

The Silent Treatment*

Geoffroy de Clippel
Brown University

Kfir Eliaz
Tel-Aviv University

Kareen Rozen[†]
Brown University

August 2017

Abstract

The attention of a principal is sought by multiple agents eager to have their ideas implemented. Such attention-seeking imposes an externality on the principal, who can only implement one idea per period and may overlook valuable proposals. Her only means of providing incentives is her rule for selecting among proposals, which must be rational for her to follow. Can she design an idea-selection mechanism that circumvents this problem? This paper argues that in repeated interactions, the principal can ensure agents refrain from communicating ideas unless they are of the highest quality. The principal may achieve her first-best outcome even when she is fully attention constrained and has minimal information about the ability of agents. Whether her first best is achievable hinges on the the worst possible agent.

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

*This paper benefitted from seminar audiences at the Barcelona GSE Summer Forum, Brown University, Hebrew University, New York University, Pennsylvania State University, the University of Wisconsin-Madison and the Zurich Workshop on Applied Behavioral Theory. Special thanks to Nageeb Ali, Heski Bar-Isaac, Johannes Hörner, Phil Reny, Larry Samuelson and Lones Smith for their helpful suggestions. Geoffroy de Clippel and Kfir Eliaz thank the BSF for financial support under award no. 003624.

[†]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

Both in their personal life and at work, most people repeatedly find themselves engaged in trying to meet other people’s demands for their attention, or demanding others’ attention themselves. We send emails and texts expecting a prompt response, while being ourselves flooded with incoming communications requiring immediate attention. In organizations, 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. News outlets get contaminated with articles based on noncredible sources (‘fake news’), an issue affecting both the readers choosing among outlets and editors considering which reporter to assign a story.

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? Can the decision maker achieve her first best outcome: finding good ideas without expending any costly attention? Stated differently, can the decision maker construct an idea-selection mechanism to find the best ideas even when she has no attention left to assess them at all?

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. We use the term ‘idea’ in a generic sense: depending on the setting, an idea could be new information or a new source of information, a project, a program, etc. 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 . We consider a principal who is fully attention constrained: in each period, she is unable to review any of the proposals to infer their quality before choosing one to implement. 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

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

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 agent i is characterized by an ability $\theta_i \in [\underline{\theta}, 1]$, 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.

The paper is organized as follows. Section 2 discusses related 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 shows that all the results can be adapted to the limiting case in which good ideas surely yield high profit. Section 7 extends the analysis to general profit distributions, showing how the first best can still be achieved through a variation of the Silent Treatment

²Piccione (2002) and Ely and Välimäki (2002) introduce belief-freeness in the repeated prisoner’s dilemma to generate robustness against the private-monitoring structure, carefully constructing mixed strategies where each player is indifferent between all his actions, no matter his opponent’s private history. Hörner, Lovo and Tomala (2011) consider dynamic games of incomplete information, examining existence of belief-free equilibrium payoffs under certain information structures and reward functions. They also use sophisticated strategies involving randomization to delimit what is achievable. By contrast, 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 first best is impossible to attain in a belief-free way.

with an optimally chosen set of profit levels that trigger punishment. Section 8 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.³ 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.⁴

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 preclude such contracts, as in most public organizations where subordinates may suggest ideas and improvements to an executive decision-maker. It is also a feature of political environments, where voters (principals, or principal in the case of a median voter) elect one of multiple candidates (agents) to an office. The literature on infinitely repeated elections in which candidates have privately known types has remained small, according to Duggan and Martinelli (forthcoming)'s survey of dynamic voting models, due to the "difficult theoretical issues related to updating of voter beliefs." They note that this small literature

³Bergemann 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.

⁴Some 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).

has examined restrictions to simplify belief updating.⁵ In our own work, belief updating is entirely unnecessary when the Silent Treatment strategy profile constitutes an equilibrium. There are also structural differences between our framework and this literature. Banks and Sundaram (1993, 1998), for instance, include moral hazard (the candidate takes an action after election) and model private information as being persistent. By contrast, we have two levels of adverse selection: the agent’s underlying ability, which is persistent, and the quality of the agent’s idea, which varies over time.

The recent game-theoretic literature on dynamic mechanism design with neither transfers nor commitment includes Lipnowski and Ramos (2016) and Li, Matouschek and Powell (2017). Both study an infinitely repeated game between a principal and a single agent. While using different game structures, both of their models have the feature that the principal decides whether to entrust a task to the agent, who is better informed. They each predict different and interesting non-stationary dynamics in equilibrium. Among other 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 non-stationary 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 Board’s 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

⁵Schwabe (2011) exogenously fixes some types’ behavior. Banks and Sundaram (1993,1998) analyze a restricted class of equilibria under assumptions that mitigate the belief-updating problem: their 1993 work assumes an infinite pool of candidates with known type distribution (so that a candidate, once taken out of office, would never be elected again), and their 1998 work assumes two-period term limits.

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 between the latter two papers and our own. 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.

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 in each period $t = 1, 2, 3, \dots$. An idea’s profit to the principal is either high (H) or low (L), where $H > L \geq 0$. (Section 7 extends the analysis to non-binary profit levels.) The principal’s profit depends stochastically on the quality of the idea that she implements. An idea’s quality is either *good* or *bad*.⁶ 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 .⁷ There is a commonly known probability $\theta \in (0, 1)$ that an agent’s idea in any given period is good.

⁶For example, a reporter may or may not have a good network of credible sources for some matter; a proposed fix to a software issue may or may not have bugs itself; a researcher’s results may or may not be robust; a paid ‘expert’ may or may not be knowledgeable about a specific question; a candidate for a task may or may not be experienced in it.

⁷Note that β may be zero. We discuss the case $\gamma = 1$ in Section 6.

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 a stage of reviewing proposals, along with a cost function for doing so, 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.⁸ 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 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

⁸Our 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.

period, a good idea whenever one exists and a bad idea otherwise.

Our model accommodates, in a stylized way, situations with the following features. A decision-maker repeatedly receives proposals to implement from either subordinates, consultants or independent contributors. It is too costly or infeasible for the decision maker to fully analyze each of the proposed ideas (e.g., the decision maker may operate under tight deadlines). There is a conflict of interest in that the decision maker wants to choose a “home run”, while the proposers of ideas want their idea to be selected. There is no institutional device that enables the decision maker to credibly commit to a policy, and the proposers’ payoff cannot be made contingent on the payoff to the decision maker. This could be due to the fact that the decision maker’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 preclude such contracts (as in most public organizations where subordinates may suggest ideas and improvements to an executive decision maker).

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 a 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.⁹ We say that the principal’s first best is *achievable in 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 in equilibrium.*

(ii) *If $\theta > \theta^*$, then the principal’s first best is achievable in equilibrium as soon as profits are sufficiently informative of quality (in particular, if $\frac{1-\beta}{1-\gamma} > 1 + \frac{n-1}{1-n(1-\theta)^{n-1}}$).*

⁹Of 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.

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 $i(h)$.

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 in 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 shortly define 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. ■

Two comments on this result are in order. First, as can be seen from the proof of part (i), the inability to reach the first best when $\theta < \theta^*$ holds not just for PBNE, but any Nash equilibrium. Second, the threshold ability θ^* in Proposition 1 is monotone in n , the number of agents. Indeed, 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, for any number of agents n , if we have $\theta > \theta^*$, then the principal's first best can be achieved through the following strategy profile.

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.

The principal achieves her first best if she and all the agents follow the Silent Treatment strategy profile. She 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. Does this strategy profile, however, constitute an equilibrium? To address this question, we 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 may now 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 ,*

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

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

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

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

$$\underbrace{\sigma}_{i \text{ selected}} \left((1 - \delta_i)u_i + \underbrace{\gamma \delta_i V_i^D}_{\text{high profit}} + \underbrace{(1 - \gamma)\delta_i V_i^{LR}}_{\text{low profit, } i \text{ becomes last resort agent}} \right) + \underbrace{(1 - \sigma)\delta_i V_i^D}_{i \text{ not selected}} \geq \delta_i V_i^D. \quad (3)$$

Let us first examine incentive condition (2). We subtract $\delta_i V_i^{LR}$ from both sides of the inequality (2), and let Δ_i represent $V_i^D - V_i^{LR}$. Then incentive condition (2) 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

$$\Delta_i \geq \frac{(1 - \delta_i)u_i}{(1 - \beta)\delta_i}. \quad (4)$$

Similar computations show that inequality (3) is equivalent to

$$\Delta_i \leq \frac{(1 - \delta_i)u_i}{(1 - \gamma)\delta_i}. \quad (5)$$

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:

$$\begin{aligned} V_i^D &= \overbrace{\theta\sigma}^{\text{good idea, selected}} \left((1 - \delta_i)u_i + \gamma\delta_i V_i^D + (1 - \gamma)\delta_i V_i^{LR} \right) + \overbrace{(1 - \theta\sigma)}^{\text{bad idea, or not selected}} \delta_i V_i^D, \\ V_i^{LR} &= \underbrace{\rho}_{\text{selected}} \left((1 - \delta_i)u_i + \delta_i V_i^{LR} \right) + \underbrace{(1 - \rho)}_{\text{not selected}} \left(\gamma\delta_i V_i^{LR} + \underbrace{(1 - \gamma)\delta_i V_i^D}_{\text{low profit, switch to discerning}} \right). \end{aligned} \quad (6)$$

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 (6), 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 (5) for proposing good ideas is always satisfied, and that the incentive condition (4) 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

$$\pi > \frac{1 - \gamma}{\gamma - \beta} = \frac{1}{\tau - 1}, \quad (7)$$

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 $[\underline{\theta}, 1]$. 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.

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.

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. 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 equilibrium depends on their ability. A discerning agent’s payoff is thus also 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 varies over time as low profits are realized.

We 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. Agents' payoffs and incentives depend on the probability with which an agent's idea is selected, assuming that all others 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 if 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, the incentive condition for a discerning agent i *not to propose a bad idea* when ℓ is

the agent of last resort, 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 a discerning agent i to *propose a good idea* when ℓ is the agent of last resort, is:

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

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. Moreover, it turns out that IC_b and IC_g 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 lemma 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 agents other than i and ℓ all 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*

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

Equilibrium Payoffs and Payoff Differences. Lemma 1 provides a preliminary characterization of the equilibrium conditions as a function of the average discounted payoff differences. We now aim to characterize these 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 $\ell \neq i$:

$$\begin{aligned} V_i^{LR}(\vec{\theta}) &= \overbrace{\rho_i(\vec{\theta})}^{i \text{ is chosen}} \left((1-\delta_i)u_i + \delta_i V_i^{LR}(\vec{\theta}) \right) + \sum_{j \neq i} \overbrace{\theta_j \sigma_j(\vec{\theta}, i)}^{j \text{ chosen when } i \text{ is last resort}} \left(\gamma \delta_i V_i^{LR}(\vec{\theta}) + \overbrace{(1-\gamma)\delta_i V_i^D(\vec{\theta}, j)}^{\text{low profit, } j \text{ becomes last resort}} \right), \\ V_i^D(\vec{\theta}, \ell) &= \overbrace{\theta_i \sigma_i(\vec{\theta}, \ell)}^{i \text{ chosen when } \ell \text{ is last resort}} \left((1-\delta_i)u_i + \gamma \delta_i V_i^D(\vec{\theta}, \ell) + \overbrace{(1-\gamma)\delta_i V_i^{LR}(\vec{\theta})}^{\text{low profit, } i \text{ becomes last resort}} \right) \\ &\quad + \sum_{j \neq i, \ell} \overbrace{\theta_j \sigma_j(\vec{\theta}, \ell)}^{j \text{ chosen when } \ell \text{ is last resort}} \left(\gamma \delta_i V_i^D(\vec{\theta}, \ell) + \overbrace{(1-\gamma)\delta_i V_i^D(\vec{\theta}, j)}^{\text{low profit, } j \text{ becomes last resort}} \right) + \underbrace{\rho_\ell(\vec{\theta})}_{\ell \text{ is chosen}} \delta_i V_i^D(\vec{\theta}, \ell). \end{aligned} \tag{8}$$

In the Appendix, we manipulate the system of equations (8) to derive the average discounted payoff differences. As it turns out, 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. In 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 knew that the matrix $B_i(\vec{\theta})$ were invertible, we could solve for $\Delta\vec{V}_i(\vec{\theta})$ using Lemma 2 and then use Lemma 1 to characterize when the Silent Treatment forms an equilibrium of the ex-post game with abilities $\vec{\theta}$. In the Appendix, we establish that when $\underline{\theta} > \theta^*$, the matrix $B_i(\vec{\theta})$ has a special property that implies invertibility. Namely, the matrix is strictly diagonally dominant: for every row, the absolute value of the diagonal element is strictly larger than the sum of the absolute values of the off-diagonal elements.

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

$$M_i^g(\vec{\theta})B_i(\vec{\theta})^{-1}\vec{\pi}_i(\vec{\theta}) \leq \vec{\sigma}_i(\vec{\theta}) \leq M_i^b(\vec{\theta})B_i(\vec{\theta})^{-1}\vec{\pi}_i(\vec{\theta}). \quad (9)$$

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 $[\underline{\theta}, 1]^A$ 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 $i \in \mathcal{A}$.

The entries of the matrix $B_i^{-1}(\vec{\theta})$ involve complicated expressions for most vectors $\vec{\theta}$. One may thus expect the equilibrium conditions to be much more complex than (1) for the homogenous case. A surprisingly simple characterization result emerges, however, for the belief-free equilibrium. In fact, the equilibrium condition is the same as (1) provided π is replaced by the minimal probability premium $\pi_{\ell i}(\vec{\theta})$ for agent i when considering all possible ability levels and last resort agents. As formally shown in the Appendix, and discussed further below, the minimal probability premium is the following function of $\underline{\theta}$:

$$\underline{\pi} = \begin{cases} \frac{\underline{\theta}}{n-1} & \text{if } n \geq 3 \text{ and } \underline{\theta} \geq 1 - \sqrt[n-2]{\frac{1}{n}} \\ \frac{1-n(1-\underline{\theta})^{n-1}}{n-1} & \text{otherwise.} \end{cases} \quad (10)$$

This characterization allows us to derive the agents' minimal discount factor that sustains the Silent Treatment as a belief-free equilibrium.

Proposition 3. *The Silent Treatment forms a belief-free equilibrium if and only if for each agent i ,*

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

where $\underline{\pi}$ is positive if and only if $\underline{\theta} > \theta^*$.

Remember that $\frac{1-n(1-\underline{\theta})^{n-1}}{n-1}$ corresponds to the probability premium in the homogenous case where all agents have an ability $\underline{\theta}$. Thus, by Proposition 3, the set of discount factors sustaining the Silent Treatment as a belief-free equilibrium for ability profiles in $[\underline{\theta}, 1]^A$ is the *same* set that sustains it as an equilibrium with homogenous abilities known to be $\underline{\theta}$ when there are two agents or $\underline{\theta}$ falls below $1 - \sqrt[n-2]{\frac{1}{n}}$. Otherwise, the range of discount factors supporting the belief-free equilibrium is *smaller* than in the case where the agents are

commonly known to be $\underline{\theta}$. Why is this so? In view of Proposition 3, to answer this question we need to understand at which profile of abilities the probability premium is minimized. Agent i 's probability premium $\pi_{\ell i}(\vec{\theta})$ is increasing in both θ_i and θ_ℓ , so it is minimized by setting both equal to $\underline{\theta}$. On the other hand, the abilities of discerning agents other than i have two opposing effects on $\pi_{\ell i}(\vec{\theta})$. When these discerning agents have higher abilities, they reduce the probability $\sigma_i(\vec{\theta}, \ell)$ that i is selected when he proposes (which lowers the premium), but they also reduce the probability $\rho_i(\vec{\theta})$ that i is picked when he is the agent of last resort (which raises the premium). The effect associated to $\sigma_i(\vec{\theta}, \ell)$ becomes relatively more important as $\underline{\theta}$ grows because $\sigma_i(\vec{\theta}, \ell)$ is premultiplied by $\theta_i = \underline{\theta}$ in the definition of the probability premium, while $\rho_i(\vec{\theta})$ is independent of θ_i . Thus the ability vector minimizing the probability premium has all agents with ability $\underline{\theta}$ when it is relatively low, but involves some high-quality opponents otherwise.

The proof of Proposition 3, available in the Appendix, is significantly more challenging than the proof of Proposition 1. Complications arise from the fact that all possible combinations of abilities must be considered, and that inverting $B_i(\vec{\theta})$ is far from trivial with heterogenous abilities. Fortunately, Lemma 3 shows that the equilibrium conditions depend directly on the vector $B_i(\vec{\theta})^{-1}\vec{\pi}_i(\vec{\theta})$. That vector can be shown to satisfy the relationship

$$B_i(\vec{\theta})^{-1}\vec{\pi}_i(\vec{\theta}) = [Id - \frac{1 - \delta_i\gamma}{\delta_i(1 - \gamma)}B_i(\vec{\theta})^{-1}]\vec{1},$$

because the sum over any row ℓ of the matrix $B_i(\vec{\theta})$ is equal to $1 + \frac{1 - \delta_i}{\delta_i(1 - \gamma)} + \pi_{\ell i}(\vec{\theta})$. This reduces the problem at hand to understanding the vector $B_i(\vec{\theta})^{-1}\vec{1}$, that is, the vector of row sums of $B_i(\vec{\theta})^{-1}$. Next, a power series development of $B_i(\vec{\theta})^{-1}$ establishes that $B_i(\vec{\theta})^{-1}\vec{1}$ is decreasing in θ_i , or that $B_i(\vec{\theta})^{-1}\vec{\pi}_i(\vec{\theta})$ is increasing in θ_i . Since M_i^b is a positive matrix, the equilibrium constraint for discerning agents not to report bad ideas is most challenging when $\theta_i = \underline{\theta}$. After observing that the matrix $B_i(\vec{\theta})$ is an M-matrix¹⁰ in that case, we can apply the Ahlberg-Nilson-Varah bound to provide a sharp upper-bound the row sums of $B_i(\vec{\theta})^{-1}$. Some algebra then establishes that a discerning agent does not want to report bad ideas when his discount factor is above the bound stated in Proposition 3. Similar techniques establish that discerning agents always want to report good ideas, independently of their discount factors. As for necessity in Proposition 3, we can just look at the equilibrium conditions stated in Lemma 3 for the ability vector that achieves $\underline{\pi}$. Although abilities are heterogenous when $\underline{\theta}$ is higher than $1 - \sqrt[n-2]{\frac{1}{n}}$, the matrix $B_i(\vec{\theta})$ remains easy to invert in that case because agents other than i are all symmetric.

Propositions 3 and part (i) of Proposition 1 together imply the following result.

¹⁰I.e., a strictly diagonally dominant matrix with positive diagonal entries and negative off-diagonal entries.

Corollary 1. *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 for all (β, γ) with $\frac{1-\beta}{1-\gamma} \geq \frac{1+\pi}{\pi}$, 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.

6 When Good Ideas Give High Profit for Sure

So far, the principal has been unable to definitively conclude that an idea was bad from observing low profit. In this section, we consider the special case of $\gamma = 1$. Thus good ideas surely deliver high profit. 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.

When $\gamma = 1$ the principal's first-best can be achieved by 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.¹¹ 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. However, the threat of severe (off-path) punishment may enable the principal to achieve his first-best in a belief-free equilibrium even when the lowest possible ability falls below θ^* .

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

¹¹Such 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.

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 $i \in A$.

With this definition, following the same reasoning as for proving Proposition 1(i) and Corollary 1(i) tells us that the principal's first best cannot be achieved by a belief-free equilibrium in a credible way when $\underline{\theta} < \theta^*$. One can also show that the Silent Treatment forms a belief-free equilibrium if and only if $\delta_i \geq 1/(1 + (1 - \beta)\underline{\pi})$, meaning that Proposition 3 remains true when $\gamma = 1$. In fact, computations becomes a bit easier in the case $\gamma = 1$ as the matrix $(1 - \gamma)B(\vec{\theta})$ (which determines $\Delta\vec{V}(\vec{\theta})$, see Lemma 2) is diagonal, and thus can easily be inverted.

One might conjecture that a more economical way to prove Proposition 1(ii) and Corollary 1(ii) would be to first establish the results for $\gamma = 1$, and then use some continuity argument to extend them to cases where γ close to one. After all, the inequalities that make the Silent Treatment an equilibrium can be satisfied strictly if agents are patient enough when $\gamma = 1$. One caveat of this approach, though, is that we wouldn't know how informative profits need to be for the Silent Treatment to form an equilibrium when agents are patient enough, nor would we know what patience threshold is required for the Silent Treatment to be an equilibrium for given β and γ . More importantly, the limit of (average or plain) discounted payoff gains as the discount factor goes to one is discontinuous at $\gamma = 1$. Consider the simpler case of homogenous ability (as in Section 4). As follows from our computations in the proof of Proposition 2, the average discounted payoff gain (ΔV) converges to 0 when γ is a given number less than one, and to πu_i when $\gamma = 1$. The limit of the plain discounted payoff gain display a discontinuity at $\gamma = 1$ as well ($\frac{\Delta V}{1 - \delta_i}$ converges to $\frac{\pi u_i}{(1 - \gamma)(1 + \pi)}$ when γ is a given number less than one, and to ∞ when $\gamma = 1$). There is a simple intuition for this discontinuity. When $\gamma < 1$, being discerning or last resort today has no impact in the long run on that agent's future position (as discerning or last resort) when implementing the Silent Treatment. The average discounted gain thus vanishes as patience increases, while the plain discounted gain converges to a finite number. By contrast, when $\gamma = 1$, agents remain in their position of last resort or discerning forever when implementing the Silent Treatment. The average discounted gain converges to a finite number as patience increases, while the plain discounted gain goes to infinity.

7 When to Punish, with General Profit Distributions

Our simple model supposes that there are only two possible profit levels. If the payoff to the principal can take multiple values, it is less clear for which profit levels an agent should be punished. Consider the following extension of our model. In any period t , the principal's profit y_t is drawn from $[\underline{y}, \bar{y}]$ according to a cumulative distribution function G (resp., B) when the idea is good (resp., bad). We include the possibility that $\underline{y} < 0$. We only require that the expected profit from a bad idea is positive, and strictly lower than the expected profit from a good idea. The ranking of expected profits from good and bad ideas lends meaning to the terms 'good' and 'bad', but otherwise imposes no restriction on the nature of the distributions. Our framework includes, among others, environments where good ideas first-order stochastically dominate bad ones, or where good ideas have a higher variance in profit. There is also no restriction on the presence of atoms. The Silent Treatment strategy profile may be adapted to this setting by endogenizing β and γ . A discerning agent still proposes his idea if and only if it is good, and the agent of last resort still proposes any idea. The principal chooses an idea uniformly among the set of proposals from discerning agents, if any exist; and otherwise chooses the idea of the agent of last resort. The main difference is that a discerning agent becomes the new agent of last resort when his idea generates a profit in some *punishment set* $Y \subset [\underline{y}, \bar{y}]$ which has positive measure according to B . Define $P_G(Y) = \int_{y \in Y} dG(y)$ and $P_B(Y) = \int_{y \in Y} dB(y)$ to be the probability that good and bad ideas, respectively, yield a profit in Y . In analogy to our earlier analysis, let $\gamma^* = 1 - P_G(Y)$ be the probability that a good idea generates a payoff outside the punishment set, and let $\beta^* = 1 - P_B(Y)$ be the probability that a bad idea does so. Observe that if this adjusted Silent Treatment strategy profile is an equilibrium, then the principal still obtains her first best. How should Y be chosen to sustain the equilibrium, if at all possible?

Consider the model with uncertain abilities. As can be seen from Proposition 3, the punishment set Y must be more likely under a bad idea than a good one (i.e., $\gamma^* > \beta^*$), else the incentive conditions would be impossible to satisfy. Furthermore, Proposition 3 states that the adjusted Silent Treatment is a belief-free equilibrium if and only if for all agents i ,

$$\delta_i \geq \frac{1}{\gamma^* + (\gamma^* - \beta^*)\underline{\pi}} = \frac{1}{1 - P_G(Y) + (P_B(Y) - P_G(Y))\underline{\pi}}, \quad (11)$$

where the minimal probability premium $\underline{\pi}$ is defined just as before. The adjusted Silent Treatment is thus a belief-free equilibrium for sufficiently patient agents when the denominator is strictly larger than one, or equivalently:

$$\underline{\pi} > \frac{P_G(Y)}{P_B(Y) - P_G(Y)} = \frac{1}{P_B(Y)/P_G(Y) - 1}. \quad (12)$$

The smallest $\underline{\theta}$ for which this is possible is obtained by picking Y to maximize $P_B(Y)/P_G(Y)$. If there exists a profit level y that is in the support of B but not of G , then $P_B(Y)/P_G(Y)$ is clearly maximized (and equal to infinity) by setting $Y = [y - \varepsilon, y + \varepsilon]$ for small enough ε . Such is the case in Section 6, where only bad ideas can generate a low profit.

What happens when bad ideas cannot be identified with certainty (i.e., the support of B is contained in the support of G)? Suppose, for instance, that B and G have continuous densities b and g satisfying the monotone likelihood ratio property, with $b(y)/g(y)$ decreasing in y .¹² Assuming \underline{y} is indeed in the support of b , the maximum of $P_B(Y)/P_G(Y)$ can be shown to be $\lim_{y \rightarrow \underline{y}} b(y)/g(y)$.¹³ Hence belief-free equilibrium can be sustained for a given $\underline{\theta}$ if $\underline{\pi} > 1/(\lim_{y \rightarrow \underline{y}} b(y)/g(y) - 1)$. If the likelihood ratio goes to infinity as y decreases to \underline{y} , then for any $\underline{\theta} > \theta^*$, one can find a y^* low enough to guarantee that the adjusted Silent Treatment with $Y = [y, y^*]$ forms an equilibrium for sufficiently patient agents.

Now that we understand, for given distributions b and g satisfying the monotone likelihood ratio property, starting from which $\underline{\theta}$ the adjusted Silent Treatment can be a belief-free equilibrium for sufficiently patient agents, we may also wish to answer another question. Given a particular $\underline{\theta}$, what is the largest possible range of discount factors for which the Silent Treatment is sustained as a belief-free equilibrium? In view of condition (11), this can be found by choosing the punishment set Y to maximize the objective

$$-P_G(Y) + [P_B(Y) - P_G(Y)]\underline{\pi} = \int_{y \in Y} (b(y)\underline{\pi} - g(y)(1 - \underline{\pi})) dy.$$

Thus, to have the widest possible range of discount factors, a profit level y should be included in the punishment set if and only if

$$\frac{b(y)}{g(y)} \geq \frac{1 + \underline{\pi}}{\underline{\pi}}. \quad (13)$$

Given the monotone likelihood ratio property, the punishment set ensuring equilibrium for the widest possible range of discount factors will take the form of an interval $Y = [y, y^*]$,

¹²The case of probability mass functions b, g satisfying the monotone likelihood ratio property is similar.

¹³For $Y = [y, y]$, we have $\lim_{y \rightarrow \underline{y}} P_B(Y)/P_G(Y) = \lim_{y \rightarrow \underline{y}} B(y)/G(y) = \lim_{y \rightarrow \underline{y}} b(y)/g(y)$ by l'Hôpital's rule. Moreover, for any other Y with positive measure under B (and thus G , by the inclusion of the support),

$$\frac{P_B(Y)}{P_G(Y)} = \frac{\int_{y \in Y} b(y) dy}{\int_{y \in Y} g(y) dy} = \frac{\int_{y \in Y} \frac{b(y)}{g(y)} g(y) dy}{\int_{y \in Y} g(y) dy} \leq \lim_{y \rightarrow \underline{y}} \frac{b(y)}{g(y)}.$$

where y^* satisfies condition (13) with equality.

The same reasoning as above shows that we can also extend to settings where an idea's outcome may be judged through lenses other than profit (an invention, a work of art, a research article) and may depend on the principal's perception. The principal may have gradations in her assessment of the outcome, but it matters only how she pools those differing assessments into 'high outcome' and 'low outcome' categories to determine when to trigger punishment of a discerning agent. Her perception of outcomes need only ensure a sufficiently high informativeness ratio to sustain the Silent Treatment as an equilibrium. In such settings, the probability distribution of the principal's possible assessments, conditional on idea quality, must be common knowledge. The principal's assessment itself, however, need not be observed by agents. It suffices to allow the principal the opportunity to publicly announce the next agent of last resort, as she has an incentive to speak truthfully.

8 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. 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.

Our work focuses on the first-best payoff to the principal, and examines when it can be achieved in equilibrium. This is an important benchmark, which allows us to understand the constraints imposed by incomplete information, imperfect monitoring and incentive compatibility. Our paper may be viewed 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. We leave it to future research to address interesting follow-up questions, such as what is the best the principal can do when her first best is unachievable.

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

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

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

$$\begin{aligned} V_i^D(\vec{\theta}, \ell) &= \theta_i \sigma_i(\vec{\theta}, \ell) \left((1-\delta_i)u_i + \gamma \delta_i V_i^D(\vec{\theta}, \ell) + (1-\gamma) \delta_i V_i^{LR}(\vec{\theta}) \right) \\ &+ \sum_{j \neq i, \ell} \theta_j \sigma_j(\vec{\theta}, \ell) \left(\gamma \delta_i V_i^D(\vec{\theta}, \ell) + (1-\gamma) \delta_i V_i^D(\vec{\theta}, j) \right) + \rho_\ell(\vec{\theta}) \delta_i V_i^D(\vec{\theta}, \ell), \end{aligned} \tag{14}$$

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

$$\begin{aligned} V_i^{LR}(\vec{\theta}) &= \rho_i(\vec{\theta}) \left((1 - \delta_i)u_i + \delta_i V_i^{LR}(\vec{\theta}) \right) \\ &+ \sum_{j \neq i} \theta_j \sigma_j(\vec{\theta}, i) \left(\gamma \delta_i V_i^{LR}(\vec{\theta}) + (1 - \gamma) \delta_i V_i^D(\vec{\theta}, j) \right). \end{aligned} \quad (15)$$

Subtracting $\delta_i V_i^{LR}(\vec{\theta})$ from both sides of Equation (14), we find that

$$\begin{aligned} V_i^D(\vec{\theta}, \ell) - \delta_i V_i^{LR}(\vec{\theta}) &= \theta_i \sigma_i(\vec{\theta}, \ell) \left((1 - \delta_i)u_i + \gamma \delta_i \Delta V_i^D(\vec{\theta}, \ell) \right) \\ &+ \sum_{j \neq i, \ell} \theta_j \sigma_j(\vec{\theta}, \ell) \left(\gamma \delta_i \Delta V_i^D(\vec{\theta}, \ell) + (1 - \gamma) \delta_i \Delta V_i^D(\vec{\theta}, j) \right) + \rho_\ell(\vec{\theta}) \delta_i \Delta V_i^D(\vec{\theta}, \ell), \end{aligned} \quad (16)$$

In view of Remark 1, Equation (14) simplifies to

$$\begin{aligned} V_i^D(\vec{\theta}, \ell) - \delta_i V_i^{LR}(\vec{\theta}) &= \theta_i \sigma_i(\vec{\theta}, \ell) (1 - \delta_i)u_i + (1 - \gamma) \delta_i \sum_{j \neq i, \ell} \theta_j \sigma_j(\vec{\theta}, \ell) \Delta V_i^D(\vec{\theta}, j) \\ &+ \delta_i \Delta V_i^D(\vec{\theta}, \ell) \left(\gamma + (1 - \gamma) \rho_\ell(\vec{\theta}) \right). \end{aligned} \quad (17)$$

Similarly, subtracting $\delta_i V_i^{LR}(\vec{\theta})$ from both sides of Equation (15), we find that

$$V_i^{LR}(\vec{\theta}) - \delta_i V_i^{LR}(\vec{\theta}) = \rho_i(\vec{\theta}) (1 - \delta_i)u_i + (1 - \gamma) \delta_i \sum_{j \neq i} \theta_j \sigma_j(\vec{\theta}, i) \Delta V_i^D(\vec{\theta}, j). \quad (18)$$

Subtracting Equation (18) from Equation (17), and using the definition of $\pi_{\ell\ell'}(\vec{\theta})$, we find that:

$$\begin{aligned} \Delta V_i^D(\vec{\theta}, \ell) &= \pi_{\ell i}(\vec{\theta}) (1 - \delta_i)u_i + \delta_i \Delta V_i^D(\vec{\theta}, \ell) \left(\gamma - (1 - \gamma) \pi_{i\ell}(\vec{\theta}) \right) \\ &+ (1 - \gamma) \delta_i \sum_{j \neq i, \ell} \left(\theta_j \sigma_j(\vec{\theta}, \ell) - \theta_j \sigma_j(\vec{\theta}, i) \right) \Delta V_i^D(\vec{\theta}, j). \end{aligned} \quad (19)$$

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 (19) 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. ■

C. Proof of Proposition 3

We start by establishing properties of selection probabilities and probability premiums. We let $\sigma^* = \sigma_i(\theta^*, \dots, \theta^*, \ell)$ for any $i \neq \ell$ (the selection probability does not vary on i and ℓ when all agents have the same ability).

Lemma 4. (a) For each agent $\ell \neq \ell'$, $\frac{\rho_{\ell'}(\vec{\theta})}{\sigma_{\ell'}(\vec{\theta}, \ell)}$ is decreasing in θ_k , for all $k \in \mathcal{A}$.

(b) $\pi_{\ell\ell'}(\vec{\theta}) > 0$ for all $\vec{\theta} \in [\underline{\theta}, 1]$ and any $\ell \neq \ell'$ in \mathcal{A} , if and only if $\underline{\theta} > \theta^*$.

(c) $(1 - \theta^*)\sigma^* \leq 1/2$.

(d) Suppose $\underline{\theta} > \theta^*$, and $\ell \neq i$ is such that $\theta_\ell \leq \theta_i$. Then $\pi_{\ell i}(\vec{\theta}) - \pi_{i\ell}(\vec{\theta}) \leq 1/2$.

(e) The minimal probability premium $\underline{\pi} := \min_{\ell \in \mathcal{A} \setminus \{i\}} \min_{\vec{\theta} \in [\underline{\theta}, 1]^n} \pi_{\ell i}(\vec{\theta})$ is given by

$$\underline{\pi} = \begin{cases} \frac{\underline{\theta}}{n-1} & \text{if } n \geq 3 \text{ and } \underline{\theta} \geq 1 - \sqrt[n-2]{\frac{1}{n}} \\ \frac{1-n(1-\underline{\theta})^{n-1}}{n-1} & \text{otherwise.} \end{cases}$$

Proof. (a) Observe that

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

This function is indeed decreasing in θ_k , for all $k \in \mathcal{A}$.

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

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

From (a), the expression on the right-hand side 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 $\vec{\theta} \in [\underline{\theta}, 1]$ and any two distinct ℓ, ℓ' in \mathcal{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, $\underline{\theta} > \theta^* = \sqrt[n-1]{\frac{1}{n}}$.

(c) 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 $(1 - \theta^*)\sigma^* \leq 1/2$ if and only if

$$\frac{2}{n} \sqrt[n-1]{\frac{1}{n}} \leq 1 - \sqrt[n-1]{\frac{1}{n}},$$

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

$$1 \leq \frac{n^n}{(n+2)^{n-1}} \quad (20)$$

Taking natural logs on both sides, and adding and subtracting $\ln(n+2)$, (20) is equivalent to

$$n(\ln n - \ln(n+2)) + \ln(n+2) \geq 0. \quad (21)$$

The inequality in (20), and thus (21), is satisfied for $n \in \{2, 3, 4\}$ (i.e., $1 \geq 1$, $27/25 \geq 1$ and $256/216 \geq 1$ respectively), and we now show it holds for all larger n by proving that the derivative of the LHS of (21) is positive for all $n \geq 4$. Indeed, that derivative is

$$\frac{3}{n+2} + \ln n - \ln(n+2) > \frac{3}{n+2} - \frac{2}{n} = \frac{n-4}{n(n+2)},$$

where the inequality follows using strict concavity of $\ln n$, so that $\frac{\ln(n+2) - \ln n}{2} < \frac{d}{dn} \ln n = \frac{1}{n}$.

(d) 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 (c).

(e) Notice that $\pi_{\ell i}(\vec{\theta})$ is increasing in θ_i , so that one should take $\theta_i = \underline{\theta}$ to find the minimum. If $n = 2$, then the minimum is reached by taking $\theta_{-i} = \underline{\theta}$ as well. Suppose $n \geq 3$. The expression $\pi_{\ell i}(\vec{\theta})$ is linear in θ_k for all $k \neq i, \ell$. Thus one need only consider the cases $\theta_k \in \{\underline{\theta}, 1\}$ for all k . Notice, however, that $\rho_i(\vec{\theta}) = 0$ as soon as one such $\theta_k = 1$, in which case $\pi_{\ell i}(\vec{\theta})$ is decreasing in θ_j for $j \neq i, \ell, k$, and independent of θ_ℓ . In addition, if $\theta_k = \underline{\theta}$ for all $k \neq i, \ell$, then $\pi_{\ell i}(\vec{\theta})$ is strictly increasing in θ_ℓ and the minimum will be reached at $\theta_\ell = \underline{\theta}$. To summarize, the minimal $\pi_{\ell i}(\vec{\theta})$ is reached at a profile $\vec{\theta}$ where $\theta_i = \underline{\theta}$, and other agents' abilities are either all $\underline{\theta}$ or all 1. The probability premium is¹⁴ $\frac{1-n(1-\underline{\theta})^{n-1}}{n-1}$ in the former case, and $\frac{\underline{\theta}}{n-1}$ in the latter case. It is then easy to check that the former expression is smaller than the latter if and only if $\underline{\theta} \leq 1 - \sqrt[n-2]{\frac{1}{n}}$ (which is larger than θ^*). ■

Now we establish some properties of the matrix $B_i(\vec{\theta})$.

¹⁴Indeed, agents other than i are symmetric and the fact that one must be chosen implies $(n-1)\underline{\theta}\sigma_i(\underline{\theta}, \ell) + \rho_\ell(\underline{\theta}) = 1$, or $\underline{\theta}\sigma_i(\underline{\theta}, \ell) = \frac{1-(1-\underline{\theta})^{n-1}}{n-1}$.

Lemma 5. (a) $B_i(\vec{\theta})\vec{1} = \frac{1-\delta_i\gamma}{\delta_i(1-\gamma)}\vec{1} + \vec{\pi}_i(\vec{\theta})$.

(b) Diagonal entries of $B_i(\vec{\theta})$ are positive. Off-diagonal entries are positive on any row ℓ such that $\theta_i > \theta_\ell$, negative on any row ℓ such that $\theta_i < \theta_\ell$, and zero on any row ℓ such that $\theta_i = \theta_\ell$.

(c) For each $\ell \neq i$, let z_ℓ be the difference between row ℓ 's diagonal entry and the sum of the absolute value of its off-diagonal entries:

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

If $\theta_i \leq \theta_\ell$, then $z_\ell = \frac{1-\delta_i\gamma}{\delta_i(1-\gamma)} + \pi_{\ell i}$. If $\theta_i \geq \theta_\ell$, then $z_\ell = \frac{1-\delta_i\gamma}{\delta_i(1-\gamma)} + 2\pi_{i\ell} - \pi_{\ell i}$.

(d) $B_i(\vec{\theta})$ is (row) strictly diagonally dominant, and thus invertible.

(e) $\|B_i(\vec{\theta})^{-1}\|_\infty \leq \frac{1}{\min_{\ell \neq i} z_\ell}$.

(f) $B_i(\vec{\theta})^{-1}\vec{\pi}_i(\vec{\theta}) = [Id - \frac{1-\delta_i\gamma}{\delta_i(1-\gamma)}B_i(\vec{\theta})^{-1}]\vec{1}$.

(g) $B_i(\vec{\theta})^{-1} = \sum_{k=0}^{\infty} (-1)^k (\theta_i - \theta^*)^k (X_i^{-1}Y_i)^k X_i^{-1}$, where X_i be the matrix $B_i(\vec{\theta})$ evaluated at $\theta_i = \theta^*$, and Y_i is the positive matrix whose $\ell\ell'$ -entry is $\frac{\rho_\ell(\vec{\theta})}{1-\theta_i}$ if $\ell = \ell'$, and $-\theta_{\ell'} \frac{d\sigma_{\ell'}}{d\theta_i}(\vec{\theta}, \ell)$ if $\ell \neq \ell'$.

(h) Each component of the vector $B_i(\vec{\theta})^{-1}\vec{\pi}_i(\vec{\theta})$ is increasing in θ_i , and each component of the vector $B_i(\vec{\theta})^{-1}\mathbb{1}$ is decreasing in θ_i , for $\theta_i \in [\theta^*, 1]$.

Proof. (a) Notice that

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

Thus the sum over the columns of the entries of $B_i(\vec{\theta})$ appearing on row ℓ is equal to

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

Thus $B_i(\vec{\theta})\vec{1} = \frac{1-\delta_i\gamma}{\delta_i(1-\gamma)}\vec{1} + \vec{\pi}_i(\vec{\theta})$, as desired.

(b) The fact that diagonal entries are positive is obvious. Off-diagonal entries on row ℓ are of the form $\pi_{i\ell'}(\vec{\theta}) - \pi_{\ell\ell'}(\vec{\theta})$, which is equal to $\theta_{\ell'} (\sigma_{\ell'}(\vec{\theta}, i) - \sigma_{\ell'}(\vec{\theta}, \ell))$. The result about the

sign of off-diagonal entries then follows as the likelihood for a discerning ℓ' to be picked diminishes when part of a better pool of discerning agents.

(c) By (b), off-diagonal entries on a row ℓ are non-positive when $\theta_i \leq \theta_\ell$, in which case z_ℓ is simply the sum of the elements appearing on row ℓ , whose value is given in (a). Suppose now $\theta_i \geq \theta_\ell$. The first computation in the proof of (a) shows that the sum of the off-diagonal elements on row ℓ (which are all positive, by (b)) is equal to $\pi_{\ell i}(\vec{\theta}) - \pi_{i\ell}(\vec{\theta})$. Thus $z_\ell = \frac{1-\delta_i\gamma}{\delta_i(1-\gamma)} + \pi_{i\ell} - (\pi_{\ell i}(\vec{\theta}) - \pi_{i\ell}(\vec{\theta}))$, and the result follows.

(d) We just need to check that $z_\ell > 0$ for all ℓ . Since $\frac{1-\delta_i\gamma}{\delta_i(1-\gamma)} > 1$, the result follows from the fact that $\pi_{\ell i} \geq 0$ for the case $\theta_\ell \geq \theta_i$, and from the fact that $\pi_{i\ell} \geq 0$ and $\pi_{\ell i} < 1$ for the case $\theta_\ell \leq \theta_i$.

(e) This follows from the Ahlberg-Nilson-Varah bound (see e.g. Varah (1975)) since $B_i(\vec{\theta})$ is strictly diagonally dominant.

(f) Since $B_i(\vec{\theta})$ is invertible by (e), the identity follows from (a) by premultiplying both sides of the equality by $B_i(\vec{\theta})^{-1}$.

(g) Notice that the entries of $B_i(\vec{\theta})$ are affine functions of θ_i . Indeed, the matrix Y_i is obtained by taking the derivative with respect to θ_i of the entries of $B_i(\vec{\theta})$, and is independent of θ_i . Thus

$$B_i(\vec{\theta}) = X_i + (\theta_i - \theta^*)Y_i.$$

The result then follows from the power series expansion of matrix inverses, after showing that $\|X_i^{-1}Y_i\|_\infty < 1$. To check this, first notice that $\|X_i^{-1}\|_\infty < 1$ by (e) given that $\theta_\ell \geq \theta^*$ for all $\ell \neq i$. Consider Y_i next. It is a positive matrix, so its infinite norm is obtained by computing for each row the sum of its entries, and then taking the maximum of these sums over the rows. Observed that Y_i is the derivative with respect to θ_i of the matrix $B_i(\vec{\theta})$. Using the computations from (a), the sum of the elements on row ℓ of Y_i is simply the derivative with respect to θ_i of $\pi_{\ell i}(\vec{\theta})$, which is equal to $\sigma_i(\vec{\theta}, \ell)$. This expression is decreasing in $\vec{\theta}$ for each ℓ , and thus lower or equal to σ^* , which is less than 1. Then $\|X_i^{-1}Y_i\|_\infty \leq \|X_i^{-1}\|_\infty \|Y_i\|_\infty < \sigma^* < 1$, as desired.

(h) By (a), the derivative of $B_i(\vec{\theta})^{-1}\vec{\pi}(\vec{\theta})$ with respect to θ_i is equal to the opposite of the derivative of $B_i(\vec{\theta})^{-1}\vec{1}$, which by (g) is equal to

$$\sum_{k=1}^{\infty} (-1)^{k+1} k (\theta_i - \theta^*)^{k-1} (X_i^{-1}Y_i)^k X_i^{-1}\vec{1}.$$

Notice that $2(\theta_i - \theta^*)Y_iX_i^{-1}\vec{1} \leq 2(1 - \theta^*)Y_iX_i^{-1}\vec{1} < \vec{1}$. The first inequality follows from the facts that Y_i and X_i^{-1} (inverse of an M -matrix) are positive, and $\theta_i \leq 1$. The strict

inequality follows from (c) in Lemma 4, since each component of the vector $Y_i X_i^{-1} \vec{1}$ is lower or equal to $\|Y_i X_i^{-1}\|_\infty$, which is strictly less than σ^* (see (g) above).

Being a product of positive matrices, the matrix $X_i^{-1} Y_i X_i^{-1}$ is positive. Hence we know $X_i^{-1} Y_i X_i^{-1} \vec{1} - 2(\theta_i - \theta^*)(X_i^{-1} Y_i)^2 X_i^{-1} \vec{1}$ is a strictly positive vector. This corresponds to the first two terms in the above expression for the derivative of $B_i(\vec{\theta})^{-1} \vec{\pi}(\vec{\theta})$ with respect to θ_i .

A fortiori, $3X_i^{-1} Y_i X_i^{-1} \vec{1} - 4(\theta_i - \theta^*)(X_i^{-1} Y_i)^2 X_i^{-1} \vec{1}$ is a strictly positive vector, and hence $3(\theta_i - \theta^*)^2 (X_i^{-1} Y_i)^3 X_i^{-1} \vec{1} - 4(\theta_i - \theta^*)^3 (X_i^{-1} Y_i)^4 X_i^{-1} \vec{1}$ is a strictly positive vector as well (since $(\theta_i - \theta^*)^2 (X_i^{-1} Y_i)^2$ is a positive matrix). This corresponds to the next two terms in the above expression for the derivative of $B_i(\vec{\theta})^{-1} \vec{\pi}(\vec{\theta})$ with respect to θ_i . Iterating the argument this way, we conclude that this derivative is strictly positive. ■

Lemma 6. *Discerning agents are always willing to propose good ideas.*

Proof. Remember that discerning agents report good ideas if and only if

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

To establish this inequality, it is sufficient to show that $\|B_i(\vec{\theta})^{-1} \vec{\pi}_i(\vec{\theta})\|_\infty \leq 1$, since $M_i^g(\vec{\theta})$ is a positive matrix with the sum of entries on any row ℓ equal to $\sigma_i(\vec{\theta}, \ell)$. It is sufficient to establish the upper-bound on the infinite norm under the assumption that $\theta_i = 1$, because of (h) in Lemma 5. Using $\|B_i(\vec{\theta})^{-1} \vec{\pi}_i(\vec{\theta})\|_\infty \leq \|B_i(\vec{\theta})^{-1}\|_\infty \|\vec{\pi}_i(\vec{\theta})\|_\infty$, combined with (c) and (e) from Lemma 5, it is sufficient to check that

$$\pi_{ki}(\vec{\theta}) < \frac{1 - \delta_i \gamma}{\delta_i (1 - \gamma)} - \pi_{\ell i}(\vec{\theta}) + 2\pi_{i\ell}(\vec{\theta}), \quad (22)$$

where k is an agent $j \neq i$ that maximizes $\pi_{ji}(\vec{\theta})$ and ℓ is an agent $j \neq i$ that minimizes $2\pi_{ij}(\vec{\theta}) - \pi_{ji}(\vec{\theta})$. Inequality (22) holds when $k = \ell$, since $\pi_{\ell i}(\vec{\theta}) - \pi_{i\ell}(\vec{\theta}) \leq 1/2$ by (c) in Lemma 4, and $\frac{1 - \delta_i \gamma}{\delta_i (1 - \gamma)} > 1$. Suppose then that $k \neq \ell$. Inequality (22) becomes (remember $\theta_i = 1$)

$$\sigma_i(\vec{\theta}, k) - \theta_\ell \sigma_i(\vec{\theta}, \ell) - 2\rho_i(\vec{\theta}) + (1 - \theta_\ell) \sigma_i(\vec{\theta}, \ell) < \frac{1 - \delta_i \gamma}{\delta_i (1 - \gamma)}.$$

It is sufficient to check that $\sigma_i(\vec{\theta}, k) - \theta_\ell \sigma_i(\vec{\theta}, \ell) + (1 - \theta_\ell) \sigma_i(\vec{\theta}, \ell) \leq 1$. Notice that the expression on the LHS 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 $\sigma_i((\theta^*, \vec{\theta}_{-\ell}), k) - \theta^* \sigma_i(\vec{\theta}, \ell) + (1 - \theta^*) \sigma_i(\vec{\theta}, \ell) \leq 1$. Remember that θ^* is less than $1/2$ when $n \geq 2$, so the total weight on $\sigma_i(\vec{\theta}, \ell)$ is positive. The expression on the LHS is thus lower or equal to $(2 - 2\theta^*) \sigma^*$. The desired inequality then follows from (c) in Lemma 4. ■

Lemma 7. *Discerning agents are willing to withhold bad ideas if $\delta_i \geq \frac{1}{\gamma + (\gamma - \beta)\underline{\pi}}$.*

Proof. Remember that discerning agents withhold bad ideas if and only if

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

By (f) from Lemma 5, this is equivalent to

$$\vec{\sigma}_i(\vec{\theta}) + \frac{1 - \delta_i\gamma}{\delta_i(1 - \gamma)}M_i^b(\vec{\theta})B_i(\vec{\theta})^{-1}\vec{1} \leq M_i^b(\vec{\theta})\mathbb{1} = \vec{\sigma}_i(\vec{\theta}) + \frac{\gamma - \beta}{1 - \gamma}\vec{\sigma}_i(\vec{\theta}),$$

or

$$\frac{1 - \delta_i\gamma}{\delta_i(1 - \gamma)}M_i^b(\vec{\theta})B_i(\vec{\theta})^{-1}\vec{1} \leq \frac{\gamma - \beta}{1 - \gamma}\vec{\sigma}_i(\vec{\theta}). \quad (23)$$

The RHS is independent of θ_i , while all the components of the LHS vector are decreasing in θ_i (by (h) from Lemma 5, using the fact that M_i^b is a positive matrix). It is thus sufficient to prove this inequality for $\theta_i = \underline{\theta}$, which we assume from now on. Since M_i^b is positive, the LHS vector is smaller or equal to $\frac{1 - \delta_i\gamma}{\delta_i(1 - \gamma)}\|B_i(\vec{\theta})^{-1}\|_{\infty}M_i^b(\vec{\theta})\vec{1}$. Using (e) from Lemma 5 and the fact that $M_i^b(\vec{\theta})\vec{1} = \frac{1 - \beta}{1 - \gamma}\vec{\sigma}_i(\vec{\theta})$, it is sufficient to check that

$$\frac{1 - \delta_i\gamma}{\delta_i(1 - \gamma)} \frac{1}{\frac{1 - \delta_i\gamma}{\delta_i(1 - \gamma)} + \min_{\ell \neq i} \pi_{\ell i}(\vec{\theta})} \frac{1 - \beta}{1 - \gamma} \leq \frac{\gamma - \beta}{1 - \gamma},$$

or that

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

The result then follows from $\underline{\pi} \leq \min_{\ell \neq i} \pi_{\ell i}(\vec{\theta})$, for all $\vec{\theta} \in [\underline{\theta}, 1]^n$ such that $\theta_i = \underline{\theta}$. ■

Lemma 8. *If the Silent Treatment is a belief-free equilibrium then $\delta_i \geq \frac{1}{\gamma + (\gamma - \beta)\underline{\pi}}$ for all i .*

Proof. The proof of Lemma 7 shows that condition (23) is necessary and sufficient for discerning agents to withhold bad ideas. Given $\underline{\theta}$, consider the ability vector $\vec{\theta}$ for which the minimal probability premium $\underline{\pi}$ is achieved. For the Silent Treatment to be a belief-free equilibrium, it is necessary that it is an ex-post equilibrium for this $\vec{\theta}$. By Lemma 4(e), this ability vector either has all agent abilities equal to $\underline{\theta}$, or there is some agent i with ability $\underline{\theta}$ and all others have ability 1. In both cases, the value of $\pi_{\ell i}(\vec{\theta})$ is constant in ℓ . By the characterization in Lemma 5(a), for this $\vec{\theta}$ we have that

$$B_i(\vec{\theta})\vec{1} = \left(\frac{1 - \delta_i\gamma}{\delta_i(1 - \gamma)} + \underline{\pi} \right) \vec{1}.$$

If a matrix has constant row sums equal to s , then the inverse has constant row sums equal

to $1/s$. Thus

$$B_i^{-1}(\vec{\theta})\vec{1} = \frac{1}{\frac{1-\delta_i\gamma}{\delta_i(1-\gamma)} + \pi}\vec{1}.$$

Applying this expression as well as the fact that $M_i^b(\vec{\theta})\vec{1} = \frac{1-\beta}{1-\gamma}\vec{\sigma}_i(\vec{\theta})$ in the necessary condition (23), we immediately obtain the desired condition on δ_i . ■

References

- Andrews, I. and D. Barron**, 2016. “The Allocation of Future Business.” *American Economic Review* **106**, 2742–2759.
- Banks, J. and R. Sundaram**, 1993. “Adverse Selection and Moral Hazard in a Repeated Elections Model,” in: William Barnett et al. (Ed.), *Political Economy: Institutions, Competition and Representation*. Cambridge University Press.
- Banks, J. and R. Sundaram**, 1998. “Optimal Retention in Agency Problems.” *Journal of Economic Theory* **82**, 293–323.
- 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.
- Duggan, J. and C. Martinelli**, forthcoming. “The Political Economy of Dynamic Elections: Accountability, Commitment and Responsiveness.” *Journal of Economic Literature*, forthcoming.
- Ely, J. and J. Välimäki**, 2002. “A Robust Folk Theorem for the Prisoner’s Dilemma.” *Journal of Economic Theory* **102**, 84–105.
- 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.

- 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**, 2017. “Power Dynamics in Organizations.” *American Economic Journal: Microeconomics*, **9**, 217–241.
- Lipnowski, E. and J. Ramos**, 2016. “Repeated Delegation.” *Working Paper*.
- Piccione, M.**, 2002. “The Repeated Prisoners’ Dilemma with Imperfect Private Monitoring.” *Journal of Economic Theory* **102**, 70–83.
- Schwabe, R.**, 2011. “Reputation and accountability in repeated elections.” *Working Paper*.
- Varah, J.M.**, 1975. “A Lower Bound for the Smallest Singular Value of a Matrix.” *Linear Algebra and Its Applications* **11**, 3-5.