Optimal Delegation and Information Transmission under Limited Awareness*

Sarah Auster[‡] and Nicola Pavoni[§]

April 2019

Abstract

We study the delegation problem between a principal and an agent, who not only has better information about the performance of the available actions but also has superior awareness of the set of actions that are actually feasible. The agent decides which of the available actions to reveal and which ones to hide. We show that it is optimal for the agent to make the principal aware of extremes options, while leaving her unaware of intermediate ones. We further show that unawareness has important effects on strategic information transmission: when principal and agent play a cheap talk game, reducing the principal's awareness can expand the set of implementable equilibrium actions and thereby benefit both parties.

JEL Codes: D82, D83, D86

*The authors thank Sandeep Baliga, Pierpaolo Battigalli, Sylvain Chassang, Wouter Dessein, Antonio Guarino, Lugi Guiso, Marina Halac, Johannes Hoerner, Navin Kartik, Alessandro Pavan, and Yuval Salant for very helpful comments. The present draft benefited from several comments received by the seminar participants at the Toulouse School of Economics, Collegio Carlo Alberto, University College London, EIEF in Rome, Bonn University and at the Max Plank Institute, at the University of Bristol, the University of Cardiff, at the Bocconi and Cattolica universities in Milan, at the University of Verona, University of Malaga, Oxford University, Stockholm School of Economics and the Cowles Conference on Economic Theory. Nicola Pavoni acknowledges financial support from the MIUR-PRIN grant 20157NH5TP.

‡Department of Decision Sciences, Bocconi University: auster.sarah@gmail.com

§Department of Economics, Bocconi University: pavoni.nicola@gmail.com

1 Introduction

In many situations economic agents need to rely on the advice of experts whose preferences are not perfectly aligned with their own. Headquarters depend on division managers who have superior information about the profitability of available projects but also a desire to attract additional resources to their own division, voters rely on politicians whose preferences may reflect a political bias or the interest of certain lobbies, financial investors seek advice from non-neutral financial professionals with a better understanding of the risks and returns of the available portfolios. The tension underlying these situations has been formalized in the delegation model - first introduced by Holmström (1977) - where an uninformed principal specifies a set of permissible actions to the informed agent and contingent transfers are infeasible.

In most of the described situations, it is plausible that the informed party not only has a better understanding of what the most suitable action is but also of the set of actions that are actually available. For instance, corporate headquarters are more detached from the day-to-day business of the different divisions and may thus not be aware of all options the division managers could pursue. Similarly, voters tend to have a limited understanding of available political instruments and legal constraints compared to politicians.[1] Also financial investors differ widely in their financial literacy. They not only face limits in their ability to assess the profitability of particular investments but also have limited awareness of the available investment opportunities.[2]

[1]Somin (2013) and Carpini and Keeter (1992) present the results of a number of surveys on US voters over various decades and document the lack of knowledge of basic institutional rules and of the set of policies available to local governments. For example, Somin documents that 34% of US voters cannot name the three branches of the federal government, a similar percentage do not know which government officials are responsible for which issues; Carpini and Keeter document that less than 50% of US voters know whether the local governors have to approve the decisions made by their higher state court; 25% do not know whether states can pass a law prohibiting abortion.

[2]For example, Guiso and Jappelli (2005) document the lack of awareness of financial assets among the 1995 and 1998 waves of the survey of Italian households (SHIW). Only 65% of potential investors were aware of stocks and only 30% of investment accounts; mutual funds and corporate bonds were known by only 50% of the sample. The share of wealth in the hand of unaware agents was also substantial. The share of wealth owned by households that were not aware of corporate bonds was approximately 20%, and so was the share owned by those unaware of mutual funds.

This paper studies the implications of such asymmetry by incorporating unawareness into the canonical delegation model. We consider the problem of a principal (she) who needs to take an action and delegates the task to an agent (he). The agent has private information about the payoffs of each of the available actions and the principal's problem is to determine a set of actions from which the agent can choose (see for example Alonso and Matouschek, 2008). We depart from this traditional framework of optimal delegation by considering a situation where the agent not only has private information about the suitability of the different actions, but also about the set of actions that are actually available to the principal. This second dimension of asymmetry is captured by the assumption that the principal is only *partially aware of the feasible actions*. Before the delegation stage the agent has the possibility to enrich the principal's awareness by revealing additional actions.

We are interested in the questions of whether the agent expands the principal's awareness, which actions the agent reveals, and what the properties of the realized actions are.

We address these questions in an environment with a continuum of states and a continuum of actions, some of which the principal is aware of. Contingent on the state of the world the agent always wants to take a higher action than the principal, which creates a tension with regard to how much flexibility the agent to grant. In the benchmark case of full awareness the optimal delegation set for the principal is an interval: the principal effectively imposes an upper cap below which the agent is free to choose. In the case of partial awareness, the principal's delegation choice depends on the set of actions she is aware of. Anticipating this, the agent chooses which actions to reveal and which ones to hide. Despite the fact that we impose very little structure on the principal's initial awareness set, we are able to obtain a clean description of the equilibrium: our main result shows that generically the agent leaves the principal unaware of an interval of intermediate options around the optimal upper cap under full awareness.[3] In other words, the agent makes the principal *aware of actions at the extremes*. The awareness gap is chosen in a way such that the principal—who still cares about the agent's information—finds it optimal to permit actions at both extremes. Thus,

[3] More precisely, this always happens unless the principal is exactly aware of the action at the optimal upper bound under full awareness.

by leaving the principal unaware of intermediate options, the agent is able to select actions that would be precluded if the principal was fully aware.

To prove the result we proceed in two steps. First, we characterize the optimal delegation set for an arbitrary awareness set of the principal. Here we extend the characterization result of the existing literature on optimal delegation to the case where the set of actions available to the principal must not be an interval. We show that, within the constraints of her limited awareness, the principal chooses the closest approximation of the optimal delegation interval under full awareness. The optimal upper bound of this approximation is simply the action in the principal's awareness closest to the optimal cap under full awareness. In the second step, we turn to the question of which actions the agent optimally reveals to the principal. The agent must chose over subsets of feasible actions, taking into account the principal's initial awareness set. Solving this optimization problem becomes tractable due to our characterization of the optimal delegation set for a given awareness set: the agent can induce the principal to permit an action strictly higher than the optimal cap under full awareness if and only if the principal remains unaware of an interval of intermediate actions around the cap. We then show that, whenever it is feasible with respect to the principal's initial awareness set, the agent can strictly gain by leaving the principal unaware of such interval. The larger the bias of the agent is, the greater is the size of the interval.

When the agent does not reveal all actions to the principal, there are states of the world where the principal and the agent can be made better off by taking an intermediate action. This gives rise to the question of what happens if after learning the state of the world, the agent can reveal additional actions to the principal. We address this question by extending the baseline model to allow for renegotiation. We show that the agent can use renegotiation to fill the gap in the initial delegation set by proposing an additional action after observing the state of the world. From the modeller's viewpoint the agent's strategy is fully revealing on the interval of intermediate states. The principal, however, only sees the action the agent reveals, while she cannot contemplate actions the agent would have revealed in different states. In the principal's eyes, any proposed action is consistent with several states of the world, which implies that she infers information only partially. Given this, the principal finds it optimal to accept the agent's proposal.

To explore the effect of limited awareness on strategic information transmission further, we also consider the situation where the principal cannot commit to a delegation set (cheap talk). It is well known that under full awareness, equilibria of the cheap talk game partition the state space into a finite number of intervals on which the equilibrium action is constant. Hence, equilibria necessarily entail a loss of information. We show that limited awareness can mitigate the problem: starting from an arbitrary awareness set, the removal of actions from the principal's awareness set expands the set of implementable equilibrium actions, something that not only benefits the agent but can also make the principal better off. Hence again, limited awareness may improve information transmission in this class of games.

Finally, we discuss the robustness of our results. While the baseline model assumes that the set of feasible actions is an interval, we show that the characterization can be extended to the case where the set of feasible actions is an arbitrary subset of the reals. We also discuss more general utility functions and specifications of the bias. Finally, we show how our main result extends when the agent faces uncertainty about the principal's initial awareness.

Related Literature: This paper is first of all related to the literature on optimal delegation. Starting with Holmström (1984), who first defines the delegation problem and provides conditions for the existence of its solution, this literature, which includes Melamud and Shibano (1991), Martimort and Semenov (2006), Alonso and Matouschek (2008), Armstrong and Vickers (2010) and Amador and Bagwell (2013), Halac and Yared (2017), and others, studies optimal delegation problems in environments of increasing generality. None of them consider limited awareness in this framework. Szalay (2005) considers a delegation problem (with full awareness) where the agent has costly access to valuable information and obtains a result that somewhat resembles ours, though the mechanism is completely different. He shows that—in order to motivate information acquisition—it might be optimal to remove intermediate actions from the delegation set so as to increase the stakes for the agent.

Furthermore, our work is related to a relatively small literature on contract theory and unawareness. The application of the concept of unawareness to contracting problems is still at its beginnings. In contrast to our setting, existing work considers contracting problems where contingent transfers are feasible and where the agent is unaware - either of possible actions (Von Thadden and Zhao, 2012 and 2014) or of possible states (Zhao, 2011; Filiz-Ozbay, 2012; Auster, 2013) - while the principal is fully aware.

Finally, we contribute to the literature on signalling and strategic information transmission (e.g. Crawford and Sobel, 1982). To the best of our knowledge, we are the first to study the equilibrium implications of limited awareness on strategic information transmission in a cheap talk environment. Heifetz et al. (2011) study the strategic disclosure of hard information and find that if the information is multidimensional and the receiver is unaware of some dimension, unraveling is not a necessary outcome of the game. Hence, unlike in our framework where information is not verifiable, they find that unawareness of the receiver can lead to less information transmission by the sender. Li and Schipper (2017) test the predictions of their model in a laboratory experiment.

In a companion paper, Auster and Pavoni (2019), we apply our delegation model—without signalling—to financial intermediation, considering a market with multiple fully aware brokers (agents) and a continuum of partially aware investors (principals). We study the effects of competition and investor heterogeneity on the market outcome. Self-reported data from customers in the Italian retail investment sector support the key predictions of the model: the menus offered to less knowledgable investors contain fewer products, most of them nevertheless perceived to be at the extremes.

The paper is organized as follows. The next section presents the delegation model with limited awareness. In Section 3 we derive the equilibrium awareness and delegation set and consider renegotiation. Section 4 analyses the cheap talk game and Section 5 discusses a number of interesting robustness checks. Section 6 concludes.

2 Environment

There is a principal and an agent. The agent has access to a set of actions, the payoffs to which depend on the state of the world. The principal is only aware of a subset of those actions, denoted by $Y^P \subseteq Y^A$.[4] We think of Y^A as a general subset of $\mathbb{R}$, for instance a finite collection of points. For expositional purposes, however, we will initially assume that the underlying set of actions is an interval, i.e. $Y^A = [y_{min}, y_{max}]$. This will simplify notation considerably and we discuss the extension to more general sets in Section 5. We assume the set Y^P is closed but impose no assumptions otherwise.

Let $\Theta = [0, 1]$ be the set of states and let $F(\theta)$ denote the cumulative distribution function on Θ, assumed to be twice differentiable on the support.[5] Both the principal and the agent have von-Neumann-Morgenstern utility functions that take the quadratic form

$$u(y, \theta) = -(y - \theta)^2 \quad \text{and} \quad v(y, \theta) = -(y - (\theta - \beta))^2.$$

The agent's preferred policy is $y = \theta$, while the principal's preferred policy is $y = \theta - \beta$. We assume $\beta > 0$, hence the agent has an upward bias of size β.

The agent is privately informed about the state of the world θ. We rule out monetary transfers and assume that the agent's participation constraint is always satisfied. The contracting problem of the principal then reduces to the decision over the set of actions from which the agent can choose.[6] However, the principal's unawareness restricts the language with which she can write a contract. In particular, we assume that the principal can only include actions into the contract that she can name explicitly. This implies that her delega-

[4]Hence, unawareness in our framework does not take the form of unforeseen contingencies but concerns the set of available actions: while the principal knows the state space, she has an incomplete understanding of the set of actions that are feasible. Karni and Vierø (2013 and 2017) formalize this idea in a decision theoretic model that not only allows for unawareness of contingencies and outcomes but also of acts.

[5]For $\theta = 0$ and $\theta = 1$, this condition holds for, respectively, the right and left derivative.

[6]Formally, the principal commits to a mechanism that specifies the action which will be implemented as a function of the agent's message. Alonso and Matouschek (2008) show that this contracting problem is equivalent to delegating a set of actions $D \subseteq Y^A$ from which the agent can choose freely after observing the state of the world. Their argument continues to hold in our setting. There are thus two possible interpretations: after observing the state of the world, the agent might directly choose an action or make a recommendation which the principal has committed to follow.

tion set must be a subset of her awareness set.[7] Thus, the larger the principal's awareness set is, the richer is the set of contracts she can write.

Before the principal makes her delegation choice and the agent observes the state of the world, the agent can make the principal aware of additional actions. The principal fully understands the options that are revealed to her and accordingly updates her awareness to the union of whatever she knew initially and what the agent reveals. Given her updated awareness, the principal determines a delegation set. Finally the agent learns the state of the world and chooses an action from those permitted by the principal. The timing of the game can be summarized as follows:

1. The principal's initial awareness Y^P is realized and observed by all parties.
2. The agent reveals a set of actions $X \subseteq Y^A$ and the principal updates her awareness to $Y = Y^P \cup X$.
3. Given Y, the principal chooses a delegation set $D \in \mathcal{D}(Y)$, where $\mathcal{D}(Y)$ is the collection of closed subsets of Y.[8]
4. The agent observes the state of world θ and chooses an action from set D.
5. Payoffs are realized.

At this stage we do not need to make any explicit assumption on whether or not the principal is aware of her unawareness. The principal might take the world at face value or she might understand that there exist actions outside her awareness. Since she cannot include such actions in the delegation set, awareness of their possible existence neither affects her expected payoff nor optimization problem.[9] Furthermore, within the constraints of her awareness, the principal is perfectly rational: she anticipates correctly the expected payoff associated to each feasible delegation set and will not be surprised ex-post.

Formally the game between principal and agent can be represented by a family of partially ordered subjective game trees (see Feinberg (2012) and Heifetz et al. (2013)). Such

[7] Alternatively, the principal might specify those actions that the agent is *not* allowed to take. It can be shown that in this case the agent will never have incentives to reveal any actions to the principal.

[8] As discussed in Alonso and Matouschek (2008), the restriction to closed sets is without loss of generality.

[9] We discuss the principal's sophistication and awareness of unawareness further at the end of Section 5.

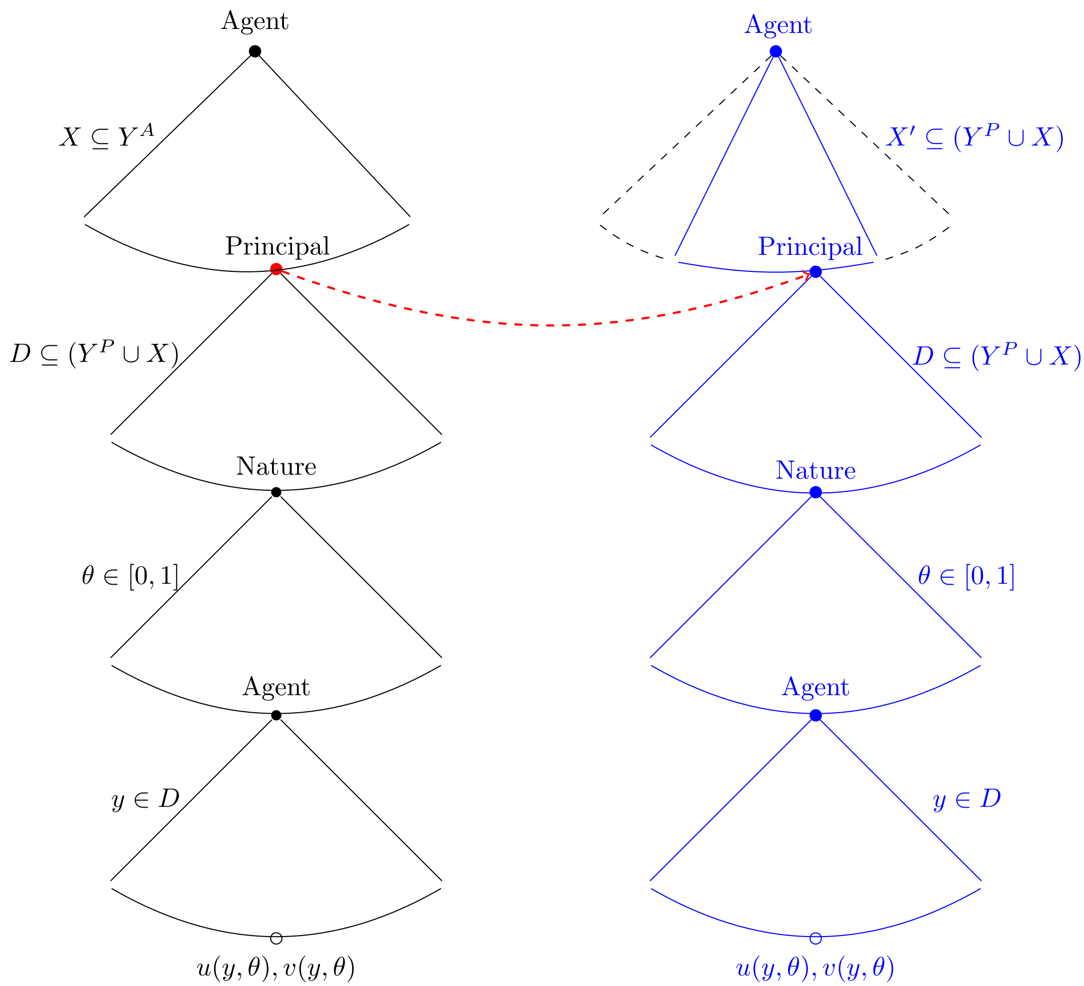

Figure 1: The left side shows the game tree as the agent perceives it. The right side shows the principal's perception of the game, induced by the agent's choice in the first stage (red dot)

family includes the modeler's view of the objectively feasible paths of play but also the feasible paths of play as subjectively viewed by some players, or as the frame of mind attributed to a player by other players or by the same player at a later stage of the game (Heifetz et al., 2013). While we provide a more detailed description of the formalization for our environment in Appendix B, Figure 1 shows an example of the principal's subjective game tree induced by the announcement of the agent.

3 Equilibrium Analysis

We will now proceed with the analysis of the awareness and delegation sets that obtain in equilibrium. Our equilibrium notion is Perfect Bayesian, suitably extended to the class of extensive form games with unawareness (see Feinberg, 2012, Definition 16). We will start our analysis by first describing the benchmark case of full awareness and then turn to the subject of our interest: the case of partial awareness. Before entering the equilibrium analysis, it is useful to mention that optimal awareness sets and optimal delegation sets will typically not be unique since different awareness sets may induce the same delegation set and different delegation sets may induce the same implemented actions for each state of the world. In what follows, we will assume that, if the principal is indifferent between two delegation sets D and D' such that $D' \subset D$, she chooses the larger set D. Similarly, if the agent is indifferent between two revelation strategies that yield awareness sets Y and Y' such that $Y' \subset Y$, we will assume that he expands the principal's awareness to Y. That is, we will consider the sets that yield maximal awareness and maximal discretion.[10]

Throughout the analysis, we will adopt a couple of regularity conditions on the distribution that are common in the delegation literature. Furthermore, we will assume that in each state of the world both the principal's and the agent's ideal actions are available.

Assumption 1. $f'(\theta)\beta + f(\theta) > 0$ *for all* $\theta \in (0, 1)$; *and* $\mathbb{E}[\theta - \beta] > 0$.[11]

Assumption 2. $y_{min} < -\beta$ *and* $y_{max} > 1$.

3.1 Full Awareness

For the specification $Y^P = Y^A$, the existing literature shows that if the density function $f(\theta) \equiv F'(\theta)$ satisfies the first regularity condition in Assumption 1, the optimal delegation set is an interval of the form $[y_{min}, \hat{y}]$ (Martimort and Semenov, 2006, and Alonso and Matouschek, 2008). The second condition in Assumption 1 guarantees that there exists

[10] The restriction only serves to select one of multiple outcome-equivalent equilibria.

[11] All expectations are taken with respect to F.

some $\hat{y} > 0$ that solves[12]

$$\hat{y} = \mathbb{E}[\theta - \beta | \theta \geq \hat{y}]. \tag{1}$$

In that case, the agent chooses his preferred action $y = \theta$ for all $\theta < \hat{y}$ and the action $\hat{y}$ in all remaining states. The following proposition summarizes these results and shows that $\hat{y}$ decreases with the bias β.

Proposition 1. *Under Assumptions 1 and 2 the maximizing solution to the optimal delegation problem with full awareness is an interval of the form* $[y_{min}, \hat{y}]$ *where the (unique) upper bound* $\hat{y} \in (0, 1)$ *solves equation* (1)*.*
Moreover, under the same assumptions, if we let $\hat{y}(\beta)$ *be the cap of the optimal delegation set when principal's preferences parameter is* $\beta \in (0, \mathbb{E}[\theta])$*, then* $\hat{y}(\cdot)$ *is decreasing and continuously differentiable.*

Proof. See Appendix A.1. □

3.2 Partial Awareness: Main Result

Our main result shows that, maintaining the regularity condition on the state distribution, it is strictly optimal for the agent to leave the principal partially unaware if and only if the principal is initially unaware of the action at the optimal threshold under full awareness, $\hat{y}$. In that case, the agent optimally reveals actions at the extremes but leaves the principal unaware of intermediate options.

Proposition 2. *Let Assumptions 1 and 2 be satisfied.*

- *If* $\hat{y} \in Y^P$*, the principal becomes fully aware and the optimal delegation set is* $[y_{min}, \hat{y}]$*.*
- *If* $\hat{y} \notin Y^P$*, the principal remains unaware of actions in* $(\hat{y} - \Delta, \hat{y} + \Delta)$ *for some* $\Delta > 0$ *and the optimal delegation set is* $[y_{min}, \hat{y} - \Delta] \cup \{\hat{y} + \Delta\}$*.*

Proof. See Proposition 3 and Section 3.4. □

[12] If instead $\mathbb{E}[\theta - \beta] < 0$, the optimal delegation set is $[y_{min}, \mathbb{E}[\theta - \beta]]$. In this case, however, the agent will choose the upper bound of the set for all θ, so it is effectively the singleton $\{\mathbb{E}[\theta - \beta]\}$ and delegation has no value.

Proposition 2 shows that whether the principal is made aware of all actions by the agent is determined only by her awareness of $\hat{y}$, the optimal cap under full awareness. If she is unaware of $\hat{y}$, the agent optimally leaves the principal unaware of an interval of actions around $\hat{y}$. As we will show, this makes it optimal for the principal to choose a delegation set that includes an action to the right of $\hat{y}$. By leaving the principal unaware of intermediate actions, the agent thus incentivizes the principal to permit actions that the agent is biased towards and that would be precluded under full awareness. As a result, the equilibrium delegation set is no longer an interval, illustrated in Figure 2.

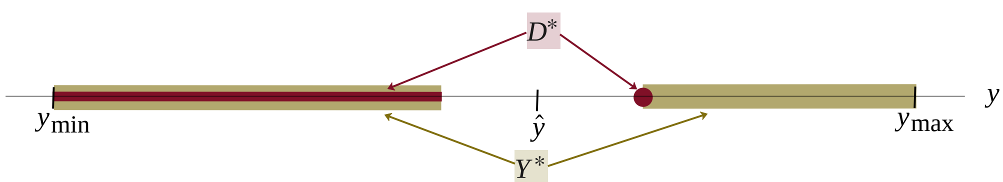

Figure 2: Equilibrium awareness and delegation. The yellow area represents a typical example of equilibrium awareness set Y^* when the principal has limited awareness. The red subset represents the resulting optimal delegation set. When $\hat{y} \notin Y^P$, the principal will be kept unaware of an interval of actions around $\hat{y}$, the cap of the optimal delegation set under full awareness. This way the agent will be allowed to choose the action represented by the red bullet to the right of $\hat{y}$, an action that would be excluded from the delegation set under full awareness.

To prove the statement of the proposition, we need to solve the agent's problem of choosing an awareness set from all subsets of Y^A that include the principal's initial awareness set Y^P. Given that we impose very little structure on Y^P, this optimization problem seems rather complicated. We show, however, that by deriving the properties of the principal's delegation choice for arbitrary awareness sets, we can conveniently restrict our attention to a very simple class of awareness sets from which the agent optimally chooses. We thus proceed recursively: first, we build on the existing literature on optimal delegation and derive the principal's optimal delegation set for an arbitrary awareness set Y; with the solution to this problem, we can then turn our focus to the main interest of this paper, the agent's choice of the awareness set.

3.3 Delegation Choice

Let $D^*(Y)$ denote the optimal delegation set when the awareness set is Y. We can show the following.

Proposition 3. *Let Assumption 1 be satisfied and define* $\hat{y}_Y \equiv \arg\min_{y\in Y} |y - \hat{y}|$ *as the element of* Y *that is closest to* $\hat{y}$*. The optimal delegation set for awareness* Y *is*

$$D^*(Y) = \{y \in Y : y \leq \hat{y}_Y\}.$$

Proof. See Appendix A.3. □

Proposition 3 shows that several known properties of the optimal delegation set extend to the situation where Y is an arbitrary subset of $\mathbb{R}$. First, the principal has no incentives to restrict the agent's choice from below. The agent is upward biased, so whenever he prefers $\min Y$ to another action in the principal's awareness, the principal prefers $\min Y$ as well.

Next, the optimal delegation set $D^*(Y)$ has no "holes" with respect to Y. This result builds on Alonso and Matouschek (2008), who derive conditions under which, in the benchmark case of full awareness, the optimal delegation set is an interval and therefore has no gaps. As we show in the Appendix, their argument perfectly generalizes to generic sets Y that may be non-connected.[13] To gain some intuition, suppose the delegation set includes three actions, y_1, y_2, y_3 with $y_1 < y_2 < y_3$, and consider removing the intermediate action y_2. Then there is an interval of states where the agent switches from y_2 to the low action y_1 and an interval of states where he switches from y_2 to the high action y_3. Given that the principal prefers a lower action than the agent, the first switch benefits the principal but the second one does not, so the question is which effect prevails. The property that the cost of moving away from the bliss point is convexly increasing implies that the cost of the agent switching to y_3 outweighs the gain of the agent switching to y_1, provided that the relative probability mass on the latter event is not too large. The regularity condition on the state distribution assures that this is indeed the case.

[13] In our setting Alonso and Matouschek's (2008) conditions correspond to Assumption 1.

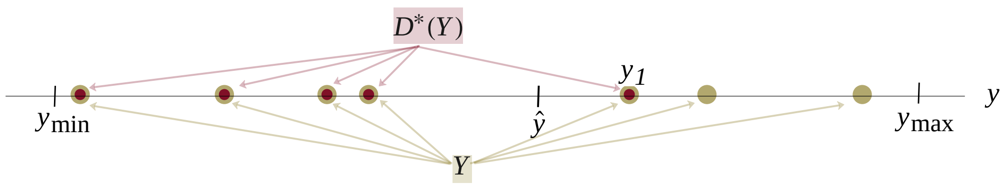

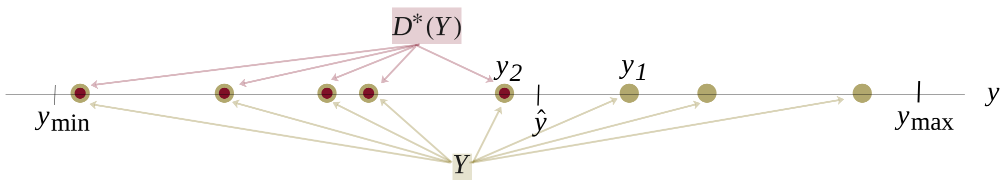

Figure 3: Optimal delegation set $D^*(Y)$. The figures represent two examples of the principal's awareness set Y. In both figures, the yellow bullets represent the set Y while the red bullets represent the resulting optimal delegation set $D^*(Y)$. In the upper figure, the principal includes action y_1 in the delegation set, as it is the closest action to $\hat{y}$. In the lower figure, the principal is aware of action y_2 as well and, for this reason, she excludes action y_1 from $D^*(Y)$. Differently put, the awareness of action y_2 by the principal 'crowds out' action y_1 from the resulting delegation set.

It then remains to determine the upper bound of the optimal delegation set. Proposition 3 shows that the optimal upper bound is given by the element of Y that is closest to $\hat{y}$. This result has two important implications: first, the optimal delegation set includes all actions belonging to Y that are weakly smaller than $\hat{y}$; second, it includes at most one action strictly greater than $\hat{y}$. The optimal delegation set under partial awareness can thus be seen as the closest approximation of the optimal interval under full awareness, $[y_{min}, \hat{y}]$, that is available to the principal given her restricted awareness. This approximation includes an element $y > \hat{y}$ if and only if y is closer to $\hat{y}$ than any element of Y smaller than $\hat{y}$. For a graphical illustration see Figure 8.

3.4 Awareness Choice

We can now turn our attention to the agent's optimal strategy of expanding the principal's awareness. As a first observation, notice that if the principal is aware of the threshold action $\hat{y}$, the agent optimally reveals all other actions. Since there is no action closer to $\hat{y}$ than $\hat{y}$ itself, the upper bound of the optimal delegation set will always be $\hat{y}$. Disclosing actions above $\hat{y}$ is thereby irrelevant; the principal will never allow the agent to implement any of them. On the other hand, revealing actions below the threshold $\hat{y}$ is strictly optimal since they will be included in the optimal delegation set, therefore expanding the agent's choice.

Starting now from an arbitrary set Y^P, the above argument implies that the optimal awareness set Y^* is such that the upper bound of the corresponding delegation set $D^*(Y^*)$ is at least $\hat{y}$. Moreover, the only reason for the agent to leave the principal unaware of certain actions is to induce the principal to permit some action strictly greater $\hat{y}$. By Proposition 3 this is optimal for the principal if and only if the principal is not aware of any action closer to $\hat{y}$. Letting $\hat{y} + \Delta, \Delta \geq 0$ denote the upper bound of the induced delegation set, we thus require $(\hat{y} - \Delta, \hat{y} + \Delta) \cap Y = \emptyset$. At the same time, revealing actions below $\hat{y} - \Delta$ and above $\hat{y} + \Delta$ either does not affect the induced delegation set or strictly expands it. It follows that the optimal awareness set is of the form

$$Y^* = [y_{min}, \hat{y} - \Delta] \cup [\hat{y} + \Delta, y_{max}],$$

with the corresponding delegation set

$$D^*(Y^*) = [y_{min}, \hat{y} - \Delta] \cup \{\hat{y} + \Delta\}.$$

The agent is thus permitted to choose from an interval of actions strictly to the left of the full awareness threshold $\hat{y}$ and one action to the right. Given such delegation set, the agent's optimal policy is as follows. In states below $\hat{y} - \Delta$ the agent uses his flexibility and implements his preferred action $y = \theta$. In states above $\hat{y} - \Delta$ the preferred action is not available, so the agent chooses the one closest to his bliss point. For states in the interval $(\hat{y} - \Delta, \hat{y})$ this is the action $\{\hat{y} - \Delta\}$, for the remaining states it is $\{\hat{y} + \Delta\}$. The agent's optimal policy can thus be sumarized by

$$y^*(\theta; \Delta) = \begin{cases} \theta & \text{if} \quad \theta \leq \hat{y} - \Delta \\ \hat{y} - \Delta & \text{if} \quad \hat{y} - \Delta < \theta < \hat{y} \\ \hat{y} + \Delta & \text{if} \quad \theta \geq \hat{y}. \end{cases}$$

Taken together, the previous analysis provides us with a very simple description of the class of delegation and awareness sets that are candidates for an equilibrium in our environment: when deciding which actions to reveal to the principal, the agent implicitly chooses an awareness gap, parametrized by Δ.

To complete the proof of Proposition 2 it remains to show that whenever a gap is feasible, it is also optimal. We can find the optimal awareness gap by considering the agent's reduced form problem of choosing Δ. The feasible values of Δ are determined by the initial level of awareness of the principal Y^P. In particular, the implementable values of Δ are weakly smaller than $\bar{\Delta}(Y^P) \equiv \min_{y \in Y^P} |y - \hat{y}|$, the smallest distance between an action in the principal's awareness set and $\hat{y}$. For each implementable Δ, the agent then anticipates the principal's optimal delegation choice and his own optimal policy. Substituting $y^*(\theta; \Delta)$ into the agent's expected payoff, his optimization problem amounts to

$$\max_{\Delta \in [0, \bar{\Delta}(Y^P)]} - \int_{\hat{y}-\Delta}^{\hat{y}} (\hat{y} - \Delta - \theta)^2 dF(\theta) - \int_{\hat{y}}^{1} (\hat{y} + \Delta - \theta)^2 dF(\theta). \tag{2}$$

The following proposition characterizes the solution to this problem.

Proposition 4. *Let Assumptions 1 and 2 be satisfied. The solution to problem* (2) *is given by* $\min\{\bar{\Delta}(Y^P), \Delta^*\}$*, where* $\Delta^* > 0$ *solves*

$$\int_{\hat{y}+\Delta^*}^{1} [\theta - (y + \Delta^*)] \, dF(\theta) = \int_{\hat{y}-\Delta^*}^{\hat{y}} [\theta - (y - \Delta^*)] \, dF(\theta) - \int_{\hat{y}}^{\hat{y}+\Delta^*} [\theta - (y + \Delta^*)] \, dF(\theta). \tag{3}$$

Proof. See Appendix A.4. □

The proof of Proposition 4 shows that the agent's payoff as a function of Δ is strictly concave and attains its maximum at Δ^*, as determined by (3). In (3), the left-hand side represents the gain from increasing Δ, while the right-hand side represents the cost of such change. For a graphical illustration see Figure 7 in Appendix A.4. In all states $\theta > \hat{y}+\Delta^*$, the agent gains from a marginal increase in the gap because the new action $\hat{y} + \Delta^*$ is uniformly closer to his ideal point. The cost of increasing the gap is the utility loss in the states $[\hat{y} - \Delta^*, \hat{y} + \Delta^*]$, where the agent moves away from his ideal action.

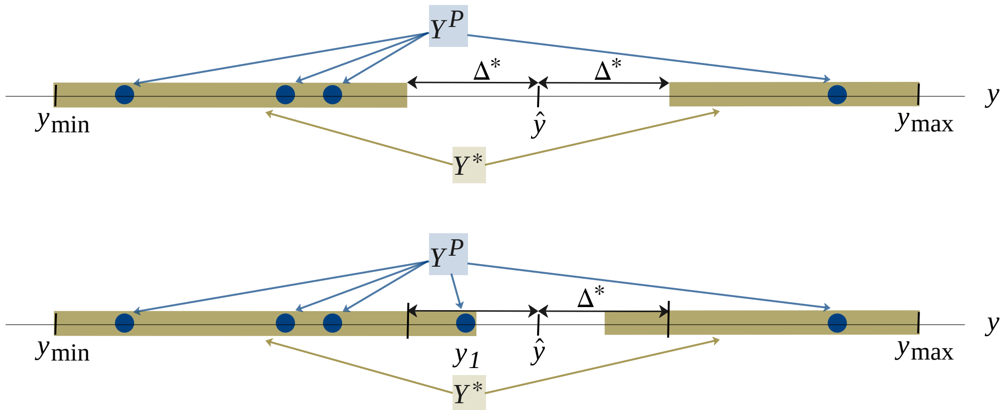

Figure 4: Optimal awareness set Y^*. The figures represent two examples of the principal's initial awareness set Y^P and associated awareness sets $Y^* = Y^P \cup X^*$ after including disclosed actions X^*. In both figures, the blue bullets represent the set Y^P, while the yellow set represents the resulting optimal awareness set Y^*. In the upper figure, the agent keeps the principal unaware of the interval $(\hat{y} - \Delta^*, \hat{y} + \Delta^*)$. In the lower figure, the principal is also aware of action y_1 and, for this reason, the agent finds it optimal to increase the principal's awareness.

The proposition states that the unconstrained solution Δ^* is strictly positive. This can be easily understood by considering the net effect of increasing the gap at $\Delta = 0$. While the left-hand side is always positive, the right-hand side at $\Delta = 0$ equals zero: the marginal cost of moving away from the bliss point at the bliss point is zero.

Proposition 5. *Let $\Delta^*(\beta)$ be the unrestricted solution to problem* (2)*, as described in Proposition 4, when the principal's preferences parameter is $\beta \geq 0$. Then $\Delta^*(\cdot)$ is an increasing function.*

Proof. See Appendix A.5. □

The proposition shows an intuitive result: the larger the divergence between the principal's and the agent's preferred action is, the more actions the agent wants to hide from the principal. The solution Δ^* is implemented whenever the principal's initial awareness does not constrain the agent in his choice of the gap. If, however, the principal is aware of some action in the interval $(\hat{y} - \Delta^*, \hat{y} + \Delta^*)$, the agent's optimal strategy is to simply choose the largest feasible gap, as shown in Figure 4.

3.5 Renegotiation

Thus far, we have assumed that the agent can only reveal actions to the principal before he learns the state of the world. This entails that even in states where the agent knows of an action that makes both parties better off, additional communication is not possible. One interesting question is how the equilibrium outcome changes if after learning the state of the world, the agent can reveal a set of additional actions to the principal, who then decides whether to permit a new action or to maintain the original contract. We thus consider a model where the agent can renegotiate with the principal, at least with a certain probability σ. If the agent gets the opportunity to renegotiate and proposes to replace the original delegation set with a new action, the principal understands that the agent's choice signals something about the state of the world. In particular, the principal can infer that the agent only reveals an action if its inclusion benefits him. However—due to the principal's limited awareness—she cannot conceive of alternative actions the agent could have disclosed. This implies that the principal cannot learn from particular actions not being disclosed, an important difference to the case of full awareness.

Proposition 6. *Let Assumptions 1 and 2 be satisfied and suppose the agent can renegotiate with probability $\sigma \in [0, 1]$ after learning the state of the world. The best equilibrium outcome for the agent is parametrised by Δ, and has the following properties:*

- *before learning the state, the agent reveals actions that do not belong to the set $(\hat{y} - \Delta, \hat{y} + \Delta)$ and the principal chooses delegation set $[y_{min}, \hat{y} - \Delta] \cup \{\hat{y} + \Delta\}$;*
- *if θ belongs to $(\hat{y} - \Delta, \hat{y} + \Delta)$, the agent proposes his preferred action $y = \theta$, which is then permitted by the principal.*
- *if θ does not belong to $(\hat{y} - \Delta, \hat{y} + \Delta)$, the agent takes his preferred action in the original delegation set;*

The equilibrium parameter Δ is weakly increasing in σ. When $\sigma = 1$, then $\Delta = \bar{\Delta}(Y^P)$.

Proof. See Appendix A.6. □

Proposition 6 shows that there is an equilibrium in which the agent leaves the principal unaware of a gap around $\hat{y}$, as in the case without renegotiation. Due to her limited awareness, the principal does not expect the agent to reveal other actions in the future and therefore chooses the delegation set $D = [y_{min}, \hat{y} - \Delta] \cup \{\hat{y} + \Delta\}$. If the realized state θ falls in the gap, the agent additionally reveals (and recommends) his preferred action $y = \theta$. The principal infers that the agent prefers y over all other actions in the original delegation set and decides whether, conditional on the agent preferring y over any other element in D, she prefers y over the action in D which the agent implements if not permitted y. The proof of Proposition 6 shows that the principal prefers to permit action y if and only if the principal would have permitted y initially, had she been aware of y. Indeed, when choosing the initial delegation set, the principal evaluates the profitability of including an additional action only on those states where the action is preferred by the agent. From Proposition 3 we then know that permitting y is optimal for the principal when there is no other action closer to $\hat{y}$. It follows that conditional on having the chance to renegotiate, the agent is able to implement any action below $\hat{y} + \Delta$ and is unable to implement any action above $\hat{y} + \Delta$.

Dynamic awareness. The model with renegotiation highlights two important aspects of games with limited awareness. The first concerns the dynamics of unawareness. Much like information, unawareness is not reversible. This means that if a player becomes aware of an action today, he remains aware of that action in the future (similarly for outcomes, events, etc.). Hence, the more a player reveals at an early stage of the game, the smaller is the collection of awareness sets from which he can choose later on. When there is uncertainty about the future, this creates incentives to hide feasible actions from the other player until later stages of the game. In our model with renegotiation, this is reflected in the fact that, whenever feasible, the optimal gap of the initial awareness set is strictly larger when there is a positive probability of renegotiation compared to when that probability is zero. The agent compromises between choosing the optimal on-shot awareness set, parameterised by Δ^*, and retaining flexibility in the future. The more weight the agent assigns to the option of renegotiation in the future, the larger is the initial gap of actions he hides from the principal .

Unawareness and strategic information transmission. In the renegotiation stage, principal and agent play a signalling game. One striking feature of the equilibrium described in Proposition 6 is that in the interval $(\hat{y}-\Delta, \hat{y}+\Delta)$ the implemented equilibrium action is linearly increasing in θ. In other words, there is no pooling of types below the threshold $\hat{y}+\Delta$. This would not be possible under full awareness: in any candidate equilibrium where types below $\hat{y}+\Delta$ separate themselves through their announcement, the fully aware principal learns the state of the world and has incentives to deviate to a strictly lower action. In the case of limited awareness, the principal cannot contemplate moves of the agent she is unaware of and this limits the extent to which she infers information from the agent recommendation. In particular, if the state is θ and the agent proposes $y=\theta$, the subjective game tree that represents the principal's frame of mind after updating does not include moves of the agent involving an action $y' \in (\hat{y}-\Delta, \hat{y}+\Delta), y' \neq y$. As a consequence, the principal cannot conceive of the fact that she would have permitted action y' if the agent had proposed y' instead. In her subjective game, there is an equilibrium where the agent reveals y in all states that are closer to y than to any element of the initial delegation set. Conditional on that information, the principal indeed prefers action y to the initial delegation set.

The commitment of the principal to the initial delegation set is important for our equilibrium. When the agent reveals an action y in the interval $(\hat{y}-\Delta, \hat{y}+\Delta)$ and the principal updates accordingly, she might in fact prefer the lower action $\hat{y}-\Delta$ over y. Since, however, the principal is committed to the initial contract, she cannot force the agent to take $\hat{y}-\Delta$. Effectively she only has two options, either to permit the new action y or to maintain the original delegation set, which also includes $\hat{y}+\Delta$. If instead the principal could renege on the original contract, she would sometimes choose a lower action than y. Anticipating this, the agent would only reveal additional actions to the principal that he expects the principal to approve. When F is uniform, the condition under which the principal prefers a newly revealed action $y \in (\hat{y} - \Delta, \hat{y} + \Delta)$ over all actions in her awareness set is $\Delta \geq 2\beta$.[14] Thus, in order for the principal to approve a new action, the gap has to be sufficiently large with respect to the bias. This suggests that in *the absence of commitment, effective communication may require a sufficient degree of unawareness,* an idea we explore further in the following section.

4 Cheap Talk

So far the principal was able commit to a delegation set from which the agent could choose once he observed the state of the world. An alternative interpretation is that the principal offers an incentive compatible mechanism that specifies an action as a function of the agent's message. We now remove the assumption that the principal can commit to such a mechanism and only allow principal and agent to communicate after the agent learns the state of the world.[15] Principal and agent thus play a cheap talk game à la Crawford and Sobel (1982), where a biased sender (the agent) transmits a message to an uninformed receiver (the principal). With respect to the previous analysis, this means that the principal will always take her preferred action in her awareness set after updating. Moreover, the agent's outside option of not revealing additional actions after learning the state of world will not be fixed by an existing contract but will depend on the equilibrium of a cheap talk game.

Equilibrium notion. We start by defining the strategies of the principal and the agent. Let $\mathcal{Y}^P$ be the set of all closed subsets of $[y_{min}, y_{max}]$ that include Y^P. The agent's strategy maps each state of the world $\theta \in [0, 1]$ into an awareness set $Y \in \mathcal{Y}^P$ and a message $m \in \mathcal{M}$, which we can interpret as a recommendation. The agent's strategy is thus described by a pair of functions $(Y(\cdot), m(\cdot))$. The principal's strategy is a mapping $\rho : \mathcal{Y}^P \times \mathcal{M} \to [y_{min}, y_{max}]$, assigning to each awareness/message pair an action. Assuming that the agent plays a pure strategy, we can define a belief function for the principal, T, mapping each

[14] This condition is obtained by a simple manipulation of $y - \mathbb{E}[\theta - \beta] \leq \mathbb{E}[\theta - \beta] - (\hat{y} - \Delta)$ with $\mathbb{E}[\theta - \beta] = (y + \hat{y})/2 - \beta$.

[15] All our results hold if principal and agent can additionally communicate before the state is realized.

awareness/message pair (Y, m) to a subset of $[0, 1]$. The set $T(Y, m)$ describes the states which the principal considers possible when the agent discloses set Y and sends message m. A pure-strategy equilibrium for our game is then defined as follows.

Definition 1. *The strategy profile $((Y^*, m^*), \rho^*)$ and belief function T^* constitute an equilibrium in the cheap-talk game with limited awareness if the following conditions hold:*[16]

- *Principal optimality: for all $(Y, m) \in (\mathcal{Y}^P \times \mathcal{M})$,*

$$\rho^*(Y, m) \in \arg\max_{y \in Y} \mathbb{E}[v(y, \theta) | \theta \in T(Y, m)];$$

- *Agent optimality: for all $\theta \in [0, 1]$,*

$$(Y^*(\theta), m^*(\theta)) \in \arg\max_{(Y,m)} u(\rho^*(Y, m), \theta);$$

- *Consistency of beliefs: for all pairs (Y, m) that are chosen on path,*

$$T^*(Y, m) = \{\theta \in [0, 1] : u(\rho^*(Y, m), \theta) \geq \max_{Y' \subseteq Y} \max_{m' \in \mathcal{M}} u(\rho^*(Y', m'), \theta)\}.$$

The consistency condition assures that the principal's beliefs are coherent with the agent's equilibrium actions, as perceived through the principal's awareness. According to the condition, the principal evaluates payoffs for lower levels of awareness using her own equilibrium strategy. Hence, with respect to lower levels of awareness that are reached on path, the principal's beliefs are correct.

Notice that in a pure-strategy equilibrium, the principal's and agent's expected payoffs are determined by the set of on-path equilibrium actions: in equilibrium the agent chooses his preferred action from the set of actions the principal is willing to take. Hence, equilibrium payoffs are the same as when the principal directly delegates the set of equilibrium actions.

In the case of full awareness ($Y^P = [y_{min}, y_{max}]$), there exist multiple equilibria, all of which partition the state space into a finite number of intervals on which the equilibrium

[16] As before, all expectations are taken with respect to F.

action is constant. Moreover, the information conveyed in the most informative equilibrium is decreasing in the bias of the expert (see Crawford and Sobel, 1982). Hence, any equilibrium entails a loss of information. The following proposition shows that one surprising remedy to the problem is limited awareness. In particular, starting from an arbitrary awareness set, more actions can be implemented in equilibrium when the principal's initial awareness set is smaller. The following proposition shows that this may imply a welfare improvement for both the principal and the agent.

Proposition 7. *Let Y^P be a non-empty, non-singleton awareness set. Consider a cheap talk game with the principal's initial awareness being Y^P and fix a pure-strategy equilibrium. There exists an awareness set $\tilde{Y}^P \subset Y^P$ for the principal given which the cheap talk game has a pure-strategy equilibrium which, compared to the equilibrium with Y^P, leads to weakly more actions taken on path and generates a weakly higher expected payoff for the principal and the agent.*

In Proposition 7, we assume that the initial awareness set of the principal includes at least two actions. If instead Y^P is a singleton, the only strict subset of Y^P is the empty set. When the principal is not aware of any actions, the agent can always reveal his preferred action and thereby force the principal to take it. In order for the principal to have some bargaining power, she needs to have an outside option, determined by the actions in her initial awareness. Proposition 7 shows that provided such outside option can be kept, both the principal and the agent can be made better off by limiting the principal's initial awareness.

The proof of Proposition 7 consists of two steps. First, we show that given an awareness set Y^P, in any pure-strategy equilibrium the highest equilibrium action, called $\bar{y}$, must be weakly smaller than $\hat{y} + \bar{\Delta}(Y^P)$ (recall that $\bar{\Delta}(Y^P)$ is the minimal distance between $\hat{y}$ and some element of Y^P). We show next that restricting the principal's awareness to a singleton $\{y\}, y \in Y^P$ with $y \leq \bar{y}$, there is an equilibrium in which the set of equilibrium actions is an interval, $[0, \bar{y}]$. Hence, by eliminating some elements of the principal's initial awareness set, it is possible to fill the gaps between original equilibrium actions. This not only benefits the agent but, as we argue in Section 3, also the principal.

To see why the described equilibrium can be sustained, consider the case where the principal is initially aware of a single action y. When the agent reveals some action $y' \neq y$, the principal's frame of mind is described by a partial game which includes the agent's choice between revealing and not revealing y' but does not include subtrees that follow a different revelation of the agent, illustrated in Figure 5. This subjective game has an equilibrium in which the agent proposes his preferred action in every state of the world. The principal believes that the agent proposes a new action if and only if the agent prefers that action over y. Given these beliefs, the principal follows the agent's advice if and only if the proposed action is closer to $\hat{y}$ than y, a result that directly follows from Proposition 3. In the described equilibrium, the set of equilibrium actions is maximal given the principal's awareness. Evidently, we can find suitable out-of-equilibrium beliefs, that support smaller sets of equilibrium actions. For instance, if $y < \hat{y}$, we can sustain an equilibrium with actions $[0, \hat{y}]$, simply by having the principal believe that the state is y after learning about actions above $\hat{y}$. This equilibrium yields the same expected payoffs that obtain under full awareness and commitment. However, it involves the use of rather unreasonable out-of equilibrium beliefs.

Neologism proofness. A natural requirement is that the agent reveals additional actions whenever such revelation benefits him and that this is reflected in the principal's beliefs. To formalise the idea, we consider neologism proofness, suitably adapted for a game with limited awareness. The original definition, introduced by Farrell (1993), rules out equilibria in which there exists a set $G \subset [0, 1]$ such that the agent prefers the principal's best response to the state belonging to G over his equilibrium payoff if and only if his type is in G. Such a set is called "self signalling". Farrell (1993) argues that when a self signalling set G exists and players have a common language capable of conveying the literal meaning of "the state belongs to G," such message should be credible.

In our game with limited awareness, the principal's best response depends on the set of actions the agent reveals. Moreover, when assessing the credibility of the agent's message,

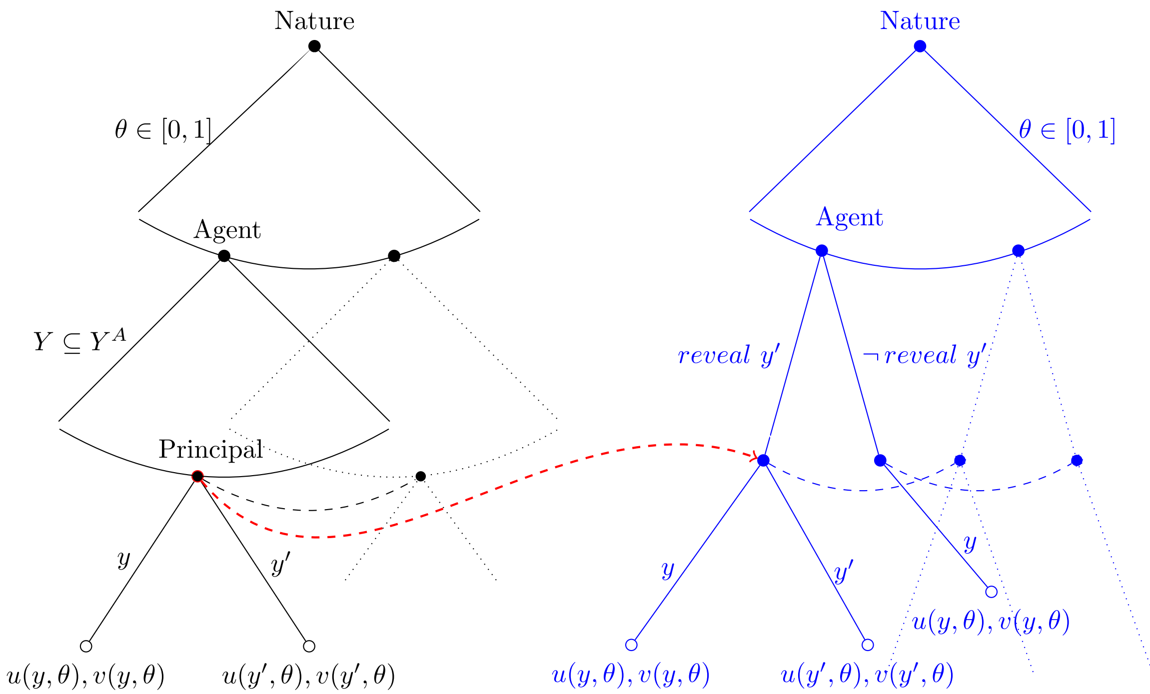

Figure 5: The left game tree represents the modeller's view of a cheap talk game with the principal's initial awareness set being $\{y\}$. The left game tree represents the principal's view of the game induced by the agent revealing action y' (red dot)

the principal evaluates the agent's equilibrium payoff according to her level of awareness. Hence, whether (or not) a set G is self signalling depends on the awareness set the agent induces. To formalise these requirements, fix an equilibrium and let

$$u^*(Y, \theta) \equiv \max_{Y' \subseteq Y} \max_{m' \in \mathcal{M}} u(\rho^*(Y', m'), \theta)$$

describe the principal's perception of the agent's highest attainable payoff at state θ when the principal's awareness is Y. Furthermore, let $BR(Y, G) \equiv \arg\max_{y \in Y} \mathbb{E}[v(y, \theta)|\theta \in G]$ be the principal's optimal action in Y, conditional on the state belonging to the set G. When there are two maximisers in Y, let $BR(Y, G)$ take the value of the larger maximiser. Hence, the principal's best-response function breaks ties in the agent's favour, strengthening the refinement. We extend Farrell's (1993) notion of neologism proofness as follows.

Definition 2. *Given a strategy profile $((Y^*, m^*), \rho^*)$ and a belief function T^*, the pair (Y, G)*

is self signalling if

$$G = \{\theta : u(BR(Y,G),\theta) > u^*(Y,\theta)\}.$$

An equilibrium is neologism proof if there exist no self signalling sets relative to the equilibrium.

We show that our previous result (Proposition 7) continues to hold when we require equilibria to be neologism proof. In particular, if there exists a neologism proof equilibrium for some awareness Y^P, we can find a smaller awareness set and a neologism proof equilibrium that induces more actions and makes both the principal and the agent better off.

Proposition 8. *Let Y^P be a non-empty, non-singleton set, describing the principal's initial awareness, and assume there is a neologism proof equilibrium. There exists an awareness set $\tilde{Y}^P \subset Y^P$ for the principal given which the cheap talk game has a pure-strategy neologism proof equilibrium which, compared to the equilibrium with Y^P, leads to weakly more actions taken on path and generates a weakly higher expected payoff for the principal and the agent.*

When out-of-equilibrium beliefs are not restricted, the highest equilibrium action is weakly smaller than $\hat{y} + \bar{\Delta}(Y^P)$ (see above). When we require neologism proofness, the highest equilibrium action is exactly $\hat{y} + \bar{\Delta}(Y^P)$. Given this property, the statement of Proposition 8 can be shown by eliminating all but the closest action to $\hat{y}$ from the awareness set Y^P. Given the new awareness set, we can construct a neologism proof equilibrium with equilibrium actions $[0, \hat{y} + \bar{\Delta}(Y^P)]$. As before, the new set of equilibrium actions fills any potential gaps of the original set, which benefits both the agent and the principal.

Recall that when the principal has commitment power (as in the baseline model), being aware of additional actions can only make her better off. The previous two results show that when the principal cannot commit, a lack of awareness may lead to equilibria that improve the welfare of both players.

Unawareness helps the principal for two reasons. 1) The principal cannot deviate to actions outside her awareness. Hence, unawareness can work as a commitment device. 2) The principal does not infer the full extent of information the agent's equilibrium actions convey.

To understand this point better, note that the principal only observes one realisation of the agent's revelation strategy. As a consequence, the principal cannot conceive of actions that the agent would have revealed in different states of the world nor can she contemplate what information she would have inferred from such revelation. This means that—in contrast to the case of full awareness—the principal's possible information may not be representable by a partition of the state space into pairwise disjoint information sets. Indeed, in the equilibrium described above with initial awareness $\{y\}$, the principal's information is represented by *overlapping intervals*. For proposals higher than y, these intervals are of the form $[t, 1]$. Hence, after the agent proposes an action greater than y the principal always considers the highest states possible and, as a consequence, has no incentives to choose y instead.

While being aware of more actions can be detrimental to the principal's welfare, being aware of particular actions helps her. As we have seen, the principal's closest action to $\hat{y}$ in her initial awareness set Y^P determines an upper bound for the actions that are taken in equilibrium, equal to $\hat{y} + \bar{\Delta}(Y^P)$. Since the principal's preferred upper bound is $\hat{y}$, she benefits from being aware of an action close to $\hat{y}$, similar to the case of commitment. When the principal's initial awareness set is $\{\hat{y}\}$, there is in fact a cheap talk equilibrium which, by inducing actions in $[0, \hat{y}]$, maximizes the principal's payoff over all incentive compatible allocations.

5 Discussion

Set of feasible actions. In the baseline model we assumed that the set of available actions is $Y^A = [y_{min}, y_{max}]$. A possible concern is that the principal, despite being unaware of certain elements of Y^A, might understand that Y^A is an interval. The principal could then attempt to include actions outside her awareness, maybe through a description of the properties of such actions. However, as we suggested in the beginning, our analysis applies when Y^A is an arbitrary subset of $\mathbb{R}$, e.g. a finite set, so that *a priori* there is no specific structure of the set of available actions—or simply the awareness set of the agent—that might be commonly known.

To see this, assume Y^A is an arbitrary closed subset of $\mathbb{R}$. The analysis of the optimal delegation set for a given awareness set $Y \subseteq Y^A$ in Section 3.3 remains valid, so we have $D^*(Y) = \{y \in Y : y \leq \hat{y}_Y\}$. With regard to the optimal awareness set, we can first notice that if the agent reveals some $y \in Y^A$, he also reveals all those actions that have a greater distance to $\hat{y}$ than y: their inclusion will weakly expand the agent's choice set. This implies that the optimal awareness set can again be described by a gap Δ and takes the form

$$Y^* = \{y \in Y^A : |y - \hat{y}| \geq \Delta\} \quad \text{with} \quad 0 \leq \Delta \leq \bar{\Delta}(Y^P),$$

where $\bar{\Delta}(Y^P)$ is the minimal distance between $\hat{y}$ and an element of Y^P, as above.

Whether or not the agent reveals all feasible actions to the principal depends on the particular form of Y^A and the principal's initial awareness Y^P. A sufficient condition for full awareness is $\hat{y}_{Y^P} = \hat{y}_{Y^A}$, i.e. the principal is aware of the action that is closest to $\hat{y}$. When this is not the case, the agent leaves the principal unaware of intermediate actions, provided they are close enough to $\hat{y}$ and that there exists a greater action than $\hat{y}_{Y^A}$ that is implementable given the principal's initial awareness.

Quadratic loss preferences and constant bias. While the utility functions we consider are rather special, we should note that the key insights of our model—both with and without commitment—can be extended to a larger class of preferences. In particular, the main result on optimal delegation with limited awarenes—the fact that the agent has an incentive to leave the principal unaware of a set of actions around the optimal threshold under full awareness—can be considerably generalised: as long as the principal's and agent's preferences are represented by smooth, single-peaked utility functions that have the property by which the ideal action is strictly monotonic in the realized state of the world, the agent's incentives to leave an awareness gap are much the same as in the baseline model. Imposing an appropriate regularity condition on the state distribution, the optimal delegation set for a fully aware principal facing an upward biased agent is still an interval with some upper bound $\hat{y}$.[17] Since the principal cares about the agent's information, the agent can then find some awareness gap around $\hat{y}$ (not necessarily symmetric) such that the principal optimally permits an action greater than $\hat{y}$, provided that $\hat{y} \notin Y^P$. Under the assumption that the agent's utility function is differentiable, we can then replicate the argument following Proposition 4: given that the marginal cost of moving away from the bliss point at the bliss point is equal zero, the net benefit of introducing a marginal gap around $\hat{y}$ will be strictly positive. In Appendix C we show more concretely how our equilibrium characterization extends to a larger class of models.

Private awareness. The assumption that the agent knows the principal's awareness set is rather strong. For instance, a financial expert may be uncertain about the investment options an investor has encountered before their meeting. We can show that also in a situation where the agent faces uncertainty about the principal's initial awareness set, the agent optimally discloses a set of actions which has a gap around $\hat{y}$.

Proposition 9. *Let the agent's belief about the principal's awareness set be described by a probability distribution over the set of closed subsets of $[y_{min}, y_{max}]$. The set of actions that the agent optimally reveals takes the form*

$$Y^* = [y_{min}, \hat{y} - \Delta] \cup [\hat{y} + \Delta, y_{max}],$$

for some $\Delta \geq 0$.

Proof. See Appendix A.9. □

In the proof of Proposition 9, we show that the agent can improve on any set of actions by disclosing all actions that have a weakly greater distance to $\hat{y}$ than the closest one in the set. No matter what the realized awareness set of the principal is, actions that are further away from $\hat{y}$ than the closest one do not crowd out any additional actions. The optimal size

[17] The optimal delegation literature provides conditions that make interval delegation optimal for a considerably larger class of environments. For the most general treatment see Amador and Bagwell (2013).

of the awareness gap for the agent is determined by his beliefs about the principal's initial awareness. The following result shows that the agent leaves the principal unaware of some actions whenever he assigns a strictly positive probability to the event that the principal's awareness is bounded away from $\hat{y}$.

Corollary 10. *Let the agent's belief about the principal's awareness set be described by a probability distribution over the set of closed subsets of* $[y_{min}, y_{max}]$ *with finite support. If the agent assigns a positive probability to the event that the principal's awareness set does not include* $\hat{y}$*, then* $\Delta > 0$.

Proof. See Appendix A.10. □

The agent's cost of not disclosing actions around $\hat{y}$ when facing a principal who is aware of $\hat{y}$ is the (potential) loss of flexibility below $\hat{y}$. However, as we argued after Proposition 4, at $\Delta = 0$ the agent's marginal utility loss associated to the reduced flexibility equals zero, since in states below $\hat{y}$ the agent is at his bliss point. This implies that as long as the agent assigns a positive probability to the event that the principal is not aware of $\hat{y}$ and, hence, that there is a strict gain of introducing a gap, the net effect of marginally increasing Δ at $\Delta = 0$ is positive.

Principal's sophistication and awareness of unawareness A key component of our baseline model is that the agent can increase the awareness of the principal before the principal makes her delegation choice and the agent observes the state of the world. We assumed that when additional options get revealed to the principal, she simply updates her awareness to the union of whatever she knew initially and what the agent reveals. This might suggest that the principal is naive in that she does not contemplate the possibility of other actions of which she is not aware. One could think that a sophisticated individual might become suspicious and take some defensive counteractions, such as refusing the contract or actively gathering information. An important question is what would trigger such suspicion.

First of all, we should note that the agent's equilibrium announcement is justifiable for the principal in the sense that it is consistent with the principal believing that the agent acts rationally. In our game, the requirement of justifiability, introduced by Ozbay (2008), states that given the principal's updated awareness Y, the principal cannot conceive of any announcement strategy which yields a higher expected payoff for the agent than Y. In other words, in equilibrium the principal should not believe that if the agent had revealed fewer actions, he would have been better off. Since the delegation set resulting from the optimal awareness set Y^* yields a weakly greater payoff for the agent than the payoff associated to any other subset $Y \subseteq Y^*$, this requirement is always satisfied in our setting.

It might however be the case that the mere revelation of new actions triggers the principal's suspicion. A simple assumption is that, once made suspicious, the principal loses trust in the agent and outright refuses to contract with him. The solution would then be straight forward: the agent should never reveal any new actions at the initial stage and the resulting optimal delegation set $D(Y^P)$ would be characterised by Proposition 3.

Lastly, one could consider a dynamic environment—much richer then ours—where the principal has ways to expand her awareness set (for example, by using a costly technology or by sampling other agents). In this case, the initial awareness of being partially aware and, perhaps more importantly, the assessment of the value of discovering new options would be key (e.g., see Karni and Vierø, 2017).

6 Conclusion

In the first part of the paper, we study the delegation problem between an agent and a principal with limited awareness of the available actions. We show that the agent finds it optimal to make the principal aware of actions at the extremes, but leaves her unaware of intermediate options. An achievement of this paper is to formulate a flexible model of delegation with limited awareness and derive a number of properties of the optimal solution. Despite the potential complexity of the resulting double delegation problem, the solution found is remarkably simple and can easily be embedded into more complex frameworks. The key insights of the theory are robust to a number of possible extensions, in particular regarding the set of feasible actions, the player's preferences, and the information structure.

We then modify the game to allow for signalling and show that limited awareness has important implications for the outcomes of such models. In particular, we demonstrate that limited awareness can restrict inference of information on the equilibrium path and argue that, in some situations, this can help to achieve better outcomes compared to the case when all parties are fully aware. We believe that a further study of such games is a promising task for future research.

Appendix

A Proofs

For the proofs of the following results it is useful to introduce the terms

$$T(y) \equiv F(y)\left(y - \mathbb{E}[\theta - \beta|\theta \leq y]\right),$$

and

$$S(y) \equiv (1 - F(y))\left(y - \mathbb{E}[\theta - \beta|\theta \geq y]\right),$$

in the literature referred to as, respectively, backward bias and forward bias (see Alonso and Matouschek, 2008). By Assumption 1 we have

$$T''(y) = \beta f'(y) + f(y) > 0 \quad \text{and} \quad S''(y) = -(\beta f'(y) + f(y)) < 0 \quad \text{for all } y \in [0, 1]$$

Note first that - since $\beta > 0$ - we have $T(y) \geq 0$ for all $y \in [y_{min}, y_{max}]$ and $T(y) > 0$ for $y \geq 0$. The variable S may change sign. Noticing however, that $S(\hat{y}) = S(1) = 0$, strict concavity of S implies that $S(y) > 0$ for all $y \in (\hat{y}, 1)$.

A.1 Proof of Proposition 1

Alonso and Matouschek (2008) show that under Assumption 1 the class of delegation sets are intervals of the form $[y_{min}, \bar{y}]$, in particular, the delegation set does not contemplate 'holes'. We do not repeat their proof here. In Figure 6 we provide a graphical intuitive explanation for the result. In the proof of Proposition 3, we prove the result allowing for partial awareness. We here show the implicit expression defining $\hat{y}$ and its monotonicity.

Given that we can concentrate on intervals of the mentioned class, the relevant objective function is as follows:

$$-\int_{y_{min}}^{\bar{y}} (\beta)^2 \, dF(\theta) - \int_{\bar{y}}^{1} \left(\bar{y} - (\theta - \beta)\right)^2 dF(\theta).$$

Note that the objective function is continuous and the set Y is compact so a maximum exists. The derivative of the objective function with respect to $\bar{y}$ is:

$$-\bar{y}(1 - F(\bar{y})) + \int_{\bar{y}}^{1} (\theta - \beta) dF(\theta). \tag{4}$$

Equation (1) is a simple rearrangement of the first order condition for the principal. The

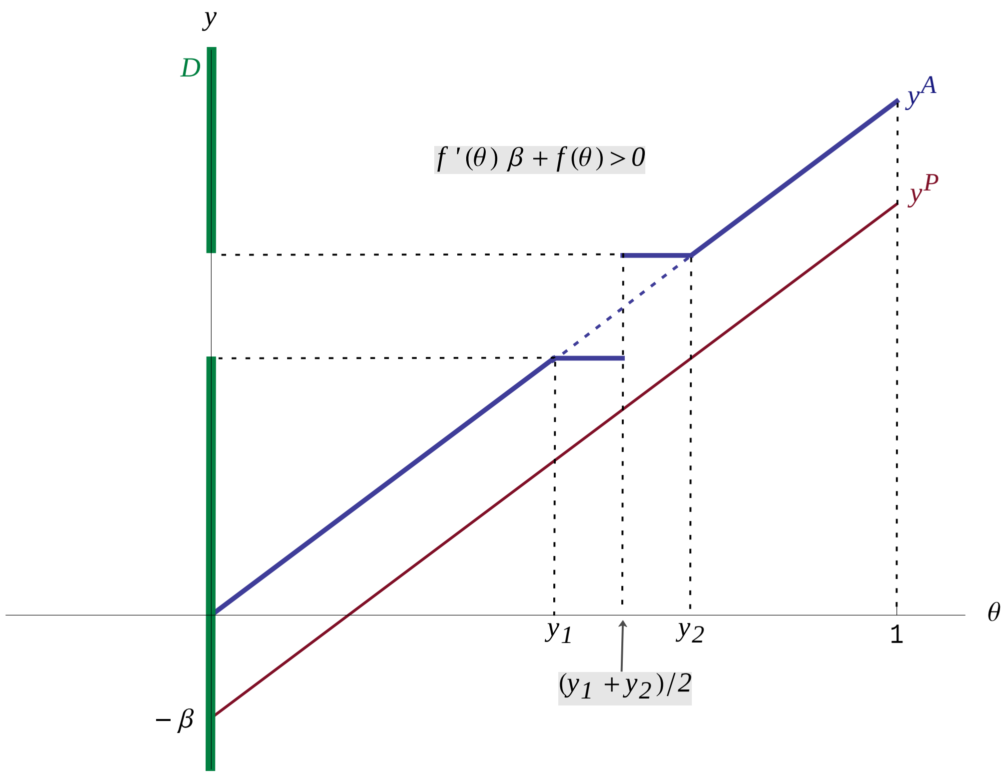

Figure 6: Alonso and Matouschek (2008) show that under Assumption 1 the delegation set has no 'holes', namely it is an interval. We now discuss graphically the key argument in the proof. In the horizontal axis the figure reports the set of states, while in the vertical axis it reports the actions. For each θ, the blue 45-degree line represents the preferred action for the agent while the red line reports the preferred actions for the principal. The green set in the vertical axis represents an example of delegation set, which contemplates a 'hole'. Consider two actions y_1 and y_2 with $y_1 < y_2$. If all actions in the interval $[y_1, y_2]$ belong to the delegation set and the realized state θ falls into that interval, the agent chooses his ideal action $y = \theta$. This is represented by the dotted line. If the principal excludes actions (y_1, y_2) from the delegation set, the agent cannot take his preferred action but chooses the one closest to his bliss point. Hence, in states below the midpoint $\frac{y_1+y_2}{2}$ the agent chooses y_1, while in states above the midpoint he chooses y_2. Given that the principal's ideal action lies strictly below the ideal action of the agent, this implies that in states below $\frac{y_1+y_2}{2}$, the implemented action moves closer to the principal's bliss point, whereas in states above $\frac{y_1+y_2}{2}$ it moves further away. Since the cost of moving away from the bliss point is convexly increasing in the distance, the principal's loss outweighs the gain, as long as the probability weight attached to the states below $\frac{y_1+y_2}{2}$ is not too large. The first condition in Assumption 1 - $\beta f'(\theta) + f(\theta) > 0$ - assures that this is indeed the case and hence the principal never finds it optimal to leave 'holes' in the delegation set.

condition $\mathbb{E}[\theta - \beta] > 0$ implies that at the point $\bar{y} = 0$ the derivative of the objective function is positive. It is easy to see that for $\bar{y} < 0$ the derivative of the objective function is larger that at $\bar{y} = 0$ so no $\bar{y} \in [y_{min}, 0)$ can be a solution to the first order condition. The second derivative of the objective function is:

$$\beta f(\bar{y}) - (1 - F(\bar{y})).$$

Assumption 1 implies that the third derivative of the objective function is positive, implying strict convexity of the first derivative. This implies that the first derivative can cross zero at most at another point. It can be checked by direct inspection of (4) that the other point where the first order condition is satisfied is $\bar{y} = 1$. Since $\hat{y} < 1$, by convexity of the first derivative, it must be that for $\bar{y} < 1$ and close enough to 1 the first derivative is negative. So $\bar{y} = 1$ cannot be a maximum. What remains is the interior $\hat{y}$ as claimed.

We now shown the monotonicity of $\hat{y}$ in β. Since a maximum exists, it must be that the following second order necessary condition is satisfied at $\hat{y}(\beta)$ (note we now write explicitly the dependence of $\hat{y}$ on β):

$$\beta f(\hat{y}(\beta)) - (1 - F(\hat{y}(\beta))) \leq 0. \tag{5}$$

By the strict convexity of the first derivative and the fact that it equals zero from below at $\bar{y} = 1$ it must be at $\hat{y}$ the first derivative crosses zero from above at it cannot be flat around $\hat{y}$. Condition (5) must hence be satisfied with strict inequality. We can hence use the implicit function theorem to (4) (Assumption 1 implies that the cumulate F is $\mathcal{C}^1$) to show that $\hat{y}(\beta)$ admits a derivative at each β, which equals:

$$\hat{y}'(\beta) = -\frac{1 - F(\hat{y}(\beta)}{\beta f(\hat{y}(\beta)) - (1 - F(\hat{y}(\beta)))} < 0,$$

where we used the necessary second order condition (5) with strict inequality. Continuous differentiability is guaranteed by the implicit function theorem and can be checked directly in the above expression.

A.2 Proof of Proposition 2

See Proposition 3 and Section 3.4.

A.3 Proof of Proposition 3

The proof is presented as three lemmas.

- **Lemma 1**: *Consider $y_1, y_2 \in Y$ with $y_1 < y_2$. If $y_1, y_2 \in D^*(Y)$, then all $y \in (Y \cap (y_1, y_2))$ belong to $D^*(Y)$.*

Proof. Towards a contradiction suppose there is some $y \in Y$ such that $y \notin D^*(Y)$ and $D^*(Y) \cap [y_{min}, y] \neq \emptyset$, $D^*(Y) \cap [y, y_{max}] \neq \emptyset$. Further, let y^- be the largest element of $D^*(Y)$ strictly smaller than y and let y^+ be the smallest element of $D^*(Y)$ strictly greater than y, that is $y^- = \max\{y' \in D^*(Y) : y' < y\}$ and $y^- = \min\{y' \in D^*(Y) : y' > y\}$. Define $s \equiv \frac{y^- + y^+}{2}$ to be the state at which the agent is indifferent between choosing action y^- and action y^+, and similarly define $r \equiv \frac{y + y^-}{2}$ and $t \equiv \frac{y^+ + y}{2}$ as the states in which the agent is indifferent, respectively, between choosing y^- and y and between y^+ and y.

Following Alonso and Matouschek (2008), we can write the change in the principal's expected payoff when including action y into the delegation set. The agent changes his choice of action only in states $[r, t]$. In states $[r, s]$ he switches from y^- to y, while in the remaining states $(s, t]$ he switches from y^+ to y. The change in the principal's expected payoff is thus given by

$$\begin{aligned} & -\int_r^t (y - \theta + \beta)^2 f(\theta) \mathrm{d}\theta + \int_r^s (y^- - \theta + \beta)^2 f(\theta) \mathrm{d}\theta + \int_s^t (y^+ - \theta + \beta)^2 f(\theta) \mathrm{d}\theta, \\ = \; & 2(y - y^-) \underbrace{F(r)\left[r - \mathbb{E}\left[\theta - \beta | \theta \leq r\right]\right]}_{=T(r)} + 2(y^+ - y) \underbrace{F(t)\left[t - \mathbb{E}\left[\theta - \beta | \theta \leq t\right]\right]}_{=T(t)} \\ & -2(y^+ - y^-) \underbrace{F(s)\left[s - \mathbb{E}\left[\theta - \beta | \theta \leq s\right]\right]}_{=T(s)}. \end{aligned}$$

Letting $y = \lambda y^+ + (1 - \lambda) y^-$ for some $\lambda \in (0, 1)$ so that $y - y^- = \lambda(y^+ - y^-)$, $y^+ - y = (1 - \lambda)(y^+ - y^-)$ and $s = \lambda r + (1 - \lambda) t$, the payoff difference can be written as

$$2(y^+ - y^-)\left[\lambda T(r) + (1 - \lambda) T(t) - T(\lambda r + (1 - \lambda) t)\right].$$

From the strict convexity of T, it then follows that the payoff difference is strictly positive. A contradiction. □

- **Lemma 2**: *The optimal delegation set satisfies* $\min D^*(Y) = \min Y$.

Proof. Consider delegation set D with $\min D(Y) > \min Y$. Letting $y = \min Y$ and $\underline{y} = \min D(\widehat{y})$, the state at which the agent is indifferent between the two actions is given by $s \equiv (y + \underline{y})/2$. If the principal includes y in the delegation set, the agent switches from $\underline{y}$ to y in all states $\theta \leq s$. The principal's change in expected payoff when including y is hence given by

$$\begin{aligned}
& -\int_0^s (y-\theta+\beta)^2 f(\theta)\mathrm{d}\theta + \int_0^s (\underline{y}-\theta+\beta)^2 f(\theta)\mathrm{d}\theta, \\
= & \int_0^s \left[(\underline{y}-y)(\underline{y}+y) - 2(\underline{y}-y)(\theta-\beta)\right] f(\theta)\mathrm{d}\theta, \\
= & \; 2(\underline{y}-y)T(s),
\end{aligned}$$

which is strictly positive. Including y in the delegation set therefore strictly increases the principal's payoff, which implies $\min D^*(Y) = \min Y$. □

- **Lemma 3**: *Let Assumption 1 be satisfied. The optimal delegation set is such that*

$$\max D^*(Y) = \arg\min_{y\in Y} |y-\hat{y}|.$$

Proof. Consider delegation set D and suppose $\max D < \max Y$. Let $\overline{y} = \max D$ and consider action $y > \overline{y}, y \in Y$. Let $t = \frac{y+\overline{y}}{2}$ denote the state at which the agent is indifferent between the two actions. The change in the principal's payoff when including action y is given by

$$\begin{aligned}
& -\int_t^1 (y-\theta+\beta)^2 f(\theta)\mathrm{d}\theta + \int_t^1 (\overline{y}-\theta+\beta)^2 f(\theta)\mathrm{d}\theta, \\
= & -\int_t^1 \left[(y-\overline{y})(y+\overline{y}) - 2(y-\overline{y})(\theta-\beta)\right] f(\theta)\mathrm{d}\theta, \\
= & -2(y-\overline{y})S(t).
\end{aligned}$$

This change is weakly positive if and only if $S(t) \leq 0$ and hence if and only if $t \leq \hat{y}$. Since t is the midpoint of $\overline{y}$ and y, this condition holds if and only if the distance between $\overline{y}$ and $\hat{y}$ is weakly greater than the distance between y and $\hat{y}$, i.e. $|\overline{y}-\hat{y}| \geq |y-\hat{y}|$. □

The previous results together conclude the proof. □

A.4 Proof of Proposition 4

Let the agent's payoff as a function of Δ be defined by (recall that for $\theta \leq \hat{y}$ the payoff equal is zero):

$$U(\Delta) = -\int_{\hat{y}-\Delta}^{\hat{y}} (\hat{y}-\Delta-\theta)^2 \mathrm{d}F(\theta) - \int_{\hat{y}}^1 (\hat{y}+\Delta-\theta)^2 \mathrm{d}F(\theta). \tag{6}$$

The first and second derivative of $U(\Delta)$ are

$$U'(\Delta) = 2\int_{\hat{y}-\Delta}^{\hat{y}} [\hat{y}-\Delta-\theta]\,\mathrm{d}F(\theta) - 2\int_{\hat{y}}^{1} [\hat{y}+\Delta-\theta]\,\mathrm{d}F(\theta), \tag{7}$$

$$U''(\Delta) = -2[1-F(\hat{y}-\Delta)] < 0. \tag{8}$$

The function $U(\Delta)$ is strictly concave in Δ and hence has a unique solution on $[0,\bar{\Delta}(Y)]$. The interior solution of the agent's optimization problem, Δ^*, is characterized by the first-order condition that equalizes the expression (7) to zero. The expression in (3) is obtained after simple rearrangements of the terms in (7).

A.5 Proof of Proposition 5

First, let us write the agent's payoff as a function of Δ and the parameter β:

$$U(\Delta;\beta) = -\int_{\hat{y}(\beta)-\Delta}^{\hat{y}} (\hat{y}(\beta)-\Delta-\theta)^2 f(\theta)\mathrm{d}\theta - \int_{\hat{y}(\beta)}^{1} (\hat{y}(\beta)+\Delta-\theta)^2 f(\theta)\mathrm{d}\theta.$$

Recall, the solution to the problem solves

$$\int_{\hat{y}(\beta)-\Delta^*}^{\hat{y}(\beta)} [\hat{y}(\beta)-\Delta^*-\theta]\, f(\theta)\mathrm{d}\theta - \int_{\hat{y(\beta)}}^{1} [\hat{y}(\beta)+\Delta^*-\theta]\, f(\theta)\mathrm{d}\theta = 0.$$

Since F is $\mathcal{C}^1$ by Assumption 1, $\hat{y}(\beta)$ is $\mathcal{C}^1$ from 1, and $U''_{\Delta,\Delta} < 0$, the conditions for applying the implicit function theorem are satisfied. There is hence a function $\Delta^*(\beta)$ describing the unrestricted solution for the agent that solves the first order condition: $U'_{\Delta}(\Delta^*(\beta);\beta) = 0$, which becomes an identity when seen as a function of β, and:

$$\Delta^{*\prime}(\beta) = -\frac{U''_{\Delta,\beta}(\Delta^*(\beta);\beta)}{U''_{\Delta,\Delta}(\Delta^*(\beta);\beta)}.$$

The statement in the proposition will hence be shown if we can prove that $U''_{\Delta,\beta}(\Delta^*(\beta);\beta) > 0$. Differentiating the expression of the first order condition (3) with respect to β keeping Δ^* as fixed, after some rearrangement, delivers:

$$U''_{\Delta,\beta}(\Delta^*(\beta);\beta) = -\hat{y}'(\beta)\left[1+F\left(\hat{y}(\beta)-\Delta^*(\beta)\right)-2F(\hat{y}(\beta))\right].$$

Since $\hat{y}'(\beta) < 0$ from Proposition 1, we would be done if $1+F\left(\hat{y}(\beta)-\Delta^*(\beta)\right)-2F(\hat{y}(\beta)) > 0$. Note that this inequality can be equivalently written as:

$$2(1-F(\hat{y}(\beta))) > 1-F\left(\hat{y}(\beta)-\Delta^*(\beta)\right).$$

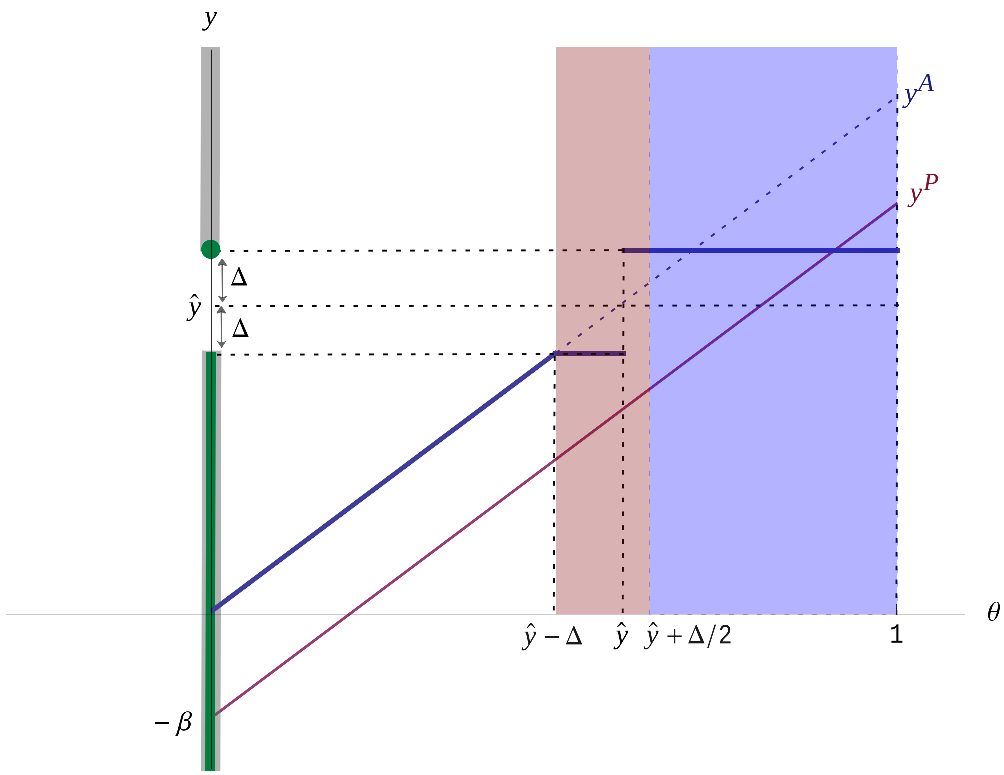

Figure 7: In the horizontal axis the figure reports the set of states, while in the vertical axis it reports the feasible actions. The green set in the vertical axis represents a typical equilibrium delegation set. For each θ, the blue 45-degree line represents the preferred action for the agent while the red line reports the preferred actions for the principal. For states from $\hat{y} + \frac{\Delta}{2}$ the agent gains as the new pooling action is uniformly closer to his ideal point compared to $\hat{y}$. This is represented by the blue area in the figure. The cost of increasing the gap is the utility loss in the states $\left[\hat{y} - \Delta, \hat{y} + \frac{\Delta}{2}\right]$, where the agent moves away from his ideal action. The states implying a loss compared to $\hat{y}$ are represented by the red area. Note that for $\Delta \approx 0$ the red area vanishes. Since at $\hat{y}$ the agent chooses his bliss point action, by increasing the gap he enjoys first order gains while losses are zero to the first order. The unrestricted optimal Δ^* equalizes marginal gains with marginal losses maximizing the overall net gain.

Now, if we use $\hat{y}(\beta) = \mathbb{E}[\theta - \beta | \theta \geq \hat{y}(\beta)]$ and rearrange the first order condition, we obtain:

$$[1 - F(\hat{y}(\beta) - \Delta^*(\beta))] \left[\mathbb{E}[\theta | \theta \geq \hat{y}(\beta) - \Delta^*(\beta)] - (\hat{y}(\beta) - \Delta^*(\beta))\right] = 2[1 - F(\hat{y}(\beta))]\beta.$$

Since $\hat{y}(\beta) - \Delta^*(\beta) < \hat{y}(\beta)$ from the definition of $\hat{y}(\beta)$, it must be that

$$\mathbb{E}[\theta | \theta \geq \hat{y}(\beta) - \Delta^*(\beta)] - (\hat{y}(\beta) - \Delta^*(\beta)) > \beta,$$

which implies $(1 - F(\hat{y}(\beta) - \Delta^*)\beta))) < 2(1 - F(\hat{y}(\beta)))$ as desired.

A.6 Proof of Proposition 6

Suppose that after the initial stage of revelation the principal's awareness is Y and the optimal delegation set is D. For later use, recall from Proposition 3:

$$y \in D \iff y \leq \arg\min_{y \in Y} |y - \hat{y}|.$$

We require Bayesian perfection and hence will check for optimality given Y and D,[18] and focus on the best equilibrium for the agent.

The game played at the renegotiation stage is a signalling game. We start by defining the strategies of the principal and the agent at this stage. We will first of all concentrate on agent's moves constituted by either 'no new proposal' (let us call it $\varnothing$) or proposal of singletons $y \notin D$. Let $X \equiv \{\varnothing\} \cup Y^A \setminus D$ be the set of possible proposals of the agent. The $\varnothing$ proposal will be interpreted as if the agent does not propose any action outside the original delegation set and hence in particular s/he does not increase the awareness of the principal. When the agent proposes a new action $y \notin D$ the principal can decide wether to replace the original delegation set D with y. It will be clear that the agent has no gain (and potential losses) in choosing more complicated strategies compared (such as proposing sets of actions) to the equilibrium we are proposing. Moreover, there is no gain in terms of payoffs in augmenting the proposal/signal x with a further message. Recall indeed that the agent can choose any action within the delegation set, this will be done at the terminal node of the game and determine the payoffs of both players.

Given D, Y, we can describe the agent's strategy as a map from each state of the world $\theta \in [0, 1]$ into a recommendation $x \in X$ and then a map from the new delegation set D' into an action $y \in D'$. In order to simplify notation, we will not indicate explicitly the dependence on (D, Y) of the objects that follow. The agent's strategy is thus a pair of functions $x(\cdot), y(\cdot, \cdot)$ where $x(\theta)$ is the proposal given θ and $y(\theta, D')$ is the final choice given (θ, D'). The principal's strategy is a mapping $\rho : Y^A \setminus D \to \{0, 1\}$, where $\rho(x) = 0$ implies $D' = D$ while $\rho(x) = 1$ implies $D' = x$. With his/her action, the principal assigns to each proposal a delegation set which can be either the original D or the singleton $\{x\}$.

We also concentrate on equilibria in pure strategies. We can hence define a belief function for the principal, T, mapping each proposal x to a subset of $[0, 1]$. The set T describes the states which the principal considers possible when the agent proposes x. The strategy profile (x^*, y^*, ρ^*) and belief function T^* constitute a pure-strategy Perfect Bayesian Equilibrium of the renegotiation game if and only if the following conditions hold:[19]

[18] Clearly, in the mechanism design interpretation of the model, the set D corresponds to 'the original contingent contract' the principal committed to after the initial updating leading to Y.

[19] All expectations are taken with respect to F.

- Principal optimality: for all $x \in Y^A \setminus D$,

$$\rho^*(x) \in \arg\max_{\rho \in \{0,1\}} \rho \mathbb{E}[v(x,\theta)\,|\theta \in T(x)] + (1-\rho)\mathbb{E}[v(y^*(D,\theta),\theta)\,|\theta \in T(x)];$$

- Agent optimality: for all $\theta \in [0,1]$,

$$x^*(\theta) \in \arg\max_{x} \rho^*(x)u(x,\theta) + (1-\rho^*(x))u(y^*(D,\theta),\theta)\,; \tag{9}$$

$$y^*(D,\theta) \in \arg\max_{y \in D} u(y,\theta); \tag{10}$$

- Consistency of beliefs: for all x that are allowed on path,

$$T^*(x) = \{\theta \in [0,1] : u(x,\theta) \geq \max_{x' \in Y} \rho^*(x')u(x',\theta) + (1-\rho^*(x'))u(y^*(D,\theta),\theta)\}.$$

The consistency condition assures that the principal's beliefs are coherent with the agent's equilibrium actions, as perceived through the principal's awareness (see that the max is taken only over Y). According to the condition, the principal evaluates payoffs for lower levels of awareness using her own equilibrium strategy. Hence, with respect to lower levels of awareness that are reached on path, the principal's beliefs are correct.

To maximise the agent's payoff, we focus on the equilibrium where the agent is able to choose as many actions as possible. We will construct a particular equilibrium where all actions y such that $y \notin D$ and $y \leq \max D$ are proposed by the agent whenever beneficial to him conditional on the state and are allowed by the principal.[20] We will also show that there is no equlibrium where the agent will be allowed to choose any $y > \max D$. The proposed equilibrium has the following strategies:

$$x^*(\theta) = \begin{cases} \theta & \text{if} \quad \theta \notin D \;\&\; \theta < \max D; \\ \varnothing & \text{otherwise.} \end{cases}$$

$$\rho^*(x) = \begin{cases} 1 & \text{if} \quad x \notin D \;\&\; x < \max D; \\ 0 & \text{otherwise.} \end{cases}$$

$$y^*(\theta, D') \text{ solves (10) for all } \theta, D'.$$

[20] This equilibrium will also satisfy the 'reasoning refinement' proposed by Ozbay (2008) for limited awareness in contingencies. Within our framework, the 'reasoning refinement' restricts the principal's beliefs over Θ when the agent proposes an action y not included in the awareness set of the principal at the renegotiation stage. It does so to induce the principal to let the agent take the proposed action whenever there is a probability distribution that both 'justifies' the agent's proposal (i.e., y is maximal for the agent according to principal's awareness) and—if chosen—improves on principal's utility compared to the status quo.

The equilibrium belief function is described below. Clearly, given ρ^*, the agent is indifferent between proposing and not proposing $x > \max D$. To compute T^* and hence complete the verification that the proposed strategies constitute and equilibrium, it will be convenient to consider three cases.

a) Consider first the possibility that after learning the state of the world the agent proposes a new action $x \notin D$ and $x \in (\min D, \max D)$ (so in particular $x \neq \varnothing$). Let y^- be the greatest element of D smaller than x and y^+ be the smallest element of D greater than x. Optimality of D with respect to the principal's awareness Y implies that there are no actions within the principal's awareness between y^- and y^+ so the mentioned proposal expands the awareness of the principal and there is not other action $y' \in (y^-, y^+)$ such that $y' \in Y$. When the agent proposes action x, if allowed in equilibrium, the principal learns that the state lies between $(\tilde{y}+y^-)/2$ and $(\tilde{y}+y^+)/2$. That is, $T^*(x) = \left(\frac{x+y^-}{2}, \frac{x+y^+}{2}\right)$ for such x. The proof of Lemma 1 in the proof of Proposition 3 shows that conditional on these states, the principal prefers to include the action x into the delegation set. That is, indeed we have an equilibrium where $\rho^*(x) = 1$ for all $x \in (\min D, \max D)$.

b) Next, suppose the agent proposes an action $x < \min D$. Since D is optimal with respect to Y, there is no element in Y smaller than $\min D$. So again, this move expands the awareness of the principal. The fact that the agent proposes action x, if allowed, it would induce the principal to conclude that $\theta < (x + \min D)/2$, that is $T^*(x) = \left[0, \frac{x+\min D}{2}\right)$ for such x. Since the principal's preferred action is strictly smaller than the agent's preferred action, this condition assures that the principal weakly prefers x over $\min D$. Hence again we have an equilibrium where $\rho^*(x) = 1$ in all these cases where $x < \min D$.[21]

c) Finally, suppose the agent proposes some action $x > \max D$.
(i) Assume first that there are no actions in Y that are strictly greater than $\max D$. In this case, if x were allowed, the principal learns that $\theta > (x + \max D)/2$, i.e., $T(x) = (\frac{x+\max D}{2}, 1)$. Allowing the agent to choose x under these beliefs is optimal if

$$\mathbb{E}[-(\max D - \theta + \beta)^2 | \theta > (x + \max D)/2] \leq \mathbb{E}[-(x - \theta + \beta)^2 | \theta > (x + \max D)/2].$$

This inequality can be rewritten as

$$(x - \max D)\left((x + \max D)/2 - \mathbb{E}[\theta - \beta | \theta > (x + \max D)/2]\right) \leq 0,$$

which is satisfied if and only if $(x + \max D)/2 \leq \hat{y}$, that is, if and only if the distance between x and $\hat{y}$ is weakly smaller than that between $\max D$ and $\hat{y}$. Hence in this case,

[21] Clearly, if $\min D \leq 0$ all arguments we made apply in a weak sense.

$\rho^*(\tilde{y}) = 1$ can be consistent with an equilibrium only if $(x + \max D)/2 \le \hat{y}$.

(ii) Now suppose the principal is aware of some additional action $y' \neq x$ greater than $\max D$. Given that those actions will never be chosen by the agent (as the principal can either allow D or x), the inference the principal makes by allowing x will be the same as above, namely $T(x) = (\frac{x+\max D}{2}, 1)$. We can hence apply the previous argument that the principal will allow such action only if $(x + \max D)/2 \le \hat{y}$ holds as well.

Taken together, this implies that there is an equilibrium with the following outcome. Given a delegation set D that is optimal with respect to the principal's awareness after the first updating, Y, at the renegotiation stage the principal permits any action that is smaller than $\max D$ and any action that is larger than $\max D$ as long as its distance to $\hat{y}$ is smaller than that between $\max D$ and $\hat{y}$. Since $\max D$ is the action in Y that is closest to $\hat{y}$, the set of actions implemented at the renegotiation stage is solely restricted by the action in the principal's awareness Y with the smallest distance to $\hat{y}$. The larger this distance is, the larger is the set of actions the agent is able to implement after learning the state of the world. This implies that at the outset of the game it is optimal for the agent to reveal a set Y in the class of awareness sets parametrised by a gap Δ, as in the game without renegotiation. The optimal value of Δ solves the following optimization problem:

$$\max_{\Delta \le \bar{\Delta}(Y^P)} \qquad U_\sigma(\Delta) \equiv -\sigma \int_{\hat{y}+\Delta}^{1} (\hat{y} + \Delta - \theta)^2 dF(\theta) + (1-\sigma)U(\Delta).$$

where $U(\Delta)$ is defined by (6). When $\sigma = 0$, we renegotiation plays no role and U_σ is maximized by $\min\{\Delta^*, \bar{\Delta}(Y^P)\}$, as we established in Proposition 4. We will thus consider the case $\sigma > 0$. We have

$$\begin{aligned} U'_\sigma(\Delta) &= -2\sigma \int_{\hat{y}+\Delta}^{1} \underbrace{(\hat{y} + \Delta - \theta)}_{<0} dF(\theta) + (1-\sigma)U'(\Delta), \\ U''_\sigma(\Delta) &= -2\sigma(1 - F(\hat{y} + \Delta)) + (1-\sigma)U''(\Delta). \end{aligned}$$

Since $U''(\Delta) < 0$ (see Section A.4), we have $U''_\sigma < 0$, so the problem is strictly concave. The unconstrained solution of the optimization problem is characterized by the first-order condition

$$2\sigma \int_{\hat{y}+\Delta}^{1} (\hat{y} + \Delta - \theta) dF(\theta) = (1-\sigma)U'(\Delta) \tag{11}$$

Let Δ^{**} denote the value of Δ solving condition (11). Since the left-hand-side is strictly negative, (11) requires $U'(\Delta^{**}) < 0$ and hence $\Delta^{**} > \Delta^*$. The agent's optimisation problem is then solved by $\min\{\Delta^{**}, \bar{\Delta}(Y^P)\}$. Finally, it is easy to that $U'_\sigma(\Delta) < U'_{\sigma'}(\Delta)$ for all $\sigma < \sigma'$ and $\Delta > \Delta^*$. Letting $\Delta^{**}(\sigma)$ be the unconstrained solution for the parameter σ, this directly implies $\Delta^{**}(\sigma) < \Delta^{**}(\sigma')$. Notice that for $\sigma = 1$, we have $U'_\sigma(\Delta) > 0$ for all $\Delta \leq 1 - \hat{y}$, so $\Delta^{**}(1) = \min\{1 - \hat{y}, \bar{\Delta}(Y^P)\}$. Given that $U'_\sigma(\Delta)$ is continuous in σ, the statement of Proposition 6 follows. □

A.7 Proof of Proposition 7

Define $\hat{y}_P \equiv \arg\min_{y \in Y^P}\{|y - \hat{y}|\}$. We start by showing that when the principal's initial awareness is Y^P, there is no equilibrium where the principal takes an action strictly greater than $\hat{y} + \Delta(Y^P)$. To see this, fix a pure strategy equilibrium and let $\bar{y}$ be the largest action taken in equilibrium. Towards a contradiction, suppose $\bar{y} > \hat{y} + \Delta(Y^P)$. We then have $(\hat{y}_P + \bar{y})/2 > \hat{y}$. By Assumption 1, conditional on $\theta \geq (\hat{y}_P + \bar{y})/2$, the principal prefers action $\hat{y}_P$ over $\bar{y}$. Hence, in order for the principal to prefer $\bar{y}$, there must exist some $t > (\hat{y}_P + \bar{y})/2$ such that upon being proposed action $\bar{y}$, the principal conditions on the realized state being weakly greater than t. This requires that when being proposed $\bar{y}$, the principal is also aware of an action $a = 2t - \bar{y}$ such at state t the agent is indifferent between a and $\bar{y}$.[22] But since $t > (\hat{y}_P + \bar{y})/2 > \hat{y}$, we have $(a + \bar{y})/2 > \hat{y}$. Hence, a is closer to $\hat{y}$ than $\bar{y}$, which means that conditional on the state being greater than t, the principal prefers action a over $\bar{y}$, a contradiction.

Having proved $\bar{y} \leq \hat{y} + \Delta(Y^P)$, we want to show next that there is an awareness set $\tilde{Y}^P \subseteq Y^P$ that gives rise to an equilibrium under which both the principal and the agent are weakly better off. The following lemma will be useful

Lemma 11. *Let the principal's awareness set be the singleton $\{y\}$. Then for all $y^* \in [y, y + |y - \hat{y}|]$, there is an equilibrium where the agent takes his preferred action $y = \theta$ for all $\theta < y^*$ and action y^* for all $\theta \geq y^*$.*

Proof. Consider the following strategies and beliefs. Upon observing state θ, the agent discloses (and recommends) his preferred action θ if $\theta < y^*$ and action y^* otherwise. When the state is $\theta = y$ and the agent reveals no additional action, the principal has no choice but to take action y. For the remaining awareness sets and messages, beliefs are specified as follows. When the agent reveals (and recommends) some $y' \leq y^*$ with $y' \neq y$, the principal believes that the agent prefers y' over y and, hence, conditions on the realized state being closer to y' than to y. By Proposition 3, we know that the principal finds it optimal to take the newly revealed action y' if and only if $y' \leq y + |y - \hat{y}|$. Since $y^* \leq y + |y - \hat{y}|$, this requirement is satisfied on path. Off path beliefs for the principal can be specified as follows. When the agent reveals a single action $y' > y^*$ or a non-singleton set $Y \subseteq [y_{min}, y_{max}]$, the principal believes that the realized state is $\theta = y + \beta$ and thus optimally takes action y. □

[22] Note that the agent cannot be indifferent among more than two actions since he knows the state.

Suppose first $\bar{y} \notin (\hat{y} - \Delta(Y^P), \hat{y} + \Delta(Y^P))$ and consider initial awareness set $\{\hat{y}_P\}$. By Lemma 11 there exists a pure-strategy equilibrium in which the set of on-path equilibrium actions is $[0, \hat{y}+\Delta(Y^P)]$. Since $\bar{y} \leq \hat{y}+\Delta(Y^P)$, this set includes the set of equilibrium actions under awareness Y^P and, hence, makes the agent better off. As to the principal, Proposition 3 shows that the principal prefers delegation set $[0, \hat{y} - \Delta(Y^P)] \cup \{\hat{y} + \Delta(Y^P)\}$ over any subset of $[0, \hat{y} - \Delta(Y^P)] \cup \{\hat{y} + \Delta(Y^P)\}$. Moreover, the proof of Proposition 3 shows that the principal can increase her expected payoff by closing gaps in the delegation set. Hence, the principal prefers delegation set $[0, \hat{y}+\Delta]$ over $[0, \hat{y}-\Delta(Y^P)] \cup \{\hat{y}+\Delta(Y^P)\}$. Hence, also the principal is better off.

Next, suppose $\bar{y} \in (\hat{y}-\Delta(Y^P), \hat{y}+\Delta(Y^P))$. Let Y^D denote the set of actions the principal takes in the equilibrium with initial awareness Y^P after no additional action is revealed. By definition of $\bar{y}$, we have $\bar{y} \geq \max Y^D$ and $|\max Y^P - \hat{y}| \geq |\bar{y} - \hat{y}|$. Consider the situation where the principal's initial awareness set is $\{\max Y^D\}$. By Lemma 11 there exists a pure-strategy equilibrium in which the set of equilibrium actions is $[0, \bar{y}]$. This set of actions has the same minimum and maximum as the set of equilibrium actions under Y^P but has no gaps. As a consequence, both the agent and the principal are better off. $\square$

A.8 Proof of Proposition 8

Consider a neologism proof equilibrium for the case where the principal's initial awareness is Y^P and let $\bar{y}$ be the highest action taken in equilibrium. In the proof of the previous proposition, we have shown $\bar{y} \leq \hat{y} + \Delta(Y^P)$. We will now show that in a neologism proof equilibrium $\bar{y} = \hat{y} + \Delta(Y^P)$.

Define $\bar{y}_P \equiv \max_{m\in\mathcal{M}} \rho^*(Y^P, m)$ as the highest action the agent can induce when revealing no additional actions. Since $\bar{y} \leq \hat{y}+\Delta(Y^P)$, we have $\bar{y}_P \leq \hat{y}_P$ (recall that $\hat{y}_P$ is the action in Y^P closest to $\hat{y}$, as defined in Section A.7). When the latter inequality is strict, there exists an action $y \in Y^P$ with $y \in (\bar{y}_P, \hat{y}_P]$ such that the pair (Y^P, G) with $G = \{\theta : \theta \geq (\bar{y}_P+y)/2\}$ is self-signalling. Hence $\bar{y}_P = \hat{y}_P$. When $\hat{y}_P \geq \hat{y}$, the property $\bar{y} = \hat{y} + \Delta(Y^P)$ follows immediately. For the other case ($\hat{y}_P < \hat{y}$), consider the pair $(Y^P \cup \{\hat{y} + \Delta(Y^P)\}, G')$ with $G' = \{\theta : \theta > \hat{y}\}$ and $\hat{y} + \Delta(Y^P) = BR(Y^P \cup \{\hat{y} + \Delta(Y^P)\}, G')$. If $\bar{y} \neq \hat{y} + \Delta$, then

$$\begin{aligned} & \{\theta : u(\hat{y} + \Delta(Y^P), \theta), \theta) > u^*(Y^P \cup \{\hat{y} + \Delta\}, \theta)\} \\ = & \{u(\hat{y} + \Delta(Y^P), \theta), \theta) > u(\hat{y} - \Delta(Y^P), \theta)\} \\ = & G', \end{aligned}$$

so the pair $(Y^P \cup \{\hat{y} + \Delta\}, G')$ is self-signalling. Hence, $\bar{y} = \hat{y} + \Delta(Y^P)$.

Next, we want to show next that there is an awareness set that gives rise to a neologism-proof equilibrium under which both the principal and the agent are weakly better off. Consider initial awareness set $\{\hat{y}_P\}$. By Lemma 11 there is an equilibrium, where the agent discloses and recommends his preferred action $y = \theta$ in all states $\theta \leq \hat{y} + \Delta(Y^P)$. This equilibrium is also neologism proof, as we show next. Strategies and beliefs can be specified as follows $\rho^*(Y, m) = \max \arg\min_{y \in Y} |y - \hat{y}|$, $Y^*(\theta) = \hat{y}_P \cup \{\theta\}$ if $\theta < \hat{y} + \Delta(Y^P)$ and $Y^*(\theta) = \hat{y} + \Delta(Y^P)$ if $\theta \geq \hat{y} + \Delta(Y^P)$.[23] Towards a contradiction, suppose there exists a pair (Y, G) that is self-signaling. In this case, we have $BR(Y, G) > \hat{y} + \Delta$ and $G = \{\theta : u(BR(Y, G), \theta) > u^*(Y, \theta)\}$. By definition of ρ^* we further have $|\rho^*(Y', m) - \hat{y}| \leq \Delta(Y^P) < |BR(Y, G) - \hat{y}|$ for all Y' including $\hat{y}_P$. Conditioning on $\theta > (\rho^*(Y', m) + BR(Y, G))/2$ the principal thus strictly prefers $\rho^*(Y', m)$ over $BR(Y, G)$, contradicting the assumption that (Y, G) is self-signalling.

Finally, we want to argue that the described equilibrium weakly improves the principal's and agent's expected utility. The equilibrium actions under initial awareness set Y^P are described by a set with maximum $\hat{y}_P + \Delta$. The set of equilibrium actions under initial awareness $\{\hat{y}_P\}$ is the interval $[0, \hat{y}_P + \Delta]$. Proposition 3 shows that the principal prefers delegation set $[0, \hat{y}_P + \Delta]$ over any subset of that interval, hence the equilibrium with initial awareness $\{\hat{y}_P\}$ yields a higher expected payoff for the principal. Since having more choice is better than having less, the same is true for the agent. □

A.9 Proof of Proposition 9

We want to show that revealing an awareness set of the form $[y_{min}, \hat{y} - \Delta] \cup [\hat{y} + \Delta, y_{max}]$ is optimal. Towards a contradiction, suppose this is not the case and let the optimal awareness set be denoted by Y. Define $\tilde{\Delta}$ to be the smallest value of Δ such that $Y \subseteq [y_{min}, \hat{y} - \Delta] \cup [\hat{y} + \Delta, y_{max}]$ and $\tilde{Y} = [y_{min}, \hat{y} - \tilde{\Delta}] \cup [\hat{y} + \tilde{\Delta}, y_{max}]$ to be the associated awareness set. Suppose the principal's realized awareness set is Y^P. According to Proposition 3, the induced delegation sets from revealing, respectively, Y and $\tilde{Y}$ are

$$
\begin{aligned}
D^*(Y \cup Y^P) &= \{y \in Y \cup Y^P : y \leq \arg \min_{y \in Y \cup Y^P} |y - \hat{y}|\}, \\
D^*(\tilde{Y} \cup Y^P) &= \{y \in \tilde{Y} \cup Y^P : y \leq \arg \min_{y \in \tilde{Y} \cup Y^P} |y - \hat{y}|\}.
\end{aligned}
$$

In order for Y to yield a strictly higher payoff for the agent than $\tilde{Y}$, there must exist some awareness set Y^P and some action y such that $y \in D^*(Y \cup Y^P)$ and $y \notin D^*(\tilde{Y} \cup Y^P)$. By its definition, $D^*(\tilde{Y} \cup Y^P)$ includes all actions in $\tilde{Y}$ weakly smaller than $\hat{y}$. Given $Y \subseteq \tilde{Y}$, it follows that $y > \hat{y}$. As well, the optimal delegation set includes at most one action strictly greater than $\hat{y}$. By definition of $\tilde{\Delta}$, the set Y includes an action whose distance to $\hat{y}$ is $\tilde{\Delta}$. This implies that the largest action in $D^*(Y \cup Y^P)$ is weakly smaller than $\hat{y} + \tilde{\Delta}$. Hence, we have $y \leq \hat{y} + \tilde{\Delta}$. Also, since y belongs to $D^*(Y \cup Y^P)$, it follows that there is no action in Y^P strictly closer to $\hat{y}$ than y. However, the property $|y - \hat{y}| \leq \tilde{\Delta}$, together with the fact that there is no action in Y^P that is closer to $\hat{y}$ than y, implies that y must also belong to

[23] The choice of m^* is not relevant, since the principal's best response only depends on the first argument.

$D^*(\tilde{Y} \cup Y^P)$. A contradiction.

□

A.10 Proof of Corollary 10

Let $EU(\Delta)$ denote the agent's expected payoff associated to the disclosure of a set of actions parametrised by Δ and let μ be the probability which the agent assigns to the event that the principal's awareness set does not include $\hat{y}$. Then, given our assumption that awareness sets are closed, there exists some $\varepsilon > 0$ such that no action in $(\hat{y} - \varepsilon, \hat{y} + \varepsilon)$ belongs to either of the principal's awareness sets in the support that do not contain $\hat{y}$. For $\Delta \in [0, \varepsilon]$, the agent's expected payoff conditional on facing a principal who is not aware of $\hat{y}$ is then described by the function $U(\Delta)$, as defined in (6). From Proposition 4 we know that $U'(0) = -2\int_{\hat{y}}^{1} (\hat{y} - \theta)\, \mathrm{d}F(\theta)$ is strictly positive.

With the complementary probability $1 - \mu$, the agent faces a principal who is aware of $\hat{y}$. In this event, the principal never permits an action greater than $\hat{y}$. A lower bound for the agent's payoff conditional on the principal being aware of $\hat{y}$ as a function of Δ is given by the payoff that obtains when the principal is unaware of all actions in $(\hat{y} - \Delta, \hat{y})$: any action in the principal's awareness set belonging to $(\hat{y} - \Delta, \hat{y})$ will be included in the delegation set and thus increases flexibility for the agent. The lower bound utility is:

$$\underline{U}(\Delta) = -\int_{\hat{y}-\Delta}^{\hat{y}-\Delta/2} (\hat{y} - \Delta - \theta)^2 f(\theta)\mathrm{d}\theta - \int_{\hat{y}-\Delta/2}^{1} (\hat{y} - \theta)^2 \mathrm{d}F(\theta),$$

with

$$\underline{U}'(\Delta) = 2\int_{\hat{y}-\Delta}^{\hat{y}-\Delta/2} (\hat{y} - \Delta - \theta)\mathrm{d}F(\theta)$$

The agent's unconditional expected payoff $EU(\Delta)$ must then be weakly greater than $\mu U(\Delta) + (1 - \mu)\underline{U}(\Delta)$. Noticing that $\underline{U}'(0) = 0$, the fact that $U'(0)$ is strictly positive implies that the first derivative of $\mu U(\Delta) + (1 - \mu)\underline{U}(\Delta)$ at $\Delta = 0$ is strictly positive. Since at $\Delta = 0$, $\mu U(\Delta) + (1 - \mu)\underline{U}(\Delta)$ is equal to $EU(\Delta)$, and it constitutes a lower bound for it elsewhere it follows that $EU(\Delta)$ is strictly increasing on a right neighbourhood of $\Delta = 0$.

□

B Generalized extensive-form games with unawareness

Heifetz, Meier and Schipper (2013, from now on HMS) define generalized extensive-form games that allow for evolving unawareness. In this section we show how our framework can be formalized as such a game. We will focus on the case where the principal is unaware of her unawareness but the framework of HMS is sufficiently flexible to also capture awareness of unawareness. To introduce the generalized extensive-form game Γ, let N be a set of decision nodes, C be a set of chance nodes, and Z be a set of terminal nodes. The nodes $\bar{N} = N \cup C \cup Z$ constitute a tree. HMS capture limited awareness via the notion of subtrees, defined as subsets of nodes of $\bar{N}$. Letting $\mathbf{T}$ be a family of subtrees of $\bar{N}$, for $T, T' \in \mathbf{T}$ the relation $T' \preccurlyeq T$ signifies that the nodes of T' constitute a subset of the nodes of T. One element of $\mathbf{T}$ represents the modeler's view of the paths of play that are objectively feasible. The other elements of $\mathbf{T}$ represent feasible paths of play as subjectively viewed by some players, or as the frame of mind attributed to a player by other players or by the same player at a later stage of the game.

HMS propose a number of natural properties for generalized extensive-form games. These include basic extensions of standard requirements of extensive form games but also new properties that are specific to unawareness. For instance, HMS require that at each information set, a player's anticipation of her future view of the game is confined to the view he currently holds. In particular, players have no expectation to forget currently conceivable paths. HMS further propose a number of requirements that mirror those of their static unawareness structures (Heifetz et al. 2006). The restrictions guarantee the coherence of the knowledge and awareness of players from more expressive game trees to subtrees of those. Letting n_T be a node in subjective tree $T \in \mathbf{T}$ and $n_{T'}$ be the copy of n_T in the less expressive tree $T' \preccurlyeq T$, the properties guarantee, for instance, that at $n_{T'}$ the player knows nothing he does not know at n_T. Similarly, at $n_{T'}$ the player is not aware of any moves he is unaware of at n_T. On the other hand, if at node n_T the player knows an event that is only based on nodes belonging to T', he also knows the event at node $n_{T'}$ in the less expressive tree T'. Likewise, at node $n_{T'}$ the player is aware of every event he is aware of at n_T, provided that these events are based on nodes belonging to T'. Finally, HMS postulate that awareness can only increase throughout the game. All of these features are satisfied in our unawareness game, which we describe in more detail now.

There are are three players: the principal, the agent, and nature. The set $\mathbf{T}$ contains the tree that provides the objective description of the game with all feasible paths of play. In our setting this description coincides with the agent's view of the game, depicted on the left of Figure 2. Let the associated game tree be denoted by T_{Y^A}. The other elements of $\mathbf{T}$ represent feasible paths of play as subjectively viewed by the principal at some node in T_{Y^A}, or as the frame of mind attributed to the principal or the agent at a node in T_{Y^A} by either the principal or the agent. We will denote T_Y the subtree associated to awareness Y.

Formally, T_Y is defined as the subtree that excludes all moves $X \nsubseteq Y$ and all nodes following those moves.

At the outset of the game, the principal's awareness is Y^P, so her subjective view of the game is described by the subtree T_{Y^P}, depicted on the left side of Figure 8. Within the confined view of T_{Y^P} the agent's announcement of investment projects in the first stage of the game is irrelevant for the rest of the game. In particular, before becoming aware of additional projects, the principal can only envisage the agent announcing subsets of Y^P. At the start of the game the principal therefore believes (falsely) that in the second stage of the game she can only choose a delegation set from the set of all subsets of Y^P, no matter what the agent announces.

Once the agent reveals additional projects to the principal, the principal updates her awareness and therefore her subjective view of the game. Given updated awareness Y, the principal's subjective game tree T_Y includes additional nodes, depicted on the right side of Figure 8. It should be noted that once the principal becomes aware of the additional nodes, she can also contemplate less expressive game trees $T_{\tilde{Y}}, Y^P \subseteq \tilde{Y} \subseteq Y$. That is, the principal can envisage how the game would have unfolded if the agent would have revealed fewer projects than he did. However, the principal cannot contemplate the paths of play that would have been feasible if the agent would have revealed more. After the initial stage, awareness no longer changes, which means that, given the constraints unawareness imposes on the players' succeeding moves, agent and principal play a standard game.

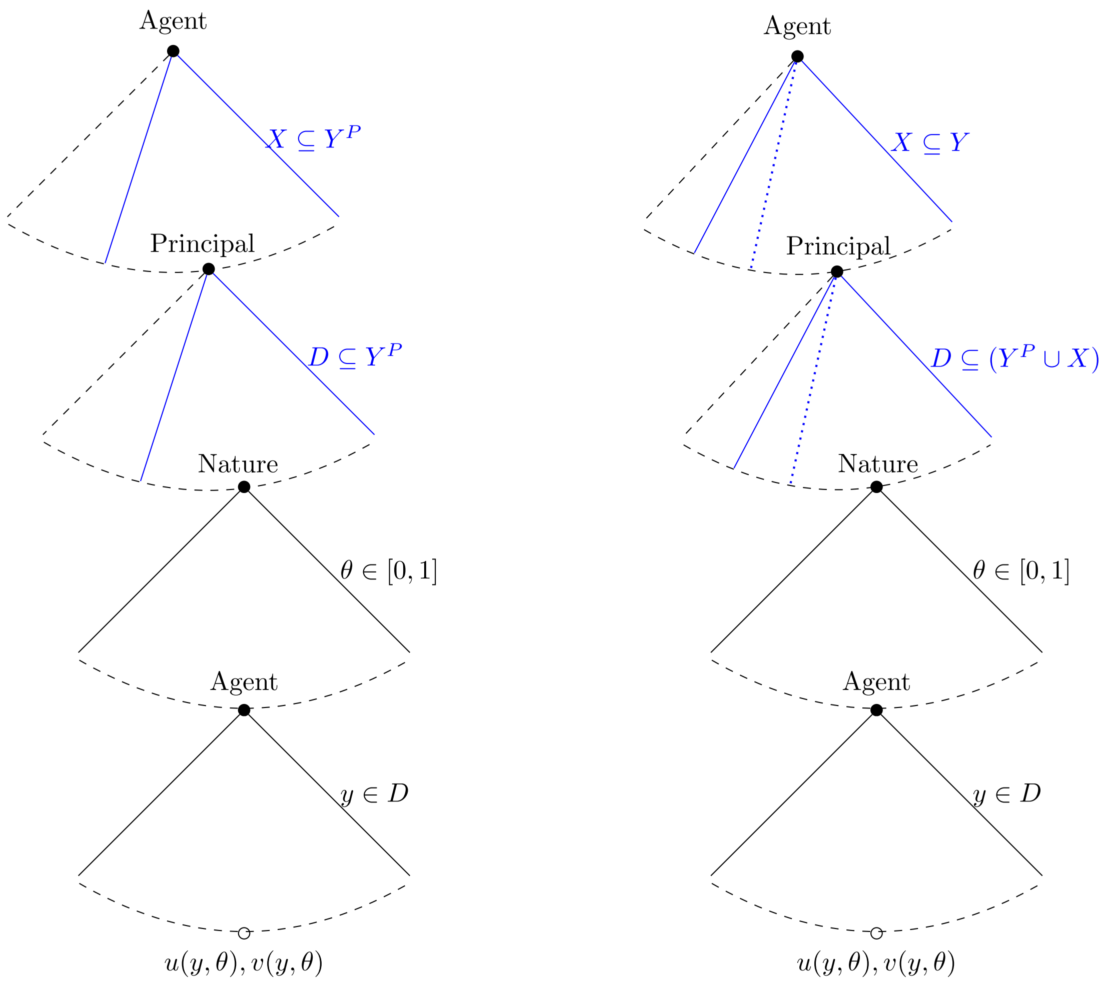

Figure 8: Subjective game trees T_{Y^P} and T_Y, respectively, before and after updating

C Extension to More General Preferences

Suppose the agent's and principal's preferences are as follows:

$$u(y, \theta) = y\theta - C(y) \quad \text{and} \quad v(y, \theta) = y(\theta - \beta) - C(y),$$

with $C(\cdot)$ strictly convex and twice differentiable. For $C(y) = \frac{1}{2}y^2$ we are back to the quadratic case. It is easy to show that under the same Assumption 1 the optimal delegation set is an interval of the form $[y_{min}, \hat{y}]$, where $\hat{y} = g(\hat{\theta})$, $\hat{\theta}$ solves $\hat{\theta} = \mathbb{E}[\theta - \beta | \theta \geq \hat{\theta}]$, and g is such that $C'(g(\theta)) = \theta$ for all θ. In this case, when the principal has limited awareness, the optimal unawareness interval might be non-symmetric around $\hat{y}$. Following the same line of proof of Proposition 3 in Appendix A.3, we can easily show that the class of sets of actions the principal is left unaware of can be described by the following intervals

$$[\hat{y} - \Delta_1, \hat{y} + \Delta_2], \qquad \text{where} \qquad \frac{C(\hat{y} + \Delta_2) - C(\hat{y} - \Delta_1)}{\Delta_1 + \Delta_2} = C'(\hat{y}).$$

The condition on the right delivers a unique value of Δ_1 for each Δ_2 and vice versa. The agent optimization problem can therefore be stated as a *function of only one variable* and the solution satisfies the analog to condition (3) in Proposition 4.

Notice that monotone increasing transformations of u and affine transformations of v can easily be allowed as well without any additional complication. State dependent bias can also be introduced with the preferences:

$$u(y, \theta) = y\theta - C(y) \quad \text{and} \quad v(y, \theta) = y(\theta - \beta(\theta)) - C(y),$$

with $\beta'(\theta) \in [0, 1]$. In this case, the analog to the first condition in Assumption 1, which implies interval delegation, is

$$\beta(\theta)f'(\theta) + (1 - \beta'(\theta))f(\theta) > 0.$$

References

[1] Alonso, R. and Matouschek, N., 2008. Optimal Delegation. *The Review of Economic Studies*, 75 (1), 259-293.

[2] Amador, M. and Bagwell, K., 2013. The Theory of Optimal Delegation with an Application to Tariff Caps. *Econometrica*, 81(4), 1541-1599.

[3] Armstrong, M. and Vickers, J., 2010. A Model of Delegated Project Choice. *Econometrica*, 78 (1), 213-244.

[4] Auster, S., 2013. Asymmetric Awareness and Moral Hazard. *Games and Economic Behavior*, 82, 503-521.

[5] Auster, S. and Pavoni, N., 2019. Limited Awareness and Financial Intermediation. Working Paper.

[6] Carpini, D. M. and Keeter, S., 1996. *What Americans Know about Politics and Why It Matters,* Yale University Press.

[7] Crawford, V. and Sobel, J., 1982. Strategic Information Transmission. *Econometrica*, 50, 1431-1451.

[8] Farrell, J., 1993. Meaning and Credibility in Cheap-Talk Games. *Games and Economic Behavior*, 5(4), 514-531.

[9] Feinberg, Y., 2012. Games with Unawareness. Working Paper.

[10] Filiz-Ozbay, E., 2012. Incorporating Unawareness into Contract Theory, *Games and Economic Behavior*, 76, 181-194.

[11] Guiso, L. and T. Jappelli, 2005. Awareness and Stock Market Participation, *Review of Finance*, 9, 1-31.

[12] Halac, M., and Yared, P., 2017. Commitment vs. Flexibility with Costly Verification, Working Paper.

[13] Heifetz, A., Meier M. and Schipper, B.C., 2006. Interactive Unawareness. *Journal of Economic Theory*, 130, 78-94.

[14] Heifetz, A., Meier, M. and Schipper, B.C., 2011. Prudent Rationalizability in Generalized Extensive-form Games, Working Paper.

[15] Heifetz, A., Meier M. and Schipper, B.C., 2013. Dynamic Unawareness and Rationalizable Behavior. Games and Economic Behavior, 81, 50-68.

[16] Holmström, B., 1977. On Incentives and Control in Organizations, Ph.D. Thesis, Stanford University.

[17] Karni, E. and Vierø, M.-L., 2013. Reverse Bayesianism. *American Economic Review*, 103(7), 2790-2810.

[18] Karni, E. and Vierø, M.-L., 2017. Awareness of Unawareness: A Theory of Decision Making in the Face of Ignorance. *Journal of Economic Theory,* 168, 301-32.

[19] Li, Y. and Schipper, B.C., 2017. Disclosure under Unawareness: An Experiment. Working Paper.

[20] Martimort, D. and Semenov, A., 2006. Continuity in mechanism design without transfers. *Economics Letters*, 93 (2), 182-189.

[21] Melamud, N.D. and Shibano, T., 1991. Communication in settings with no transfers. *RAND Journal of Economics*, 22, 173-198.

[22] Ozbay, E., 2008. Unawareness and Strategic Announcements in Games with Uncertainty, Working Paper.

[23] Szalay, D., 2005. The Economics of Clear Advice and Extreme Options, *Review of Economic Studies*, 72(4), 1173-1198.

[24] Somin, I., 2013. *Democracy and Political Ignorance: Why Smaller Government is Smarter,* Stanford University Press.

[25] von Thadden, E.-L. and Zhao X.J., 2012. Incentives for Unaware Agents, *Review of Economic Studies*, 79, 1151-1174.

[26] von Thadden, E.-L. and Zhao X. J., 2014. Multitask Agency with Unawareness, *Theory and Decisions*, 77(2), 197-222.

[27] Zhao, X.J., 2011. Framing Contingencies in Contracts, *Mathematical Social Sciences*, 61, 31-40.