have $\beta < \gamma$ to maintain the distinction that ‘good’ ideas are better for the principal than ‘bad’ ones, the probability β with which bad ideas deliver high profit does not affect our results.

The analysis in Section 4 determining when the Silent Treatment strategy profile forms a belief-free equilibrium also applies when $\gamma = 1$. Because the matrices Lemmas 1 and 2 take a diagonal form when $\gamma = 1$, it becomes possible to find a closed-form characterization of the range of parameters for which the silent-treatment strategies form an equilibrium of the ex-post game with abilities $\vec{\theta}$. The ratio $\frac{1-\beta}{1-\gamma}$ is infinitely large when $\beta < 1$ and $\gamma = 1$, and so one would expect part (ii) of Proposition 4 to hold. The following proposition, which is proved in the appendix, confirms this intuition.

Proposition 5. *Suppose $\gamma = 1$. The silent-treatment strategy profile is an equilibrium of the ex-post game with abilities $\vec{\theta}$ if and only if the following inequality holds for all i :*

$$\delta_i \geq \frac{1}{1 + (1 - \beta) \min_{\ell \neq i} \pi_{\ell i}(\vec{\theta})}.$$

There exists a patience threshold δ such that the silent-treatment strategy profile forms a belief-free PBNE when $\delta_i \geq \delta$ for all agents i if and only if $\underline{\theta} > \theta^$.*

Unlike part (ii) of Proposition 4, part (i) does not continue to hold when $\gamma = 1$. The principal’s first-best can be achieved by other belief-free equilibria that coincide with the Silent Treatment strategy profile on the equilibrium path (i.e., so long as the principal has always received high profits) but impose harsher punishments off path.¹⁰ For instance, suppose players start by following the Silent Treatment strategy with an arbitrary agent i as last resort. However, as soon as a low profit occurs by implementing the idea of some discerning agent j , the following strategy profile is played forever after: agent i proposes his idea regardless of its quality, agents other than i report only low-quality ideas, and the principal picks i ’s ideas when he proposes and otherwise picks uniformly among proposed ideas. The principal’s profit by following this strategy is clearly suboptimal should cheating actually occur.

¹⁰Such considerations did not emerge for $\gamma < 1$ since no relevant history falls off the equilibrium path when implementing the Silent Treatment strategy profile in that case.

In the context of mechanism design without commitment, one could think of the principal as a special player who can make any equilibrium focal. If the principal’s first best can be achieved on path, then why would she adhere to an inefficient off-path payoff if she could make the on-path equilibrium strategies salient once again?

We say that the principal’s first best is *achievable in a credible way* if there exists $\delta < 1$ and a strategy profile that achieves the principal’s first best in the stage game played after each history, and that forms an equilibrium whenever $\delta_i \geq \delta$, for all $i \in A$. Notice that the Silent Treatment strategy profile does achieve the principal’s first best in this sense, (indeed, for any $\gamma \leq 1$). In addition, the first part of Proposition 4 does hold for $\gamma = 1$ when adding this requirement of credibility.

Proposition 6. *If $\underline{\theta} < \theta^*$, then the principal’s first best cannot be achieved by a belief-free equilibrium in a credible way.*

Propositions 5 and 6 then allow us to conclude that when $\gamma = 1$ (and independently of β), a belief-free equilibrium achieves the principal’s first best in a credible way if and only if her first best is achieved by the Silent Treatment strategy profile.

7 Concluding Remarks

Information overload is endemic to every organization where limited cognitive resources, multiple obligations, and short deadlines can lead managers to overlook important ideas from subordinates. We propose an approach to this problem that treats the set of items demanding a manager’s attention as *endogenous*, in the sense that this set is generated by *strategic* agents. In an environment where a principal repeatedly interacts with agents, we ask how can a principal provide non-monetary dynamic incentives for agents to be discerning, so that they communicate information only when it is important.

We address this question using a simple, stylized model that allows us to lay out our arguments transparently. Our results suggest that, even in the absence of commitment and monetary incentives, the optimal outcome can be achieved with a very simple and intuitive strategy that does not require complex probabilistic inferences. Thus, we demonstrate that the concept of belief-free equilibrium, which has been applied in repeated games by constructing mixed-strategy profiles in which players are indifferent over all continuation paths, can be used to construct intuitive strategy profiles involving strict incentives in a repeated principal-multi-agent problem.

We view our framework and results as a first step towards a more general understanding of what is the best a principal can achieve, what incentives the principal should use and

how robust are these incentives to incomplete information. While our framework is very simple, showing that the first-best outcome is achievable in a belief-free equilibrium is not as simple. This suggests that analyzing a more general framework may require even more involved arguments, and we therefore leave this analysis for future research.

Appendix

A. Preliminaries

We collect here several useful definitions and observations. Remember that

$$\begin{aligned}\rho_i(\vec{\theta}) &= \prod_{k \neq i} (1 - \theta_k), \\ \sigma_i(\vec{\theta}, \ell) &= \frac{\sum_{S \subseteq \mathcal{A} \setminus \{\ell\}: i \in S} \frac{1}{|S|} \prod_{k \in S} \theta_k \prod_{k \notin S, k \neq \ell} (1 - \theta_k)}{\theta_i}, \\ p_j(\vec{\theta}, i, \ell) &= \frac{\sum_{S \subseteq \mathcal{A} \setminus \{\ell\}: i, j \in S} \frac{1}{|S|} \prod_{k \in S} \theta_k \prod_{k \notin S, k \neq \ell} (1 - \theta_k)}{\theta_i (1 - \sigma_i(\vec{\theta}, \ell))}, \\ q_j(\vec{\theta}, i, \ell) &= \frac{\sum_{S \subseteq \mathcal{A} \setminus \{i, \ell\}: j \in S} \frac{1}{|S|} \prod_{k \in S} \theta_k \prod_{k \notin S, k \neq \ell} (1 - \theta_k)}{1 - \theta_i}.\end{aligned}$$

Remark 1. Observe that $\sum_{j \neq \ell} \theta_j \sigma_j(\vec{\theta}, \ell) + \rho_\ell(\vec{\theta}) = 1$, since the principal always selects some agent, resorting to the last resort agent if no discerning agent proposes. Moreover, note that $\sum_{j \neq i, \ell} p_j(\vec{\theta}, i, \ell) = 1$, since the fact that player i has proposed means that the selected agent will come from the discerning pool. On the other hand, $\sum_{j \neq i, \ell} q_j(\vec{\theta}, i, \ell) + \frac{\rho_\ell(\vec{\theta})}{1 - \theta_i} = 1$, since it is possible that no discerning agent will propose.

For each agent $i \in \mathcal{A}$, and agent of last resort $\ell \neq i$, the aggregate discounted payoff difference under the Silent Treatment strategy profile is

$$\Delta V_i(\vec{\theta}, \ell) = V_i^D(\vec{\theta}, \ell) - V_i^{LR}(\vec{\theta}).$$

We let $\Delta \vec{V}_i(\vec{\theta})$ denote the $(n - 1)$ -column vector obtained by varying the agent of last resort in $\mathcal{A} \setminus \{i\}$. Let $\vec{\sigma}_i(\vec{\theta})$ be the $(n - 1)$ -column vector whose ℓ -th entry, for all $\ell \neq i$, equals $\sigma_i(\vec{\theta}, \ell)$. For each ℓ and i in \mathcal{A} , we define

$$\pi_{\ell i}(\vec{\theta}) = \theta_i \sigma_i(\vec{\theta}, \ell) - \rho_i(\vec{\theta})$$

as agent i 's likelihood premium, capturing the additional probability with which i is selected as a discerning agent when ℓ would be the agent of last resort versus when i himself is the

agent of last resort.

We now define three matrices. For each i and $\vec{\theta}$, let $M_i^g(\vec{\theta})$ be the $(n-1)$ -square matrix whose $\ell\ell'$ -entry, for all $\ell, \ell' \in \mathcal{A} \setminus \{i\}$, is given by

$$[M_i^g(\vec{\theta})]_{\ell\ell'} = \begin{cases} q_{\ell'}(\vec{\theta}, i, \ell) - p_{\ell'}(\vec{\theta}, i, \ell)(1 - \sigma_i(\vec{\theta}, \ell)) & \text{if } \ell \neq \ell', \\ \frac{\rho_{\ell}(\vec{\theta})}{1-\theta_i} & \text{if } \ell = \ell'. \end{cases}$$

For each i and $\vec{\theta}$, let $M_i^b(\vec{\theta})$ be the $(n-1)$ -square matrix whose $\ell\ell'$ -entry, for all $\ell, \ell' \in \mathcal{A} \setminus \{i\}$, is given by

$$[M_i^b(\vec{\theta})]_{\ell\ell'} = \begin{cases} q_{\ell'}(\vec{\theta}, i, \ell) - p_{\ell'}(\vec{\theta}, i, \ell)(1 - \sigma_i(\vec{\theta}, \ell)) & \text{if } \ell \neq \ell', \\ \frac{\rho_{\ell}(\vec{\theta})}{1-\theta_i} + \frac{\gamma-\beta}{1-\gamma}\sigma_i(\vec{\theta}, \ell) & \text{if } \ell = \ell'. \end{cases}$$

Remark 2. Note $M_i^b(\vec{\theta})$ can be derived from $M_i^g(\vec{\theta})$ by adding $\frac{\gamma-\beta}{1-\gamma}\sigma_i(\vec{\theta}, \ell)$ on each $\ell\ell$ -entry.

Finally, for each i and $\vec{\theta}$, define $B_i(\vec{\theta})$ to be the $(n-1)$ -square matrix whose $\ell\ell'$ -entry, for any ℓ, ℓ' in $\mathcal{A} \setminus \{i\}$, is given by

$$[B_i(\vec{\theta})]_{\ell\ell'} = \begin{cases} \pi_{i\ell'}(\vec{\theta}) - \pi_{\ell\ell'}(\vec{\theta}) & \text{if } \ell \neq \ell', \\ 1 + \pi_{i\ell}(\vec{\theta}) + (1 - \delta_i)/(\delta_i(1 - \gamma)) & \text{if } \ell = \ell'. \end{cases}$$

B. Proof of Proposition 3

The proof proceeds through a series of lemmas. Lemmas 1 and 2 have already been stated, but not proved, in the text.

Proof of Lemma 1. First note that the Silent Treatment strategy of the principal is first best for him, regardless of his discount factor and agents' types, so long as agents follow their strategies. Moreover, given that the principal follows this strategy, a last resort agent cannot change his probability of going back into the discerning pool of agents by his own actions. The last resort agent thus finds it optimal to propose any idea with probability one, regardless of his discount factor and agents' types. It remains to check the incentive conditions for discerning agents.

Subtracting $\delta_i V_i^{LR}(\vec{\theta})$ from both sides of the incentive condition (IC_b) for i to withhold

a bad idea when ℓ is the last resort agent, we find that

$$\begin{aligned} & \frac{\rho_\ell(\vec{\theta})}{1-\theta_i} \delta_i \Delta V_i^D(\vec{\theta}, \ell) + \sum_{j \neq i, \ell} q_j(\vec{\theta}, i, \ell) \left(\gamma \delta_i \Delta V_i^D(\vec{\theta}, \ell) + (1-\gamma) \delta_i \Delta V_i^D(\vec{\theta}, j) \right) \\ & \geq \sigma_i(\vec{\theta}, \ell) \left((1-\delta_i) u_i + \beta \delta_i \Delta V_i^D(\vec{\theta}, \ell) \right) \\ & \quad + (1-\sigma_i(\vec{\theta}, \ell)) \sum_{j \neq i, \ell} p_j(\vec{\theta}, i, \ell) \left(\gamma \delta_i \Delta V_i^D(\vec{\theta}, \ell) + (1-\gamma) \delta_i \Delta V_i^D(\vec{\theta}, j) \right). \end{aligned}$$

Collect all ΔV_i^D terms on the left-hand side, and multiply the inequality through by $\frac{1}{(1-\delta_i)u_i}$. Then, for each $j \neq \ell$, the coefficient multiplying $\frac{(1-\gamma)\delta_i}{(1-\delta_i)u_i} \Delta V_i^D(\vec{\theta}, j)$ is easily seen to be $[M_i^b(\vec{\theta})]_{\ell j}$. The coefficient multiplying $\frac{(1-\gamma)\delta_i}{(1-\delta_i)u_i} \Delta V_i^D(\vec{\theta}, \ell)$ is

$$\begin{aligned} & \frac{1}{1-\gamma} \left(\frac{\rho_\ell(\vec{\theta})}{1-\theta_i} + \gamma \sum_{j \neq i, \ell} q_j(\vec{\theta}, i, \ell) - \beta \sigma_i(\vec{\theta}, \ell) - \gamma(1-\sigma_i(\vec{\theta}, \ell)) \sum_{j \neq i, \ell} p_j(\vec{\theta}, i, \ell) \right) \\ & = \frac{1}{1-\gamma} \left(\frac{\rho_\ell(\vec{\theta})}{1-\theta_i} + \gamma \left(1 - \frac{\rho_\ell(\vec{\theta})}{1-\theta_i} \right) - \beta \sigma_i(\vec{\theta}, \ell) - \gamma(1-\sigma_i(\vec{\theta}, \ell)) \right) \\ & = [M_i^b(\vec{\theta})]_{\ell \ell}, \end{aligned}$$

where the first equality follows from Remark 1. Stacking the inequalities for $\ell \neq i$ yields the matrix inequality with $M_i^b(\vec{\theta})$.

Next, subtracting $\delta_i V_i^{LR}(\vec{\theta})$ from both sides of the incentive condition (IC_g) for agent i to propose a good idea when ℓ is the last resort agent, we find that

$$\begin{aligned} & \sigma_i(\vec{\theta}, \ell) \left((1-\delta_i) u_i + \gamma \delta_i \Delta V_i^D(\vec{\theta}, \ell) \right) \\ & \quad + (1-\sigma_i(\vec{\theta}, \ell)) \sum_{j \neq i, \ell} p_j(\vec{\theta}, i, \ell) \left(\gamma \delta_i \Delta V_i^D(\vec{\theta}, \ell) + (1-\gamma) \delta_i \Delta V_i^D(\vec{\theta}, j) \right) \\ & \geq \sum_{j \neq i, \ell} q_j(\vec{\theta}, i, \ell) \left(\gamma \delta_i \Delta V_i^D(\vec{\theta}, \ell) + (1-\gamma) \delta_i \Delta V_i^D(\vec{\theta}, j) \right) + \frac{\rho_\ell(\vec{\theta})}{1-\theta_i} \delta_i \Delta V_i^D(\ell, \theta, \delta_i, \gamma). \end{aligned}$$

Collect all ΔV_i^D -terms on the right-hand side, and multiply the inequality through by $\frac{1}{(1-\delta_i)u_i}$. Then the coefficient multiplying $\frac{(1-\gamma)\delta_i}{(1-\delta_i)u_i} \Delta V_i^D(\vec{\theta}, j)$ is easily seen to be $[M_i^g(\vec{\theta})]_{\ell j}$. The coefficient multiplying $\frac{(1-\gamma)\delta_i}{(1-\delta_i)u_i} \Delta V_i^D(\vec{\theta}, \ell)$ reduces to

$$\frac{1}{1-\gamma} \left(\gamma \sum_{j \neq i, \ell} q_j(\vec{\theta}, i, \ell) + \frac{\rho_\ell(\vec{\theta})}{1-\theta_i} - \gamma \right) = [M_i^g(\vec{\theta})]_{\ell \ell},$$

where the equality follows from Remark 1. Stacking the inequalities for $\ell \neq i$ yields the

matrix inequality with $M_i^g(\vec{\theta})$. ■

Proof of Lemma 2. The value function V_i^D is defined by the equation

$$\begin{aligned} V_i^D(\vec{\theta}, \ell) &= \theta_i \sigma_i(\vec{\theta}, \ell) \left((1 - \delta_i) u_i + \gamma \delta_i V_i^D(\vec{\theta}, \ell) + (1 - \gamma) \delta_i V_i^{LR}(\vec{\theta}) \right) \\ &+ \sum_{j \neq i, \ell} \theta_j \sigma_j(\vec{\theta}, \ell) \left(\gamma \delta_i V_i^D(\vec{\theta}, \ell) + (1 - \gamma) \delta_i V_i^D(\vec{\theta}, j) \right) + \rho_\ell(\vec{\theta}) \delta_i V_i^D(\vec{\theta}, \ell), \end{aligned} \quad (7)$$

while the value function V_i^{LR} is defined by

$$\begin{aligned} V_i^{LR}(\vec{\theta}) &= \rho_i(\vec{\theta}) \left((1 - \delta_i) u_i + \delta_i V_i^{LR}(\vec{\theta}) \right) \\ &+ \sum_{j \neq i} \theta_j \sigma_j(\vec{\theta}, i) \left(\gamma \delta_i V_i^{LR}(\vec{\theta}) + (1 - \gamma) \delta_i V_i^D(\vec{\theta}, j) \right). \end{aligned} \quad (8)$$

Subtracting $\delta_i V_i^{LR}(\vec{\theta})$ from both sides of Equation (7), we find that

$$\begin{aligned} V_i^D(\vec{\theta}, \ell) - \delta_i V_i^{LR}(\vec{\theta}) &= \theta_i \sigma_i(\vec{\theta}, \ell) \left((1 - \delta_i) u_i + \gamma \delta_i \Delta V_i^D(\vec{\theta}, \ell) \right) \\ &+ \sum_{j \neq i, \ell} \theta_j \sigma_j(\vec{\theta}, \ell) \left(\gamma \delta_i \Delta V_i^D(\vec{\theta}, \ell) + (1 - \gamma) \delta_i \Delta V_i^D(\vec{\theta}, j) \right) + \rho_\ell(\vec{\theta}) \delta_i \Delta V_i^D(\vec{\theta}, \ell), \end{aligned} \quad (9)$$

In view of Remark 1, Equation (7) simplifies to

$$\begin{aligned} V_i^D(\vec{\theta}, \ell) - \delta_i V_i^{LR}(\vec{\theta}) &= \theta_i \sigma_i(\vec{\theta}, \ell) (1 - \delta_i) u_i + (1 - \gamma) \delta_i \sum_{j \neq i, \ell} \theta_j \sigma_j(\vec{\theta}, \ell) \Delta V_i^D(\vec{\theta}, j) \\ &+ \delta_i \Delta V_i^D(\vec{\theta}, \ell) \left(\gamma + (1 - \gamma) \rho_\ell(\vec{\theta}) \right). \end{aligned} \quad (10)$$

Similarly, subtracting $\delta_i V_i^{LR}(\vec{\theta})$ from both sides of Equation (8), we find that

$$V_i^{LR}(\vec{\theta}) - \delta_i V_i^{LR}(\vec{\theta}) = \rho_i(\vec{\theta}) (1 - \delta_i) u_i + (1 - \gamma) \delta_i \sum_{j \neq i} \theta_j \sigma_j(\vec{\theta}, i) \Delta V_i^D(\vec{\theta}, j). \quad (11)$$

Subtracting Equation (11) from Equation (10), and using the definition of $\pi_{\ell\ell'}(\vec{\theta})$, we find that:

$$\begin{aligned} \Delta V_i^D(\vec{\theta}, \ell) &= \pi_{\ell i}(\vec{\theta}) (1 - \delta_i) u_i + \delta_i \Delta V_i^D(\vec{\theta}, j) \left(\gamma + (1 - \gamma) \pi_{i\ell}(\vec{\theta}) \right) \\ &+ (1 - \gamma) \delta_i \sum_{j \neq i, \ell} \left(\theta_j \sigma_j(\vec{\theta}, \ell) - \theta_j \sigma_j(\vec{\theta}, i) \right) \Delta V_i^D(\vec{\theta}, j). \end{aligned} \quad (12)$$

Note that $\theta_j \sigma_j(\vec{\theta}, \ell) - \theta_j \sigma_j(\vec{\theta}, i) = \pi_{\ell j}(\vec{\theta}) - \pi_{i j}(\vec{\theta})$. We can thus rearrange Equation (12) and

divide through by $(1 - \gamma)\delta_i$ to find that $B_i(\vec{\theta})\Delta\vec{V}_i(\vec{\theta}) = \frac{(1-\delta_i)u_i}{(1-\gamma)\delta_i}\vec{\pi}_i(\vec{\theta})$, as claimed. ■

Clearly, likelihood premiums must be smaller than one. We next show that they are strictly positive if and only if agents' abilities are all higher than θ^* .

Lemma 3. $\pi_{\ell\ell'}(\vec{\theta}) > 0$ for all $\vec{\theta} \in [\underline{\theta}, \bar{\theta}]$ and any two distinct ℓ, ℓ' in A , if and only if $\underline{\theta} > \theta^*$.

Proof. Notice that $\pi_{\ell\ell'}(\vec{\theta}) > 0$ if and only if

$$\theta_{\ell'} > \frac{\rho_{\ell'}(\vec{\theta})}{\sigma_{\ell'}(\vec{\theta}, \ell)}.$$

We now show that the expression on the right-hand side is decreasing in θ_k , for all $k \in A$. To this end, observe that

$$\begin{aligned} \frac{\rho_{\ell'}(\vec{\theta})}{\sigma_{\ell'}(\vec{\theta}, \ell)} &= \frac{\prod_{j \neq \ell'} (1 - \theta_j)}{\sum_{k=0}^{n-2} \frac{1}{k+1} \sum_{S \subseteq A \setminus \{\ell, \ell'\}, |S|=k} \prod_{j \in S} \theta_j \prod_{j \in A \setminus S, j \neq \ell, \ell'} 1 - \theta_j} \\ &= \frac{1 - \theta_{\ell'}}{\sum_{k=0}^{n-2} \frac{1}{k+1} \sum_{S \subseteq A \setminus \{\ell, \ell'\}, |S|=k} \prod_{j \in S} \frac{\theta_j}{1 - \theta_j} \prod_{j \in A \setminus S, j \neq \ell, \ell'} \frac{1 - \theta_j}{1 - \theta_j}} \\ &= \frac{1 - \theta_{\ell'}}{\sum_{k=0}^{n-2} \frac{1}{k+1} \sum_{S \subseteq A \setminus \{\ell, \ell'\}, |S|=k} \prod_{j \in S} \frac{\theta_j}{1 - \theta_j}}. \end{aligned}$$

This function is indeed decreasing in θ_k , for all $k \in A$, and thus takes its highest value at $\vec{\theta} = (\underline{\theta}, \dots, \underline{\theta})$. Recalling the analysis of Section 4 when abilities are identical, we have that $\pi_{\ell\ell'}(\vec{\theta}) > 0$ for all $\vec{\theta} \in [\underline{\theta}, \bar{\theta}]$ and any two distinct ℓ, ℓ' in A , if and only if

$$\underline{\theta} > \underline{\theta}(n-1) \frac{(1 - \underline{\theta})^{n-1}}{1 - (1 - \underline{\theta})^{n-1}},$$

or equivalently, $\underline{\theta} > \theta^* = \sqrt[n-1]{\frac{1}{n}}$. ■

Lemma 4. The matrix $B_i(\vec{\theta})$ is strictly diagonally dominant if $\underline{\theta} > \theta^*$.

Proof. Showing strict diagonal dominance requires checking for every $\ell \neq i$ that

$$|[B_i(\vec{\theta})]_{\ell\ell}| > \sum_{\ell' \neq \ell} |[B_i(\vec{\theta})]_{\ell\ell'}|.$$

Observe that all the diagonal entries of $B_i(\vec{\theta})$ are positive, and that within any given row ℓ , all the off-diagonal entries have the same sign: positive if $\theta_i > \theta_{\ell}$, negative if $\theta_i < \theta_{\ell}$ and

zero if $\theta_i = \theta_\ell$. Using Remark 1 and the definition of $\pi_{\ell\ell'}(\vec{\theta})$, we have

$$\begin{aligned} \sum_{\ell' \neq \ell} [B_i(\vec{\theta})]_{\ell\ell'} &= \sum_{\ell' \neq \ell} \left(\theta_{\ell'} \sigma_{\ell'}(\vec{\theta}, i) - \theta_{\ell'} \sigma_{\ell'}(\vec{\theta}, \ell) \right) \\ &= \rho_\ell(\vec{\theta}) - \rho_i(\vec{\theta}) + \theta_i \sigma_i(\theta, \ell) - \theta_\ell \sigma_\ell(\theta, i) \\ &= \pi_{\ell i}(\vec{\theta}) - \pi_{i\ell}(\vec{\theta}). \end{aligned}$$

Thus, given that all off-diagonal elements have the same sign within a given row, we have to check that $|\pi_{\ell i}(\vec{\theta}) - \pi_{i\ell}(\vec{\theta})| < \pi_{i\ell}(\vec{\theta}) + 1 + \frac{1-\delta_i}{(1-\gamma)\delta_i}$. This is clearly true when $\theta_i = \theta_\ell$. Suppose now $\theta_\ell > \theta_i$. In this case, we have to check that

$$\pi_{i\ell}(\vec{\theta}) - \pi_{\ell i}(\vec{\theta}) < \pi_{i\ell}(\vec{\theta}) + 1 + \frac{1-\delta_i}{(1-\gamma)\delta_i}$$

which is satisfied since $\pi_{\ell i}(\vec{\theta}) > 0$ by Lemma 3. Suppose next that $\theta_\ell < \theta_i$. In this case, we have to check that

$$\pi_{\ell i}(\vec{\theta}) - 2\pi_{i\ell}(\vec{\theta}) < 1 + \frac{1-\delta_i}{(1-\gamma)\delta_i}.$$

The right-hand side of the inequality is lowest when $\delta_i = 1$. This holds since $\pi_{\ell i}(\vec{\theta}) < 1$ and $\pi_{i\ell}(\vec{\theta}) > 0$ by Lemma 3. ■

Proof of Proposition 4

The proof proceeds through a series of lemmas. As is well-known, if $A\vec{x} \gg 0$ for some M -matrix¹¹ A , then $\vec{x} \gg 0$, since M -matrices admit positive inverses. The matrices we encounter in the characterization of the $\Delta\vec{V}$'s are diagonally dominant, but not M -matrices since they may have rows with positive entries. In that case, $A\vec{x} \gg 0$ need not imply that $\vec{x} \gg 0$. However, the next lemma shows that if \vec{x} has a non-positive component, then there exists ℓ such that $x_\ell \leq 0$ and the ℓ^{th} row of A has only positive entries.

Lemma 5. *Let $A = (a_{ij})_{1 \leq i, j \leq n}$ be a square matrix that is strictly diagonally dominant, with positive elements on the diagonal, and such that any two off-diagonal entries in the same row have the same sign. If $A\vec{x} \gg 0$ and \vec{x} has a non-positive component, then there exists ℓ such that $x_\ell \leq 0$ and all the entries in row ℓ of A are positive.*

Proof. Suppose, on the contrary, that the set J of components j such that $x_j \leq 0$ is nonempty, and that row j of A has non-positive entries off-diagonal, for all $j \in J$. Let ℓ be the element of J such that $x_\ell \leq x_j$, for all $j \in J$.

¹¹An M -matrix has positive diagonal entries and negative off-diagonal entries.

We have

$$\sum_{k=1}^n a_{\ell k} x_k > 0,$$

which is equivalent to

$$\sum_{k \in J} a_{\ell k} x_k + \sum_{k \notin J} a_{\ell k} x_k > 0. \quad (13)$$

Notice that $a_{\ell k} x_k \leq -|a_{\ell k}| x_\ell$, for all $k \in J \setminus \{\ell\}$. Hence $\sum_{k \in J} a_{\ell k} x_k \leq (a_{\ell \ell} - \sum_{k \in J \setminus \{\ell\}} |a_{\ell k}|) x_\ell$, which is non-positive since the coefficient of x_ℓ is nonnegative by diagonal dominance of A . All the coefficients of the second term in (13) fall off the diagonal of A and are thus non-positive, while the corresponding components of x are positive since $k \notin J$. Hence the second term is non-positive, reaching a contradiction. ■

The next lemma states a well-known fact in algebra, which will prove useful to establish the subsequent lemma.

Lemma 6. *A square matrix that is strictly diagonally dominant and has positive entries on the diagonal admits a positive determinant.*

Proof. See, for instance, Carnicer et al (1999). ■

The linear system characterizing equilibrium payoff gains in Proposition 2 can be rewritten as

$$\frac{\delta_i(1-\gamma)}{u_i(1-\delta_i)} B_i(\vec{\theta}) \Delta \vec{V}_i(\vec{\theta}) = \vec{\pi}_i(\vec{\theta}),$$

where $\frac{\delta_i(1-\gamma)}{u_i(1-\delta_i)} \Delta \vec{V}_i(\vec{\theta})$ is the expression that appears in the equilibrium constraints. Thus we are interested in better understanding the system $B_i(\vec{\theta}) \vec{x} = \vec{\pi}_i(\vec{\theta})$. We start with the case $\delta_i = 1$, where the matrix $B_i(\vec{\theta})$ takes a simpler form. Let thus $B_i^*(\vec{\theta})$ be the $(n-1)$ -square matrix whose $\ell \ell'$ -entry is $1 + \pi_{i\ell}(\vec{\theta})$ when $\ell = \ell'$, and $\pi_{i\ell'}(\vec{\theta}) - \pi_{i\ell}(\vec{\theta})$ when $\ell \neq \ell'$.

Lemma 7. *Consider $n \geq 3$. Then, (i) defining $\sigma^* = \sigma_i(\theta^*, \dots, \theta^*, \ell)$ for some i, ℓ , we have $(1 - \theta^*)\sigma^* < 1/2$. Moreover, (ii) if $\theta_i \geq \theta_\ell$ and $\underline{\theta} > \theta^*$ then $\pi_{\ell i}(\vec{\theta}) - \pi_{i\ell}(\vec{\theta}) < 1/2$.*

Proof. For (i), first note that the definition of σ^* is independent of the choice of i, ℓ since σ is evaluated when all abilities are equal to θ^* . Then observe $(1 - \theta^*)\sigma^* < 1/2$ if and only if

$$\frac{2}{n} {}^{n-1}\sqrt{\frac{1}{n}} < 1 - {}^{n-1}\sqrt{\frac{1}{n}},$$

since, by construction, $\theta^* \sigma^* = \rho^* := \rho_i(\theta^*, \dots, \theta^*)$ and $\theta^* = 1 - {}^{n-1}\sqrt{\frac{1}{n}}$. The desired inequality is thus equivalent to

$$\frac{1}{n} < \left(\frac{n}{n+2} \right)^{n-1}.$$

The LHS above is decreasing in n , while the RHS is increasing in n , and the inequality is satisfied for $n = 3$. For (ii), note that $\theta_i \geq \theta_\ell$ implies that

$$\begin{aligned} \pi_{\ell i}(\vec{\theta}) - \pi_{i\ell}(\vec{\theta}) &= \theta_i \sigma_i(\vec{\theta}, \ell) - \rho_i(\vec{\theta}) - \theta_\ell \sigma_\ell(\vec{\theta}, i) + \rho_\ell(\vec{\theta}) \\ &= (\theta_i - \theta_\ell) \left(\sigma_i(\vec{\theta}, \ell) - \prod_{j \neq i, \ell} (1 - \theta_j) \right) \\ &\leq (\theta_i - \theta_\ell) \sigma_i(\vec{\theta}, \ell) \\ &\leq (1 - \theta^*) \sigma^*. \end{aligned}$$

The proof concludes by applying the inequality from part (i). ■

Lemma 8. *If $B_i^*(\vec{\theta})\vec{x} = \vec{\pi}_i(\theta)$ and $\underline{\theta} > \theta^*$, then $\vec{x} \gg 0$.*

Proof. This is evident in the case $n = 2$, since $B_i^*(\vec{\theta})$ is a positive scalar. The rest of the proof pertains to $n \geq 3$, where $B_i^*(\vec{\theta})$ is a matrix.

Suppose that \vec{x} has some non-positive component. By Lemma 5, we can pick j such that $x_j \leq 0$ and all entries of $B_i^*(\vec{\theta})$ are positive. By Cramer's rule, $x_j = \det(M_j) / \det(B_i^*(\vec{\theta}))$, where M_j is the matrix obtained by replacing the j -column of $B_i^*(\vec{\theta})$ by $\vec{\pi}_i(\theta)$. By Lemma 4 and Lemma 6, $B_i^*(\vec{\theta})$ has a positive determinant, since it is strictly diagonally dominant, and has positive entries on the diagonal. We finish the proof of this lemma by showing that M_j has a positive determinant as well, which contradicts $x_j \leq 0$. We do this by showing that M_j also satisfies the conditions of Lemma 6.

Consider first the case of row j , where the diagonal entry is $\pi_{ji}(\vec{\theta})$, while the off-diagonal entries come from $B_i^*(\vec{\theta})$, and are positive by the lemma. Thus, we have to show that $\pi_{ji}(\vec{\theta}) > \sum_{k \neq i, j} (\pi_{ik}(\vec{\theta}) - \pi_{jk}(\vec{\theta}))$. The sum on the right-hand side is equal to

$$\left(1 - \rho_i(\vec{\theta}) - \theta_j \sigma_j(\vec{\theta}, i)\right) - \left(1 - \rho_j(\vec{\theta}) - \theta_i \sigma_i(\vec{\theta}, j)\right) = \pi_{ji}(\vec{\theta}) - \pi_{ij}(\vec{\theta}).$$

This is indeed strictly less than $\pi_{ji}(\vec{\theta})$, since $\pi_{ij}(\vec{\theta}) > 0$ by Lemma 3.

Next, consider a row $\ell \neq j$, and suppose first that $\theta_i \leq \theta_\ell$. Then, all the off-diagonal entries on row ℓ of $B_i^*(\vec{\theta})$ are non-positive. A similar computation to that from the previous paragraph tells that the sum of the absolute value of the off-diagonal entries of $B_i^*(\vec{\theta})$ is $\pi_{i\ell}(\vec{\theta}) - \pi_{\ell i}(\vec{\theta})$. The difference between $B_i^*(\vec{\theta})$ and M_j on row ℓ is that the ℓj -entry of M_j is $\pi_{ji}(\vec{\theta})$ instead of $\pi_{ij}(\vec{\theta}) - \pi_{\ell j}(\vec{\theta})$. Thus, the sum of the absolute value of the off-diagonal entries of M_j is

$$\pi_{i\ell}(\vec{\theta}) - \pi_{\ell i}(\vec{\theta}) - \left(\pi_{\ell j}(\vec{\theta}) - \pi_{ij}(\vec{\theta})\right) + \pi_{ji}(\vec{\theta}),$$

which has to be strictly inferior to the diagonal entry $1 + \pi_{i\ell}(\vec{\theta})$, in view of Lemma 3 and

that $\pi_{\ell j}(\vec{\theta}) \geq \pi_{ij}(\vec{\theta})$ when $\theta_\ell \geq \theta_i$. This concludes the subcase for row $\ell \neq j$ when $\theta_i \leq \theta_\ell$.

Finally, consider the case of row ℓ where $\theta_i \geq \theta_\ell$. Similar computations tell us that strict diagonal dominance is satisfied on such a row if and only if

$$\pi_{ji}(\vec{\theta}) + \pi_{\ell i}(\vec{\theta}) + \pi_{\ell j}(\vec{\theta}) < 1 + 2\pi_{i\ell}(\vec{\theta}) + \pi_{ij}(\vec{\theta}),$$

which requires the following function of $\vec{\theta}$ to be everywhere positive on the domain $[\theta^*, 1]^A$:

$$1 + 2\pi_{i\ell}(\vec{\theta}) + \pi_{ij}(\vec{\theta}) - \pi_{ji}(\vec{\theta}) - \pi_{\ell i}(\vec{\theta}) - \pi_{\ell j}(\vec{\theta}). \quad (14)$$

Note that the function in Equation (14) is multi-linear in $\vec{\theta}$, as each $\pi_{\ell\ell'}$ is multi-linear in $\vec{\theta}$. Hence it is larger or equal to its minimal value when replacing each of $(\theta_i, \theta_j, \theta_\ell)$ with one of the two boundary points, 1 or θ^* . In the four cases where $\theta_i = \theta_\ell$, that is, $(\theta_i, \theta_j, \theta_\ell)$ is either $(1, 1, 1)$, $(\theta^*, \theta^*, \theta^*)$, $(1, \theta^*, 1)$ or $(\theta^*, \theta^*, \theta^*)$, the function in Equation (14) is positive since $\pi_{ij}(\vec{\theta}) = \pi_{\ell j}(\vec{\theta})$ and $\pi_{i\ell}(\vec{\theta}) = \pi_{\ell i}(\vec{\theta})$. In the two cases where $\theta_\ell > \theta_i$, that is, $(\theta_i, \theta_j, \theta_\ell)$ is either $(\theta^*, 1, 1)$ or $(\theta^*, \theta^*, 1)$, note that we have $\theta_j \geq \theta_i$. The desired positivity follows since $\pi_{ij}(\vec{\theta}) \geq \pi_{ji}(\vec{\theta})$ and $\pi_{i\ell}(\vec{\theta}) \geq \pi_{\ell i}(\vec{\theta})$. Finally, we consider the two cases where $\theta_\ell < \theta_i$, that is, $(\theta_i, \theta_j, \theta_\ell)$ is either $(1, 1, \theta^*)$ or $(1, \theta^*, \theta^*)$. Since in this case $\pi_{ij}(\vec{\theta}) \geq \pi_{ji}(\vec{\theta})$ and $\pi_{\ell j}(\vec{\theta}) \geq \pi_{\ell i}(\vec{\theta})$, it suffices to show $2(\pi_{\ell i}(\vec{\theta}) - \pi_{i\ell}(\vec{\theta})) < 1$. This follows from Lemma 7. ■

The next lemma shows that the right-hand side inequality in Proposition 3 can be uniformly satisfied with slack when $\delta_i = 1$ provided that profits are sufficiently informative of ideas' qualities.

Lemma 9. *Let $\underline{\theta} > \theta^*$. There exists an informativeness threshold τ and an $\eta > 0$ such that*

$$\vec{\sigma}_i(\vec{\theta}) + \eta \vec{1} \leq M_i^b(\vec{\theta}) B_i^*(\vec{\theta})^{-1} \vec{\pi}_i(\vec{\theta}),$$

for all $\vec{\theta} \in [\underline{\theta}, \bar{\theta}]^A$, and all (β, γ) with an informativeness likelihood ratio $\frac{1-\beta}{1-\gamma} \geq \tau$.

Proof. Notice that both $\vec{\sigma}_i(\vec{\theta})$ and $B_i^*(\vec{\theta})^{-1} \vec{\pi}_i(\vec{\theta})$ are independent of β and γ . The vector $\vec{\sigma}_i(\vec{\theta}) \ll \vec{1}$ for all θ , and there exists $\epsilon > 0$ such that $B_i^*(\vec{\theta})^{-1} \vec{\pi}_i(\vec{\theta}) \geq \epsilon \vec{1}$ by Lemma 8 (the expression being continuous in $\vec{\theta}$, its infimum for $\vec{\theta} \in [\underline{\theta}, \bar{\theta}]^A$ will be attained on the domain). Finally, the matrix $M_i^b(\vec{\theta})$ has diagonal entries larger than $\frac{\gamma-\beta}{1-\gamma} \min_{\vec{\theta} \in [\underline{\theta}, \bar{\theta}]^A} \min_{\ell \neq i} \sigma_i(\vec{\theta}, \ell)$, and off-diagonal entries in $[0, 1]$, for all $\vec{\theta} \in [\underline{\theta}, \bar{\theta}]^A$. The result then follows after observing that $\min_{\vec{\theta} \in [\underline{\theta}, \bar{\theta}]^A} \min_{\ell \neq i} \sigma_i(\vec{\theta}, \ell) > 0$, and that $\frac{\gamma-\beta}{1-\gamma} = \frac{1-\beta}{1-\gamma} - 1$. ■

We continue our exploration of the system $B_i(\vec{\theta}) \vec{x} = \vec{\pi}_i(\theta)$. Lemma 8 focused on the case $\delta_i = 1$. The next lemma shows that, for any sequence of discount factors that converges to

1, the corresponding solutions to the system uniformly (in $\vec{\theta}$) converge to the solution when $\delta_i = 1$. To emphasize the dependence of B_i on δ_i , we write $B_i(\vec{\theta}, \delta_i)$ in the next Lemma.

Lemma 10. *Let $\vec{x}(\vec{\theta}, \delta_i) = B_i(\vec{\theta}, \delta_i)^{-1}\vec{\pi}_i(\vec{\theta})$. Consider a sequence $(\delta_i^m)_{m \geq 1}$ of positive discount factors that converge to 1. Then for all $\epsilon > 0$, there exists an integer \bar{m} such that for all $m \geq \bar{m}$ and all $\vec{\theta} \in [\underline{\theta}, \bar{\theta}]^A$, we have $\|\vec{x}(\vec{\theta}, \delta_i^m) - \vec{x}(\vec{\theta}, 1)\| \leq \epsilon$.*

Proof. Notice that $B_i(\vec{\theta}, \delta_i) = B_i^*(\vec{\theta}) + \frac{1-\delta_i}{(1-\gamma)\delta_i} Id$. So B_i is differentiable with respect to δ_i , and its derivative is equal to $\frac{-1}{(1-\gamma)\delta_i^2}$. Hence $B_i(\vec{\theta}, \delta_i)^{-1}$ is differentiable with respect to δ_i , and its derivative is equal to $\frac{-1}{(1-\gamma)\delta_i^2} B_i(\vec{\theta}, \delta_i)^{-1} B_i(\vec{\theta}, \delta_i)^{-1}$. We have:

$$\|\vec{x}(\vec{\theta}, 1) - \vec{x}(\vec{\theta}, \delta_i)\|_\infty \leq |1 - \delta_i| \sup_{\delta \in [\delta_i, 1]} \left\| \frac{\partial \vec{x}}{\partial \delta}(\vec{\theta}, \delta) \right\|_\infty \leq |1 - \delta_i| \sup_{\delta \in [\delta_i, 1]} \frac{\|B_i(\vec{\theta}, \delta)^{-1}\|_\infty^2}{(1-\gamma)\delta^2} \|\vec{\pi}_i(\vec{\theta})\|_\infty.$$

Thus,

$$\sup_{\vec{\theta} \in [\underline{\theta}, \bar{\theta}]^A} \|\vec{x}(\vec{\theta}, 1) - \vec{x}(\vec{\theta}, \delta_i)\|_\infty \leq |1 - \delta_i| \sup_{\delta \in [\delta_i, 1], \vec{\theta} \in [\underline{\theta}, \bar{\theta}]^A} \frac{\|B_i(\vec{\theta}, \delta)^{-1}\|_\infty^2}{(1-\gamma)\delta^2}.$$

The result then follows from the fact that the supremum on the right-hand side of the inequality is bounded. ■

We are now ready to prove Proposition 4. The negative result when $\underline{\theta} < \theta^*$ follows from Proposition 1, since achieving the principal's first best in a belief-free way is more demanding than achieving it for commonly known symmetric vectors of abilities. The positive result when $\underline{\theta} > \theta^*$ is proved by showing that the Silent Treatment strategy profile constitutes a belief-free PBNE in the right range of parameters. The following lemma pertains to discerning agents not reporting bad ideas, while the next focuses on discerning agents reporting good ideas.

Lemma 11. *There exists an informativeness threshold $\tau > 0$ such that, for all (β, γ) with $\frac{1-\beta}{1-\gamma} \geq \tau$, there exists $\underline{\delta} \in (0, 1)$ so that the inequality*

$$\vec{\sigma}_i(\vec{\theta}) \leq M_i^b(\vec{\theta}) B_i(\vec{\theta}, \delta_i)^{-1} \vec{\pi}_i(\vec{\theta}),$$

is satisfied for all $\delta \geq \underline{\delta}$ and all $\vec{\theta} \in [\underline{\theta}, \bar{\theta}]^A$.

Proof. Pick τ as in Lemma 9, and let (β, γ) be such that $\frac{1-\beta}{1-\gamma} \geq \tau$. If this lemma were not true, then there would exist a sequence $(\delta_i^m)_{m \geq 1}$ in $(0, 1)$ that converges to 1 and a sequence $(\vec{\theta}^m)_{m \geq 1}$ in $[\underline{\theta}, \bar{\theta}]^A$ that converges to some $\vec{\theta} \in [\underline{\theta}, \bar{\theta}]^A$ such that

$$\vec{\sigma}_i(\vec{\theta}^m) > M_i^b(\vec{\theta}^m) B_i(\vec{\theta}^m, \delta_i^m)^{-1} \vec{\pi}_i(\vec{\theta}^m),$$

for all m . By Lemma 10, for each $\epsilon > 0$, there exists \bar{m} such that

$$\bar{\sigma}_i(\theta^{\bar{m}}) + \epsilon \|M_i^b(\theta^{\bar{m}})\|_\infty \bar{\mathbf{1}} > M_i^b(\theta^{\bar{m}}) B_i^*(\theta^{\bar{m}})^{-1} \bar{\pi}_i(\theta^{\bar{m}}),$$

for all $m \geq \bar{m}$. This contradicts Lemma 9, as seen when taking m to infinity while selecting $\epsilon < \eta / \max_{\bar{\theta} \in [\underline{\theta}, \bar{\theta}]^A} \|M_i^b(\bar{\theta})\|_\infty$. ■

We now conclude the proof of Proposition 4 by checking that the Silent Treatment strategy profile forms a belief-free PBNE when δ_i is close enough to one, which holds independently of the informativeness level $\frac{1-\beta}{1-\gamma}$.

Lemma 12. *Suppose that $\underline{\theta} > \theta^*$. Then there exists a $\underline{\delta} < 1$ such that the inequality*

$$M_i^g(\bar{\theta}) B_i(\bar{\theta}, \delta_i)^{-1} \bar{\pi}_i(\bar{\theta}) \leq \bar{\sigma}_i(\bar{\theta}),$$

is satisfied for all $\delta \geq \underline{\delta}$ and all $\bar{\theta} \in [\underline{\theta}, \bar{\theta}]^A$.

Proof. The inequality to be checked can be rewritten as

$$\hat{M}_i^g(\bar{\theta}) B_i(\bar{\theta}, \delta_i)^{-1} \bar{\pi}_i(\bar{\theta}) \leq \bar{\mathbf{1}},$$

where $\hat{M}_i^g(\bar{\theta})$ be the matrix obtained by dividing each row ℓ of $M_i^g(\bar{\theta})$ by $\sigma_i(\bar{\theta}, \ell)$.

We now show that the inequality holds, and uniformly with slack, when $\delta_i = 1$, that is, there exists $\eta > 0$ such that

$$\hat{M}_i^g(\bar{\theta}) B_i^*(\bar{\theta})^{-1} \bar{\pi}_i(\bar{\theta}) \leq (1 - \eta) \bar{\mathbf{1}}, \quad (15)$$

for all $\bar{\theta} \in [\underline{\theta}, \bar{\theta}]^A$. The expression being continuous in $\bar{\theta}$ and $[\underline{\theta}, \bar{\theta}]^A$ being compact, it is sufficient to check that for all $\bar{\theta} \in [\underline{\theta}, \bar{\theta}]^A$,

$$\|\hat{M}_i^g(\bar{\theta}) B_i^*(\bar{\theta})^{-1} \bar{\pi}_i(\bar{\theta})\|_\infty < 1.$$

To see this, note that we have:

$$\begin{aligned} \|\hat{M}_i^g(\bar{\theta}) B_i^*(\bar{\theta})^{-1} \bar{\pi}_i(\bar{\theta})\|_\infty &\leq \|\hat{M}_i^g(\bar{\theta})\|_\infty \|B_i^*(\bar{\theta})^{-1}\|_\infty \|\bar{\pi}_i(\bar{\theta})\|_\infty \\ &= \|B_i^*(\bar{\theta})^{-1}\|_\infty \|\bar{\pi}_i(\bar{\theta})\|_\infty \\ &\leq \frac{\max_{k \neq i} \pi_{ki}(\bar{\theta})}{\min_{\ell \neq i} 1 + \pi_{i\ell}(\bar{\theta}) - \sum_{j \neq i, \ell} |\pi_{ij}(\bar{\theta}) - \pi_{\ell j}(\bar{\theta})|}, \end{aligned}$$

where the first inequality follows from properties of matrix norms, the equality follows from

the fact that elements on each row of $\hat{M}_i^g(\vec{\theta})$ sum up to 1, and the second inequality follows from the definition of $\vec{\pi}_i(\theta)$ and the Ahlberg-Nilson-Varah bound (see e.g. Varah (1975)) since $B_i^*(\theta)$ is diagonally dominant. In the case that $n = 2$, we immediately see that the right-hand side is strictly smaller than one, as $\pi_{i\ell}(\vec{\theta}), \pi_{\ell i}(\vec{\theta}) \in (0, 1)$ and the summation over $j \neq i, \ell$ vanishes. The rest of the proof considers the case $n \geq 3$.

Let k be the index at which the maximum in the numerator of the last expression is reached and let ℓ be the index at which the minimum in the denominator is reached. Suppose first that $\theta_\ell \geq \theta_i$. Then the denominator becomes $1 + \pi_{\ell i}(\vec{\theta})$ (see similar computations when showing that $B_i(\vec{\theta})$ is diagonally dominant). The fact that $\pi_{ki}(\vec{\theta})$ is strictly inferior follows at once since it is strictly inferior to 1 and $\pi_{\ell i}(\vec{\theta})$ is nonnegative (which follows from $\theta_i \geq \theta^*$, for all i). Suppose now that $\theta_\ell \leq \theta_i$. Now the denominator in the last expression becomes $1 - \pi_{\ell i}(\vec{\theta}) + 2\pi_{i\ell}(\vec{\theta})$ (see similar computations when showing that $B_i(\vec{\theta})$ is diagonally dominant), and we have to show that $\pi_{ki}(\vec{\theta})$ is strictly inferior to it. For this, suppose first that $k = \ell$. The inequality to check becomes $\pi_{\ell i}(\vec{\theta}) - \pi_{i\ell}(\vec{\theta}) < 1/2$, which was already shown in Lemma 7. Finally, suppose that $k \neq \ell$. We need to check that

$$\theta_i \sigma_i(\vec{\theta}, k) - \theta_\ell \sigma_i(\vec{\theta}, \ell) - 2(\theta_i - \theta_\ell) \Pi_{j \neq i, \ell} (1 - \theta_j) + (\theta_i - \theta_\ell) \sigma_i(\vec{\theta}, \ell) < 1.$$

It is sufficient to check that $\theta_i \sigma_i(\vec{\theta}, k) - \theta_\ell \sigma_i(\vec{\theta}, \ell) + (\theta_i - \theta_\ell) \sigma_i(\vec{\theta}, \ell) < 1$. Notice that the expression on the LHS is increasing in θ_i . Thus it is no larger than the same expression evaluated at $\theta_i = 1$. Next, it is linear in θ_ℓ , and it is thus maximized by taking $\theta_\ell = 1$ or θ^* . The inequality is obvious if $\theta_\ell = 1$, so let's assume that $\theta_\ell = \theta^*$. Thus it is sufficient to prove that $\sigma_i((\theta^*, \vec{\theta}_{-\ell}), k) - \theta^* \sigma_i(\vec{\theta}, \ell) + (1 - 2\theta^*) \sigma_i(\vec{\theta}, \ell) < 1$. Remember that θ^* is less than $1/2$ when $n \geq 2$. Thus the coefficient of $\pi_i(\theta_{-\ell})$ is positive, and the expression on the LHS is lower or equal to $(2 - 2\theta^*)\pi^*$. The desired inequality then follows from Lemma 7 for $n \geq 3$.

Now that equation (15) has been proved, we are ready to conclude the proof of this lemma. If the result is not true, then there exists a sequence $(\delta_i^m)_{m \geq 1}$ in $(0, 1)$ that converges to 1 and a sequence $(\vec{\theta}^m)_{m \geq 1}$ in $[\underline{\theta}, \bar{\theta}]^A$ that converges to some $\vec{\theta} \in [\underline{\theta}, \bar{\theta}]^A$ such that

$$M_i^g(\vec{\theta}^m) B_i(\vec{\theta}^m, \delta_i^m)^{-1} \vec{\pi}_i(\vec{\theta}^m) > \vec{\sigma}_i(\vec{\theta}^m),$$

for all m . By Lemma 10, for each $\epsilon > 0$, there exists \bar{m} such that

$$M_i^g(\vec{\theta}^m) B_i^*(\vec{\theta}^m)^{-1} \vec{\pi}_i(\vec{\theta}^m) + \epsilon \|M_i^g(\vec{\theta}^m)\|_\infty \vec{1} > \vec{\sigma}_i(\vec{\theta}^m),$$

for all $m \geq \bar{m}$. This contradicts Equation (15), as seen when taking m to infinity while selecting $\epsilon < \eta / \max_{\vec{\theta} \in [\underline{\theta}, \bar{\theta}]^A} \|M_i^g(\vec{\theta})\|_\infty$. ■

Proof of Proposition 5

The incentive constraint for an agent i to propose a good idea when $\ell \neq i$ is the agent of last resort is:

$$\sigma_i(\vec{\theta}, \ell)(1 - \delta_i)u_i + \delta_i V_i^D(\vec{\theta}, \ell) \geq \delta_i V_i^D(\vec{\theta}, \ell),$$

which is thus always satisfied. The incentive constraint for an agent i not to propose a bad idea when $\ell \neq i$ is the agent of last resort is:

$$\delta_i V_i^D(\vec{\theta}, \ell) \geq \sigma_i(\vec{\theta}, \ell) \left((1 - \delta_i)u_i + \beta \delta_i V_i^D(\vec{\theta}, \ell) + (1 - \beta) \delta_i V_i^{LR}(\vec{\theta}) \right) + (1 - \sigma_i(\vec{\theta}, \ell)) \delta_i V_i^D(\vec{\theta}, \ell).$$

Subtracting $\delta_i V_i^{LR}(\vec{\theta})$ on both sides, and rearranging terms, we get:

$$\Delta V_i(\vec{\theta}, \ell) \geq \frac{u_i(1 - \delta_i)}{\delta_i(1 - \beta)} \quad (16)$$

The average discounted payoffs $V_i^D(\vec{\theta}, \ell)$ and $V_i^{LR}(\vec{\theta})$ are given by the following recursive equations, as can also be seen by using $\gamma = 1$ in (7) and (8):

$$\begin{aligned} V_i^D(\vec{\theta}, \ell) &= \theta_i \sigma_i(\vec{\theta}, \ell)(1 - \delta_i)u_i + \delta_i V_i^D(\vec{\theta}, \ell), \\ V_i^{LR}(\vec{\theta}) &= \rho_i(\vec{\theta})(1 - \delta_i)u_i + \delta_i V_i^{LR}(\vec{\theta}), \end{aligned}$$

which gives

$$\Delta V_i(\vec{\theta}, \ell) = \pi_{\ell i}(\vec{\theta})u_i.$$

Thus, refraining from reporting bad ideas when ℓ is the agent of last resort imposes the following restriction:

$$\pi_{\ell i}(\vec{\theta}) \geq \frac{(1 - \delta_i)}{\delta_i(1 - \beta)},$$

or equivalently,

$$\delta_i \geq \frac{1}{1 + (1 - \beta)\pi_{\ell i}(\vec{\theta})}.$$

Varying the agent of last resort, the highest lower-bound on δ_i is obtained when $\pi_{\ell i}$ is the smallest, which gives us the desired inequality.

Thus, the silent-treatment strategy forms a belief-free PBNE if and only if $\pi_i(\vec{\theta}) > 0$ for all i and all $\vec{\theta} \in [\underline{\theta}, \bar{\theta}]^n$. The last part of this proposition then follows from Lemma 3. ■

Proof of Proposition 6

The proof is analogous to that given for Proposition 1(i) in the text.

References

- Andrews, I. and D. Barron**, 2016. “The Allocation of Future Business.” *American Economic Review*, forthcoming.
- Bar-Isaac, H.**, 2003. “Reputation and Survival: Learning in a Dynamic Signalling Model” *Review of Economic Studies* **70**, 231–251.
- Bergemann, D. and J. Välimäki**, 2008. “Bandit Problems,” in *The New Palgrave Dictionary of Economics*, 336–340.
- Board, S.**, 2011. “Relational Contracts and the Value of Loyalty.” *American Economic Review* **101**, 3349–3367.
- Carnicer, J.M., T.N.T. Goodman and J.M. Peña**, 1999. “Linear Conditions for Positive Determinants.” *Linear Algebra and Its Applications* **292**, 39–59.
- Dessein, W., A. Galeotti and T. Santos**, 2015. “Rational Inattention and Organizational Focus.” *American Economic Review*, forthcoming.
- Ely, J. and J. Välimäki**, 2002. “A Robust Folk Theorem for the Prisoner’s Dilemma.” *Journal of Economic Theory* **102**, 84–105.
- Garicano, L. and A. Prat**, 2013. “Organizational Economics with Cognitive Costs,” in *Advances in Economic Theory*, Econometric Society 10th World Congress, ed. by D. Acemoglu, M. Arellano, and E. Dekel, Cambridge University Press.
- Ghosh, A. and P. Hummel**, 2013. “Learning and Incentives in User-Generated Content: Multi-Armed Bandits with Endogenous Arms,” in *Proc. 4th Conference on Innovations in Theoretical Computer Science (ITCS)*.
- Gittins, J. and D. Jones**, 1974. “A Dynamic Allocation Index for the Sequential Allocation of Experiments,” in *Progress in Statistics*, Econometric Society 10th World Congress, ed. by J. Gani, 241–266, North Holland, Amsterdam.
- Halac, M. and A. Prat**, 2016. “Managerial Attention and Worker Performance.” *American Economic Review*, forthcoming.
- Hörner, J., S. Lovo and T. Tomala**, 2011. “Belief-Free Equilibria in Games with Incomplete Information: Characterization and Existence.” *Journal of Economic Theory* **146**, 1770–1795.
- Li, J., N. Matouscheck and M. Powell**, 2016. “Power Dynamics in Organizations.” *American Economic Journal: Microeconomics*, forthcoming.
- Lipnowski, E. and J. Ramos**, 2015. “Repeated Delegation.” *Working Paper*.
- Piccione, M.**, 2002. “The Repeated Prisoners’ Dilemma with Imperfect Private Monitoring.” *Journal of Economic Theory* **102**, 70–83.

Varah, J.M., 1975. "A Lower Bound for the Smallest Singular Value of a Matrix." *Linear Algebra and Its Applications* **11**, 3-5.