
ECONtribute
Discussion Paper No. 059

**Optimal Delegation and Information
Transmission under Limited Awareness**

Sarah Auster

Nicola Pavoni

April 2023 (updated version)

www.econtribute.de



Optimal Delegation and Information Transmission under Limited Awareness*

Sarah Auster[‡] and Nicola Pavoni[§]

February 2023

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 superior awareness of the set of actions that are actually feasible. We provide conditions under which the agent finds it optimal to leave the principal unaware of relevant options. By doing so, the agent increases the principal's cost of distorting the agent's choices and increases the principal's willingness to grant him higher information rents. We further show that the principal may use the option of renegotiation as a tool to implement actions that are not describable to her at the contracting stage. If the agent renegotiates, his proposal signals information about the payoff state. Due to her limited awareness, the principal makes a coarse inference from the agent's recommendations and, as a result, accepts a large number of the agent's proposals, which ultimately benefits both.

JEL Codes: D82, D83, D86

*The authors thank Sandeep Baliga, Pierpaolo Battigalli, Sylvain Chassang, Wouter Dessein, Antonio Guarino, Yingni Guo, Marina Halac, Johannes Hörner, Navin Kartik, Elliot Lipnowski, 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 Planck 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. Sarah Auster acknowledges funding from the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy—EXC 2126/1—390838866 and CRC TR 224 (Project B02).

[‡]Department of Economics, University of Bonn: auster.sarah@gmail.com

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

1 Introduction

In many situations, economic agents delegate decisions to experts whose preferences may not be perfectly aligned with their own. Public and private organizations use procurement managers to purchase products and services for the tasks at hand; corporate headquarters rely on division managers with superior information about the profitability of new projects; small investors seek advice from financial experts with a better understanding of the risks and returns of the available portfolios. Oftentimes, the informed party not only has a better understanding of what the most suitable action is but also of the options that are actually available. Procurement managers have superior awareness of the feasible products and potential suppliers in the market where they operate; division managers have a better understanding of the projects they could pursue; financial experts are familiar with more financial instruments than retail investors, etc.

This paper proposes a framework to study the implications of such asymmetry by incorporating unawareness into a canonical delegation model. More specifically, we consider the problem of a principal (she) who needs to take an action and delegates the task to an agent (he). The agent receives private information about the payoffs of each available action, and the principal's problem is to determine a set of actions from which the agent can choose. We depart from the traditional framework of optimal delegation by considering a situation where the principal is *unaware of some feasible actions* and where this limits the language with which she can write a contract: the principal can only permit actions in the delegation set if she can name these actions explicitly, hence, if she is aware of them. Before the delegation stage and before receiving private information, the agent can expand the principal's awareness by revealing additional actions and thereby enrich the set of feasible contracts for the principal.

We are interested in the question if and how the agent distorts the principal's awareness in order to increase his own rent. We address this question in an environment with a continuum of payoff states, a continuum of feasible actions, and an agent who prefers a higher action than the principal in each state. Given her awareness, the principal's optimal delegation set solves the usual tradeoff between minimizing distortions and limiting the agent's information rent. Since the agent has an upward bias, optimal delegation

entails that the principal limits the agent's choice from above. An optimal delegation set thus has a threshold above which no action is permitted. How high this threshold is depends on the principal's awareness set. We identify conditions under which the agent optimally leaves the principal unaware of an interval of actions around the optimal upper threshold under full awareness. By choosing the bounds of the interval appropriately, the agent makes it optimal for the principal—who still cares about the agent's information—to permit an action above the full awareness cap and, hence, an action that would be precluded if the principal was fully aware.

An important assumption of our model is that the principal cannot specify actions in the delegation set of which she is unaware. An introspective principal might, however, ask herself whether there are other contracts that can improve on the optimal delegation set without giving the agent the flexibility of taking unknown, potentially harmful options. One such possibility for the principal is to forego full commitment and add a contractual clause that allows the delegation set to be adjusted when new options or contingencies come to light. Adding such a clause and thereby allowing for ex-post renegotiation is indeed consistent with the principal's sophistication and language. In the second part of the paper, we study the implications of the use of such contracts. The agent is then allowed to propose additional actions to the principal after the initial delegation set is agreed upon and he observes the state. Subsequently, the principal decides whether to permit a new action or whether to maintain the original delegation set.

Since the disclosure of additional actions is made after the agent receives private information, the principal and agent play a signaling game at the renegotiation stage. We characterize the set of proposals the principal is willing to accept after the initial agreement is signed. Fixing the agent's initial disclosure and focusing on the equilibrium with the maximal set of acceptable proposals, we provide conditions under which adding a renegotiation clause to the initial delegation set *always benefits the principal*. The downside of renegotiation is that it limits the agent's disclosure incentives in the contracting phase. Indeed, we show that, depending on the principal's initial awareness, the agent can significantly increase his flexibility by initially holding back some available actions and tailoring additional disclosures to the information he receives. From a modeler's view-

point, the agent's strategy in this equilibrium is fully revealing. The unaware principal, however, cannot compare the agent's proposal at a given state to the actions which the agent would have proposed in a different state and hence infers strictly less information from the agent's recommendation.

An application of our setting is procurement delegation within an organization. Many organizations have procurement managers in charge of purchasing products and services according to the organization's current needs. This choice is often limited via pre-specified lists of approved products or vendors. A possible concern rationalizing such restrictions is that the procurement manager can use his discretion to further personal goals, such as career enhancement, minimization of workload, personal enrichment, etc. (Rogerson, 1994). We can think of our agent as the procurement manager and interpret the agent's actions as the different products (or suppliers) available in the market. The state captures the characteristics of the task at hand, determining the procurement manager's "ideal" product. The principal can be interpreted as a high-level manager with limited awareness of the available products.

Our results suggest that a biased procurement manager typically benefits from hiding certain products or suppliers. In particular, under the identified conditions, the procurement manager has incentives to propose products with relatively extreme characteristics for approval but wants to hide those with more moderate attributes. The more pronounced the awareness asymmetry, the larger the scope for the procurement manager to distort the procurement decision through strategic disclosure. Finally, if the procurement manager has the possibility to seek approval for additional options after observing the characteristics of the task, he can substantially expand his flexibility by tailoring the disclosure to the circumstances.

After the literature review, the paper is organized as follows. Section 2 presents the delegation model with limited awareness. In Section 3 we analyze the agent's optimal disclosure and the resulting delegation set. Section 4 analyses the game with renegotiation and Section 5 concludes.

1.1 Related Literature

The paper makes both applied and theoretical contributions. It introduces unawareness to the canonical delegation problem and shows how the agent can distort the principal’s delegation choice through strategic disclosure. The analysis builds on the literature on optimal delegation. Holmström (1980) first defines the delegation problem and provides conditions for the existence of its solution. Following the seminal paper, the literature was further developed by Melumad and Shibano (1991), Szalay (2005), Martimort and Semenov (2006), Alonso and Matouschek (2008), Kováč and Mylovanov (2009), Armstrong and Vickers (2010), Amador and Bagwell (2013) and Halac and Yared (2020), among others. None of them consider limited awareness in this framework.

The paper is also related to the smaller literature that applies unawareness to games in general and contracting problems in particular. In contrast to our setting, most of the existing work considers contracting problems where contingent transfers are feasible and where *the agent* has limited awareness, while the principal is fully aware (Von Thadden and Zhao 2012; Zhao 2011; Filiz-Ozbay 2012; Auster 2013). One exception is Francetich and Schipper (2020), which studies a screening model where the principal is unaware of certain cost types (but has full awareness over actions) and the agent decides which types to disclose.

In Auster and Pavoni (2022), we consider a finance application of our model, interpreting the agent as a financial expert and the principal as an investor with limited awareness about the available financial products. We collect self-reported data from customers in the Italian retail investment sector and find support for the key predictions of the model: the menus offered to less knowledgeable investors contain fewer products, which are perceived to be more extreme.¹ Also Lei and Zhao (2021) study a financial market application of our model but focus on the unawareness of contingencies (nature’s moves) rather than players’ actions.

On the theoretical side, the study of the disclosure problem reveals how the agent’s

¹The data we collected consists of approximately 1,400 investors reporting on their experience in the Italian retail investment sector. We regress both the number of offered products and a measure of perceived ‘extremeness’ in the menu on an index of knowledge, which is based on a number of questions eliciting the investor’s background knowledge. See Auster and Pavoni (2022) for details.

information rents depend on the set of feasible actions—or the principal’s perception thereof—in delegation settings. This question is related to a recent literature looking at the determinants of agency rents in models with full awareness, initiated by Roesler and Szentes (2017). With the second part of the paper, we also contribute to the literature on incomplete contracts and unforeseen contingencies (Grossman and Hart 1986; Hart and Moore 1988) by demonstrating the value of ex-post renegotiation in settings with partial awareness. Previous papers that study the interaction between limited awareness and the possibility of renegotiation are Tirole (2009) and Piermont (2017), albeit in rather different settings.

2 Environment

There is a principal and an agent. The agent has access to an interval of actions $Y^A = [y_{min}, y_{max}]$. The principal’s and agent’s payoffs depend on the action that is chosen and an unknown payoff parameter θ , which can be privately observed by the agent. Let $\Theta = [0, 1]$ be the set of payoff states and let $F(\theta)$ denote the cumulative distribution function on Θ , assumed to be twice differentiable on the support. The principal and the agent have twice continuously differentiable utility functions²

$$U^P(\theta, y), \quad U^A(\theta, y).$$

Fixing θ , U^i for $i = P, A$ is assumed to be strictly concave in y with an interior maximum on Y^A . The principal’s and agent’s conditionally preferred actions are described by the functions

$$y^P(\theta) := \arg \max_{y \in Y^A} U^P(\theta, y), \quad y^A(\theta) := \arg \max_{y \in Y^A} U^A(\theta, y).$$

We assume $U_{\theta y}^i > 0$, which implies that $y^P(\cdot), y^A(\cdot)$ are strictly increasing functions. Furthermore, we assume that conditional on the payoff parameter θ , the agent prefers a higher action than the principal: for all θ , $y^P(\theta) < y^A(\theta)$.

²Note that the principal does not have full access to her payoff function U^P but just to a payoff function restricted to the domain of actions of which she is aware.

Awareness. Let \mathcal{Y} denote the set of closed subsets of $[y_{min}, y_{max}]$. The principal is aware of a subset of available actions, denoted by $Y^P \in \mathcal{Y}$. Hence, unawareness in our framework does not take the form of unforeseen contingencies but concerns the set of available actions. Apart from the assumption that Y^P is closed, we impose no further structure on the principal's initial awareness set. Before the principal contracts with the agent and the agent observes θ , the agent can make the principal aware of additional actions by revealing a closed set $X \in \mathcal{Y}$. 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.³

Delegation. Given her updated awareness, the principal offers a contract to the agent. 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 once he observes the payoff parameter θ .⁴ Our substantial assumption is that the principal's unawareness restricts the language with which she can write a contract. In particular, we assume that the principal can only refer to actions in the contract which she can name explicitly. The larger the principal's awareness set, the richer the set of contracts she can write.

Given the principal's updated awareness set, she then has two natural options: the principal can either name the actions she allows the agent to take or she can name the actions she explicitly forbids. Under full awareness, these two options are clearly equivalent. With unawareness, on the other hand, specifying only the forbidden actions leaves the principal vulnerable to the agent taking actions that the principal does not anticipate. We will discuss this case and other options in Section 3.2 and concentrate now on the case where the principal specifies the actions which she permits. Since the

³Assuming that the agent discloses actions before receiving private information avoids signaling effects in the baseline model. That is, after any disclosure by the agent, the principal's beliefs about the payoff state are described by the prior F . This is consistent with 'Reverse Bayesianism' (see Karni and Vierø (2013)), which postulates that relative beliefs on events of which the decision maker was previously aware do not change when her awareness grows.

⁴The standard delegation problem is equivalent to a mechanism design problem when the principal restricts herself to deterministic allocations (see Alonso and Matouschek (2008) and Kováč and Mylovanov (2009)). Formally, the principal commits to a mechanism that specifies an action as a function of the agent's message.

principal cannot specify actions of which she is unaware, the principal's delegation set is then a subset of her awareness set. We restrict attention to closed delegation sets. The timing of the game can be summarized as follows:

1. The agent reveals a set of actions $X \in \mathcal{Y}$ and the principal updates her awareness to $Y = Y^P \cup X$.
2. Given awareness set Y , the principal chooses a delegation set $D \in \mathcal{Y}$ such that $D \subseteq Y$.
3. The agent observes θ and chooses an action from set D .
4. Payoffs are realized.

The game between the principal and the agent can be formally represented by a family of partially ordered subjective game trees. Such 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., 2021).⁵ As a solution concept, we use a strong version of Perfect Bayesian Nash Equilibrium (PBE) which implies subgame perfection, adapted to generalized extensive-form games with unawareness (e.g., see Halpern and Rêgo (2014) and Feinberg (2021)).

Remark. There is an alternative reading of our model as one of limited authority. We can think of a situation where the agent, rather than disclosing feasible actions to the principal, actually enables the principal to pursue them. The agent thus decides on the set of actions he makes available to the principal and, as before, the principal delegates some subset of those actions to the agent. By deciding which actions to make available, the agent is given commitment power not to take certain actions. Since such commitment limits the principal's choice over feasible contracts, we are ultimately faced with a double delegation game between the agent and the principal.

⁵In a working paper version of this paper, available on our websites, we provide a more extensive description of the family of game trees representing the generalized game with unawareness associated to our delegation model according to the approach proposed by Heifetz et al. (2013). We also describe the set of outcomes that satisfy a prudent version of extensive-form rationalizability and we show that whenever we restrict to pure strategies and assume the tie-breaking rules we adopt below to be commonly known, the PBE outcome we obtain is also the sole rationalizable outcome of the generalized game.

3 Equilibrium Analysis

We will work backward and start the analysis by considering the last stage of the game. Given a delegation set D and observed payoff state θ , the agent's best response for the last stage of the game is defined by

$$BR^A(\theta, D) := \arg \max_{y \in D} U^A(\theta, y). \quad (1)$$

When the agent is indifferent between two actions, let $y^*(\theta, D) := \min BR^A(\theta, D)$ be the selection that takes the smallest value (indifference is broken in favor of the principal).⁶

Delegation stage. Turning to the principal's delegation choice, we first define the principal's value of delegation set $D \in \mathcal{Y}$ given y^* :

$$V^P(D) := \int_0^1 U^P(\theta, y^*(\theta, D)) dF(\theta). \quad (2)$$

There are typically actions that the principal could permit but the agent will not implement. W.l.o.g. we will restrict attention to delegation sets D such that for any $y \in D$, there is some state $\theta \in [0, 1]$ such that $y^*(\theta, D) = y$. Let $\mathcal{D}(Y)$ be the set of delegation sets in $\{D \in \mathcal{Y} : D \subseteq Y\}$ that satisfy this requirement. For each awareness set $Y \in \mathcal{Y}$, the principal's optimal delegation set solves the problem

$$\max_{D \in \mathcal{D}(Y)} V^P(D). \quad (3)$$

Theorem 1 in Holmström (1980) guarantees existence for each closed Y (see also Proposition 12 in Appendix A.8). If problem (3) has multiple solutions, we assume that the principal chooses the agent-preferred set. For each Y , we denote by $D^*(\cdot)$ such selection from the set of maximizers. Furthermore, we assume that in the case where the principal is fully aware, delegation is valuable. A sufficient condition for valuable delegation is $y_0^* > y^A(0)$, where $y_0^* \in \arg \max_y V^P(\{y\})$. This requires the bias to be not too large and implies that the principal prefers the delegation set $[y^A(0), y_0^*]$ to the singleton $\{y_0^*\}$ (see

⁶Such selection is well defined, since the BR^A correspondence is non-empty and upper hemicontinuous (see Holmström (1980)). In addition, for each closed D , the set of θ 's for which $BR^A(\cdot, D)$ is not a singleton is at most countable, and hence of F -measure zero (see Lemma 11 in Appendix A.8).

also Alonso and Matouschek (2008), Corollary 2).

Disclosure stage. In the first stage of the game, the agent chooses an awareness set $Y \in \mathcal{Y}$. Since the agent cannot make the principal unaware of actions that the principal already knows, the induced awareness set must contain the principal's initial awareness set Y^P . The smaller Y^P , the larger the collection of awareness sets from which the agent can choose. An optimal awareness set Y^* solves the problem

$$\max_{Y \in \mathcal{Y}} \int_0^1 U^A(\theta, y^*(\theta, D^*(Y))) dF(\theta) \quad \text{s.t.} \quad Y^P \subseteq Y. \quad (4)$$

Since different awareness sets might induce the same delegation set, the solution to problem (4) is again typically not unique. Of course, this type of multiplicity does not affect the outcome. We assume that when two solutions of problem (4) are nested, the agent discloses the larger set. This assumption allows us to distinguish the actions that remain undisclosed for strategic reasons from those that are redundant. Let \mathcal{Y}^* denote the set of all solutions of (4) satisfying this requirement.

Our model is a sequential move game with infinite actions. This makes equilibrium existence a non-trivial issue. In Appendix A.8, Proposition 13, we show that a solution to problem (4) exists. Hence, there is an equilibrium where the agent discloses a set $Y \in \mathcal{Y}^*$, the principal delegates set $D^*(Y)$ and, after observing the state realization θ , the agent takes action $y^*(\theta, D^*(Y))$.

Equilibrium disclosure. The central question of this paper is whether the agent distorts the principal's delegation choice in his favor by leaving the principal unaware of some feasible actions. Due to the conflict of interest between the principal and the agent, a fully aware principal will not find it optimal to permit the agent his preferred action in every payoff state. Indeed, since the agent is upward biased, the principal can always improve on full delegation by excluding an interval of high actions, forcing the agent for high realizations of θ to take an action closer to the principal's conditionally preferred action.

Following this argument, we define $\hat{y} := \max D^*(Y^A) < y^A(1)$ as the highest action which the principal permits under the optimal delegation set in the full awareness

benchmark. The following proposition shows conditions under which unawareness of \hat{y} is sufficient to ensure that the agent benefits from the principal's limited awareness.⁷ To this end, let $s := (y^A)^{-1}$ denote the inverse of y^A , implicitly defined by the first-order condition $U_y^A(s(y), y) = 0$. Since the agent's utility function is twice continuously differentiable, the function $s(\cdot)$ is differentiable (see Lemma 8 in the Appendix).

Suppose now the upper threshold \hat{y} is a limit point of $D^*(Y^A)$. It must then satisfy the following optimality conditions:

$$\int_{s(\hat{y})}^1 U_y^P(\theta, \hat{y}) dF(\theta) = 0 \quad (5)$$

$$-U_y^P(s(\hat{y}), \hat{y})f(s(\hat{y}))s'(\hat{y}) + \int_{s(\hat{y})}^1 U_{yy}^P(\theta, \hat{y})dF(\theta) \leq 0. \quad (6)$$

The first-order condition (5) says that, conditioning on the event $\theta \geq s(\hat{y})$, action \hat{y} is optimal for the principal in expectation. The second-order condition (6) is necessary for \hat{y} to constitute a local maximizer. For the following result we will maintain that this condition holds strictly.

Proposition 1. *Assume that the full awareness problem (3) has a unique maximizer $D^*(Y^A)$ and that the upper threshold \hat{y} is a limit point of $D^*(Y^A)$, with (6) holding as strict inequality. Then:*

$$\hat{y} \notin Y^P \implies Y^A \notin \mathcal{Y}^*.$$

Proposition 1 shows that, under the stated conditions, if the principal is initially unaware of the highest action in the optimal delegation set under full awareness, then the agent finds it profitable to hide some of the feasible actions from the principal. To prove the result, we consider a simple perturbation of the full awareness set. The perturbation entails that the principal remains unaware of an interval (y^-, y^+) of actions around the upper threshold \hat{y} . In the first step, we show that the bounds of the interval, y^- and y^+ , can be chosen in a way such that the principal finds it optimal to include both y^- and y^+ in the delegation set. The principal essentially tries to find the best approximation of the full awareness delegation set that is compatible with her actual awareness. By setting y^+

⁷Since Y^P is a closed set, unawareness of \hat{y} implies that the principal is unaware of an interval around \hat{y}

sufficiently close to \hat{y} relative to y^- , the agent assures that this approximation includes the higher action y^+ . Leaving the principal unaware of actions around \hat{y} thus allows the agent to implement an action $y^+ > \hat{y}$ that would not be permitted under full awareness.

The possibility of taking the higher action y^+ comes at the cost of losing the option to take an action in the interval $(y^-, \hat{y}]$. In the second step of the proof, we show that the perturbation is profitable for the agent despite this cost. The perturbation forces the agent to move away from the bliss point in states just below $s(\hat{y})$. However, since the marginal cost of moving away from the bliss point at the bliss point is zero, the effect of losing these actions is second order and thus dominated by the agent's benefit of increasing the implemented action in states above of $s(\hat{y})$. Note indeed that, since for the principal the action \hat{y} is optimal in expectation when conditioning on $\theta \geq s(\hat{y})$, it is strictly too low for the agent.

We should emphasize that when the optimal delegation set under full awareness is not an interval, the agent may profit from perturbations around other pooling points as well. For example, if $D^*(Y^A)$ has an intermediate gap (\underline{y}, \bar{y}) , the agent would benefit from moving up the lower bound \underline{y} at the cost of losing some flexibility below \underline{y} . The main complication here is that such perturbation may affect the principal's optimal choice of \bar{y} . If the optimal value for \bar{y} decreases as a result of the perturbation, the agent strictly gains. If, on the other hand, it increases, then there are two first-order effects that need to be compared. To guarantee the profitability of the perturbation in this case, more stringent assumptions on the principal's initial awareness set would be needed in order to give the agent the necessary tools to deter the principal from undesired movements of adjacent pooling actions. In the described situation, for instance, it might be necessary to keep the principal unaware of some actions to the right of \bar{y} .

Finally, notice that if the bias is not restricted to be positive but can change, the argument proving Proposition 1 can sometimes be extended to a potential lower threshold of $D^*(Y^A)$. For instance, suppose there is a point $\bar{\theta} \in (0, 1)$ such that the agent is downward biased if $\theta < \bar{\theta}$ and upward biased if $\theta > \bar{\theta}$. In this case, it is optimal for the principal to restrict the agent's choice from both sides, so we have $y^A(0) < \min D^*(Y^A) < \max D^*(Y^A) < y^A(1)$. Assuming that, in addition to the requirements of Proposition 1,

$\min D^*(Y^A)$ is a limit point of $D^*(Y^A)$ and that the analog of (6) holds as a strict inequality, a sufficient condition for the agent to optimally leave the principal unaware is that the principal is initially unaware of either threshold $\min D^*(Y^A)$ or $\max D^*(Y^A)$.⁸

Proposition 1 is a consequence of a more general principle. Revealing an action y to the principal typically has a benefit and a cost. Conditional on the principal permitting y , the benefit of revelation is the utility gain in the states where y is preferred by the agent. The downside is that the action may crowd out other actions which the principal would permit if she remains unaware. In regions of θ where the principal gives full discretion, crowding out is not an issue, so the agent optimally discloses the relevant options. In regions where the conflict of interest is instead severe, the principal optimally restricts the agent's choice and full revelation can be detrimental to the agent. In the case of Proposition 1, the action \hat{y} crowds out all actions higher than \hat{y} . Since the agent benefits from being permitted such actions, he optimally leaves the principal unaware of \hat{y} (and some actions around it).

Interval delegation. Without additional assumptions, we are unable to characterize the agent's disclosure set further, simply because we know very little about the optimal delegation set under full awareness. The literature on optimal delegation, however, establishes sufficient conditions under which the optimal delegation set under full awareness is an interval. These conditions assure that any delegation set that has gaps can be improved upon by adding intermediate actions to the set. Assumption 1 makes this requirement explicit.

Assumption 1. *Consider a delegation set $D \in \mathcal{Y}$ and its convex hull $\text{Conv}(D)$. Then, for all $A \subseteq \text{Conv}(D)$:*

$$\int_0^1 U^P(\theta, y^*(\theta, D)) dF(\theta) \leq \int_0^1 U^P(\theta, y^*(\theta, D \cup A)) dF(\theta).$$

Consider a convex delegation set and suppose the principal removes an interval of actions in the interior of the set. Let this interval be denoted by (\underline{y}, \bar{y}) . The removal of actions in (\underline{y}, \bar{y}) means that there is an interval of states where the agent switches to the

⁸We formalize the analogous condition and provide an argument for the claim at the end of the proof of Proposition 1.

lower action \underline{y} and an interval of states where the agent switches to the higher action \bar{y} with respect to the original delegation set. Since the principal is downward biased with respect to the agent, the switch to the lower action can benefit her, whereas the switch to the higher one does not. Concavity of U^P means that the principal is risk-averse and therefore has incentives to hedge against these two possibilities. Hence, unless the principal views the scenario of the beneficial switch considerably more likely, she favors intermediate actions. The literature on optimal delegation provides conditions on the state distribution with respect to the utility functions that guarantee this property. We provide a set of sufficient assumptions below. For more general conditions we refer the reader to Alonso and Matouschek (2008, Proposition 5) and Amador and Bagwell (2013, Propositions 1 and 2).

Assuming that interval delegation is optimal, the optimal delegation set under full awareness is described by an upper cap below which the agent is free to choose his preferred action. The optimal delegation set under full awareness thus takes the form $[y^A(0), y]$ for some $y < y^A(1)$. The associated value for the principal is

$$V^P([y^A(0), y]) = \int_0^{s(y)} U^P(\theta, y^A(\theta)) dF(\theta) + \int_{s(y)}^1 U^P(\theta, y) dF(\theta) \quad (7)$$

We now show that under Assumption 1, unawareness of the full awareness upper threshold \hat{y} maximizing (7) is not only a sufficient condition for less-than-full revelation to be strictly optimal but also a necessary one. Moreover, provided this assumption holds and the principal's payoff in (7) is single-peaked, the resulting delegation set has a single gap around \hat{y} .

Proposition 2. *Let Assumption 1 be satisfied. Assume the full awareness problem (3) has a unique maximizer, with (6) holding as a strict inequality.*

- (i) *The agent optimally reveals all feasible actions to the principal if and only if the principal is aware of action \hat{y} .*
- (ii) *If in addition (7) is strictly quasi-concave in y , there exist two parameters $\Delta_1, \Delta_2 \geq$*

0 such that:

$$\begin{aligned} Y^* &= [y_{min}, \hat{y} - \Delta_1] \cup [\hat{y} + \Delta_2, y_{max}], \\ D^*(Y^*) &= [y^A(0), \hat{y} - \Delta_1] \cup \{\hat{y} + \Delta_2\}, \end{aligned}$$

with $\Delta_1, \Delta_2 > 0$ if and only if $\hat{y} \notin Y^P$.

Proposition 2 shows that when the principal is initially aware of \hat{y} the agent optimally reveals everything. The principal will not allow the agent to take any action higher than \hat{y} , so the agent maximizes his discretion by revealing all actions below \hat{y} —he cannot improve on the full awareness delegation set $[y^A(0), \hat{y}]$. On the other hand, if the principal is initially unaware of \hat{y} , it is optimal for the agent to leave the principal unaware of an interval around \hat{y} and disclose the remaining actions (recall that we assumed that the agent reveals the largest optimal set). The resulting delegation set includes all relevant actions below the interval and one action above it.

When choosing the optimal disclosure set, the agent is constrained by the principal's initial awareness. The more actions the principal initially knows, the smaller the set of awareness sets the agent can induce. Under Assumption 1, this gives rise to a monotonic relation between the principal's awareness before and after disclosure: the smaller the principal's initial awareness set Y^P , the smaller her equilibrium awareness set after the agent's disclosure. Indeed, adding an action to the principal's initial awareness set Y^P can only shrink the induced awareness gap (Δ_1, Δ_2) and hence expand the set of actions that the agent optimally proposes.

Remark. We model limited awareness as a restriction on the set of feasible contracts from which the principal can choose. It is interesting to contrast our framework to the case where the principal can describe all possible actions but faces uncertainty about their availability.⁹ The principal would then have a prior belief about the available set and perfectly understand the payoff implications of all potentially available actions. The agent's disclosure choice would be interpreted as evidence that certain actions are available and, in addition, serve as a signal about the availability of other undisclosed

⁹Adapting the setting of Armstrong and Vickers (2010), Guo and Shmaya (2021) study such disclosure problem and solve for the mechanism that minimizes the maximal regret.

actions. Assumption 1 implies that the principal benefits from closing potential gaps in the delegation set, independent of her belief about which actions are actually available. Hence, in contrast to our setting, the principal would delegate an interval following any disclosure of the agent, as long as Assumption 1 holds.¹⁰

3.1 Quadratic Utility and Uniform Bias

For a concrete illustration of the main results and an explicit solution of the agent's optimal disclosure policy, consider the specification

$$U^P(y, \theta) = -(y - (\theta - \beta))^2, \quad U^A = -(y - \theta)^2. \quad (8)$$

The agent's conditional preferred action is $y^A(\theta) = \theta$, while the principal's preferred action is $y^P(\theta) = \theta - \beta$. The agent thus has a constant upward bias equal to β . In this environment, a condition implying Assumption 1 and hence guaranteeing interval delegation to be optimal is the following regularity condition on the distribution function (Martimort and Semenov, 2006):

$$f'(\theta)\beta + f(\theta) > 0 \text{ for all } \theta \in (0, 1). \quad (9)$$

Delegation is valuable for the principal if $\mathbb{E}[\theta - \beta] > 0$. When both conditions are satisfied, the optimal delegation set under full awareness is an interval $[0, \hat{y}]$, where \hat{y} solves the following equality (Martimort and Semenov, 2006):¹¹

$$\hat{y} = \mathbb{E}[\theta - \beta | \theta \geq \hat{y}]. \quad (10)$$

To characterize the equilibrium disclosure and delegation sets for the quadratic setting, let $\bar{\Delta}(Y^P) := \arg \min_{y \in Y^P} |y - \hat{y}|$ indicate the smallest distance between \hat{y} , as defined in (10), and the actions in the principal's awareness set.

Proposition 3. *Assume preferences are as in (8), condition (9) is satisfied and $\mathbb{E}[\theta - \beta] > 0$. The equilibrium disclosure set is*

¹⁰The agent could still try to raise the upper threshold of that interval by leaving the principal uncertain about the availability of some actions around \hat{y} . Whether this is possible or not will depend on the principal's equilibrium beliefs following the agent's disclosure.

¹¹If instead $\mathbb{E}[\theta - \beta] < 0$, the optimal delegation set is $\{\mathbb{E}[\theta - \beta]\}$.

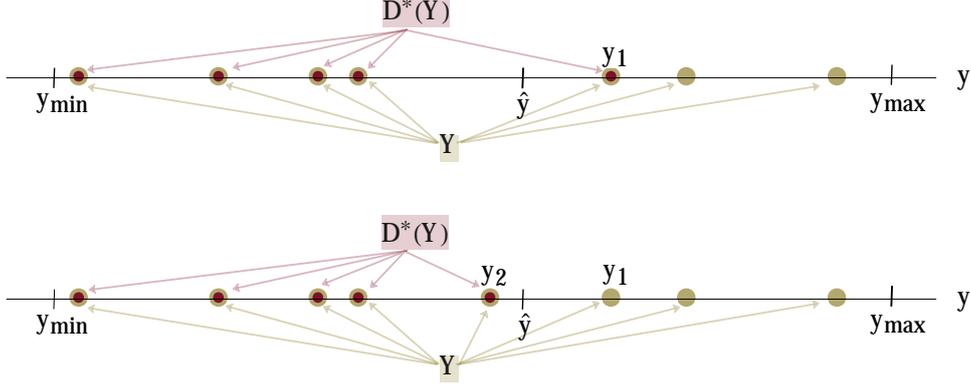


Figure 1: Optimal delegation set $D^*(Y)$. The figures represent two examples of the principal's awareness set Y . In both figures, the yellow bullets represent the set Y while the red bullets represent the resulting optimal delegation set $D^*(Y)$. In the upper figure, the principal includes action 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 $D^*(Y)$.

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

where $\Delta = \min\{\bar{\Delta}(Y^P), \Delta^*\}$ and Δ^* solves the agent's first order condition

$$\int_{\hat{y}-\Delta^*}^{\hat{y}} [\theta - (y - \Delta^*)] dF(\theta) = \int_{\hat{y}}^1 [\theta - (y + \Delta^*)] dF(\theta). \quad (11)$$

The resulting equilibrium delegation set is

$$D^*(Y^*) = [0, \hat{y} - \Delta] \cup \{\hat{y} + \Delta\}. \quad (12)$$

Under the restriction to quadratic preferences and a uniform bias, the optimal delegation set for a given awareness set Y is given by

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

The optimal delegation set under partial awareness can thus be seen as the closest approximation of the optimal delegation interval under full awareness, $[0, \hat{y}]$, which is available to the principal given her restricted awareness. The 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} , as illustrated in Figure 1.

This implies that the agent optimally chooses a gap that is symmetric around \hat{y} ,

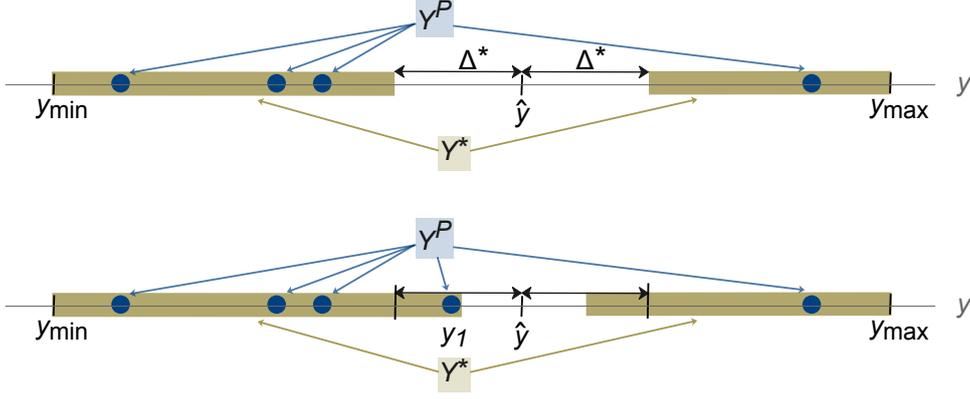


Figure 2: Optimal awareness set Y^* . The figures represent two examples of the principal's initial awareness set Y^P and associated awareness sets $Y^* = Y^P \cup X^*$ after including disclosed actions X^* . In both figures, the blue bullets represent the set Y^P , while the yellow set represents the 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 , making the unconstrained solution Δ^* infeasible.

i.e. $\Delta_1 = \Delta_2 = \Delta$. We can then write the agent's optimization problem over a single parameter Δ as follows:

$$\max_{\Delta \geq 0} - \int_{\hat{y}-\Delta}^{\hat{y}} (\hat{y} - \Delta - \theta)^2 dF(\theta) - \int_{\hat{y}}^{\hat{y}+\Delta} (\hat{y} + \Delta - \theta)^2 dF(\theta) \quad \text{s.t.} \quad \Delta \leq \bar{\Delta}(Y^P). \quad (13)$$

The first order condition (11) describes Δ^* as the unrestricted maximum of this problem. 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 choose the largest feasible gap, as we illustrate in Figure 2.

Next, we show how the optimal awareness gap from the agent's perspective, $(\hat{y} - \Delta^*, \hat{y} + \Delta^*)$, depends on the bias β .

Proposition 4. *Assume preferences are as in (8), condition (9) is satisfied and $\mathbb{E}[\theta - \beta] > 0$. Let $\Delta^*(\beta)$ be the unconstrained solution to problem (13) when the principal's preferences parameter is $\beta \in (0, \mathbb{E}[\theta])$. Then $\Delta^*(\cdot)$ is an increasing function.*

Proposition 4 shows an intuitive result: the larger the divergence between the principal's and the agent's preferred action, the more the agent wants to distort the principal's delegation choice by hiding actions from the principal. For a simple illustration, consider the case where F is uniform. The larger β is, the lower is the cap $\hat{y}(\beta)$ of the optimal delegation set under full awareness, as can be seen from condition (10). Considering the

agent's tradeoff when choosing Δ , notice that when F is uniform, the cost associated with the loss of flexibility for a given gap around $\hat{y}(\beta)$ is the same for all β . The desired consequence of generating a gap is an increase of the highest permitted action—from $\hat{y}(\beta)$ to $\hat{y}(\beta) + \Delta$ —and hence an increase of the agent's information rent in all states above $\hat{y}(\beta) + \Delta$. The lower the original cap $\hat{y}(\beta)$, the larger the range of values for θ above $\hat{y}(\beta) + \Delta$ and hence the set of types to whom this rent accrues.

Remark. While we assume that the set of available actions is an interval for tractability, the agent's disclosure incentives are qualitatively similar in the case where Y^A is an arbitrary subset of \mathbb{R} , e.g. a collection of discrete points. The analysis of the optimal delegation set for a given awareness set $Y \subseteq Y^A$ remains valid, so $D^*(Y)$ is described by (12). With regard to the optimal awareness set, notice that if the agent reveals some $y \in Y^A$, he might as well reveal all 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 parameter Δ . A sufficient condition for optimality of full disclosure is then that the principal is initially aware of the action in Y^A closest to \hat{y} . A sufficient condition for less-than-full disclosure is that the set of feasible actions includes three actions $y_1 < y_2 < y_3$ such that $Y^P \cap (y_1, y_3) = \emptyset$ and $|y_2 - \hat{y}| < |y_3 - \hat{y}| \leq |y_1 - \hat{y}| < \delta$ for some δ sufficiently small.

3.2 Limited Awareness and Contract Language

We assume that the principal specifies the actions that are *permitted* (ruling in) rather than those that are *not permitted* (ruling out). While under full awareness these two options are equivalent, in the case of limited awareness contract language matters for the outcome. If the principal's contract only specifies those actions that are not permitted, the agent has no incentives to expand the principal's awareness: since the agent is free to take any action of which the principal remains unaware, revealing them before the contracting stage can only reduce his flexibility.

The restriction to ruling-in contracts arises naturally if one views the principal's problem as designing a direct mechanism. Under this interpretation, the principal commits to a mapping from messages to actions, where unawareness imposes restrictions on the image

of such mappings. In particular, the principal cannot commit to an action that she does not know to exist. It might be interesting to consider more complex relationships between the principal’s awareness and implementable action profiles through indirect mechanisms. For instance, the principal could attempt to permit additional actions through an indirect description of those. Whether this improves the principal’s welfare or actually hurts her depends on the details of the model, of which the principal is unaware. Any aversion to such unknown possibilities might in fact call for a description of actions as specific as possible given the principal’s language.

An introspective principal—one who is aware of her unawareness—might, however, wonder whether there are any contracts that can improve on the optimal delegation set without giving the agent blanket approval to take unknown actions. One such possibility is to add a contractual clause specifying that the initial contract can be adjusted when new options come to light and parties mutually agree (see also Piermont (2017)). Specifying a contractual clause of this form would not rely on the principal’s ability to describe actions outside her awareness and would hedge her against the possibility of the agent taking harmful actions without her consent. In the spirit of the incomplete contract literature, one could then view the initial delegation set as a preliminary agreement that can be renegotiated when new, mutually beneficial options appear. We explore this possibility in the following section. In contrast to the incomplete contracts literature, we will maintain the assumption that the principal has full commitment, so renegotiation is in fact fully avoidable. We ask instead whether under limited awareness the principal can actually benefit from voluntarily giving up some of that commitment.

4 Ex-Post Renegotiation

Suppose that, rather than fully committing to the initial delegation set, the principal proposes a contract that fixes a set of permitted actions but allows for an adjustment when new options appear. We thus consider contracts under which the agent can renegotiate with the principal over actions that were not disclosed (or simply not permitted) in the initial stage of the game. Crucially, we allow the agent to propose such actions after he observes the payoff state, thereby signaling information. In particular, upon receiving a

proposal for a new action, the principal infers that the agent makes such a proposal only if taking the new action benefits him. However—due to the principal’s limited awareness—she cannot conceive of alternative actions the agent could have disclosed instead and, hence, cannot learn from particular actions not being proposed. This asymmetry arises as a consequence of the principal’s limited awareness and plays a crucial role in the results that follow.

The modified game has two phases, *the contracting phase* and *the renegotiation phase*. The contracting phase is the same as before: the agent discloses a set of actions and the principal determines a delegation set. In the renegotiation phase, the agent first observes the payoff state θ and then decides between two options. Either he picks an action from the delegation set or he proposes a different action to the principal, who can then accept the proposal or keep the original delegation set.

Strategies and beliefs. While we return to the contracting phase at the end of the section, our main focus will lie on the renegotiation phase. To this end, we fix the principal’s interim awareness set $Y \in \mathcal{Y}$ and a delegation set $D \subseteq Y$ as primitives. The set Y is interpreted as the principal’s updated awareness after the contracting phase and the set D as the corresponding delegation set. Next, we define the strategies of the principal and the agent. The agent’s possible moves are either ‘no new proposal’ (let us call it N) or singleton proposals $x \in Y^A$. The set of possible proposals for the agent is thus $X := N \cup Y^A$ and the agent’s strategy is a map $x : [0, 1] \rightarrow X$ from the possible realizations of θ to a recommendation. Upon receiving a new proposal, the principal needs to decide whether to accept or reject it. Her strategy is a mapping $\rho : X \rightarrow \{0, 1\}$, where $\rho(x) = 0$ means that the principal rejects the agent’s proposal and keeps the original delegation set D , while $\rho(x) = 1$ means that the principal accepts the agent’s proposal and the implemented action is x . To account for the fact that after ‘no new proposal’ the original delegation set must be kept, we set $\rho(N) = 0$. Whenever there is no new proposal or the proposal is rejected, the agent chooses an action from the original set D . The agent’s optimal choice in this case is described by $y^*(\theta, D)$, as introduced in Section 3. We will take this part of the agent’s strategy as given.

We restrict strategies x and ρ to be upper semicontinuous functions.¹² Intuitively, this amounts to assuming that, in case of indifference, the agent breaks ties in favor of proposing a new action, and the principal breaks ties in favor of allowing new proposals. It is easy to see that this assumption generates equilibrium sets of permitted actions that are closed. We further concentrate on outcomes in pure strategies. Finally, we denote for each $x \in X$ the set of conceivable proposals under the principal's updated awareness $Y \cup \{x\}$ by $X_x := Y \cup \{x\} \cup \{N\}$.

Definition 1. Fix an awareness set $Y \in \mathcal{Y}$ and a delegation set $D \subseteq Y$. The strategy profile (x^*, ρ^*) , together with a belief function $\mu^*(\cdot|x) \in \Delta([0, 1])$ for each $x \in X$ and a collection of strategy perceptions $(x_x^*)_{x \in X}$, with $x_x^* : [0, 1] \rightarrow X_x$, constitutes a PBE of the renegotiation game if and only if the following conditions hold:

1. *Principal optimality:* for all $x \in X \setminus N$,

$$\rho^*(x) \in \arg \max_{\rho \in \{0,1\}} \rho \mathbb{E}_{\mu^*(\cdot|x)}[U^P(\theta, x)] + (1 - \rho) \mathbb{E}_{\mu^*(\cdot|x)}[U^P(\theta, y^*(\theta, D))];$$

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

$$x^*(\theta) \in \arg \max_{x \in X} \rho^*(x) U^A(\theta, x) + (1 - \rho^*(x)) U^A(\theta, y^*(\theta, D));$$

3. *Consistency of beliefs:* for all $x \in X$, $\mu^*(\cdot|x)$ is consistent with the perceived strategy $x_x^*(\cdot)$, where

$$x_x^*(\theta) \in \arg \max_{x' \in X_x} \rho^*(x') U^A(\theta, x') + (1 - \rho^*(x')) U^A(\theta, y^*(\theta, D)). \quad (14)$$

In particular, letting $\Theta^*(x)$ denote the preimage of x for the function $x_x^*(\cdot)$, $\mu^*(\cdot|x)$ is derived via Bayes rule whenever $\int_{\Theta^*(x)} dF(\theta) > 0$. If $\int_{\Theta^*(x)} dF(\theta) = 0$ but $\Theta^*(x) \neq \emptyset$, then $\mu^*(\cdot|x)$ is an arbitrary distribution with support $\Theta^*(x)$. Finally, if $\Theta^*(x) = \emptyset$, then $\mu^*(\cdot|x)$ is unrestricted;

The main novelty in Definition 1 is the awareness-adapted consistency condition, which assures that the principal's beliefs are coherent with the agent playing optimally

¹²To define upper semicontinuity for the agent, associate a negative number to N in the codomain of his strategy.

in the game *as perceived through the principal's awareness*. To formalize this requirement we need to account for the fact that the principal's perception of the agent's set of feasible strategies depends on the principal's updated awareness set and, hence, on the agent's realized proposal. Indeed, each proposal $x \in X$ induces a different subjective game in the mind of the principal. We thus define for each $x \in X$, a perceived strategy x_x^* , which maps the state $\theta \in [0, 1]$ to a feasible recommendation $x' \in X_x$ in the principal's subjective game. Condition (14) then requires that in this game, strategy x_x^* is optimal against the principal's equilibrium strategy ρ^* . The key implication of the consistency condition is that after any change of awareness following the agent's proposal, the principal's strategy and her beliefs are part of an equilibrium in the resulting subjective game.¹³

Acceptable proposals. We now ask which proposals the principal is willing to accept in a renegotiation equilibrium. To answer the question, we restrict attention to initial delegation sets D that solve the principal's delegation problem (3): $D = D^*(Y)$. Given the principal's awareness in the contracting phase, delegation set $D^*(Y)$ is indeed optimal, since any action in Y , which the principal plans to permit in the renegotiation phase, can directly be included in the delegation set. Given this restriction, we can show that in any equilibrium of the renegotiation game, the principal permits an agent's proposal $x \notin D^*(Y)$ only if she would have preferred to add the action to the initial delegation set $D^*(Y)$ at the contracting stage. The set of actions satisfying this requirement is defined by

$$A(Y) := \{x \in Y^A : V^P(D^*(Y) \cup \{x\}) \geq V^P(D^*(Y))\},$$

where V^P is the principal's value in the 'full commitment delegation' benchmark.

Proposition 5. *Fix an awareness set $Y \in \mathcal{Y}$ and a delegation set $D = D^*(Y)$.*

(i) *In any equilibrium $(x^*, \rho^*, (\mu^*(\cdot|x), x_x^*)_{x \in X})$, if $\rho^*(x) = 1$, then $x \in A(Y)$.*

(ii) *There is an equilibrium such that $\rho^*(x) = 1$ for all $x \in A(Y)$.*

¹³We could impose the consistency condition only on beliefs following on-path proposals. None of our results would be affected. While this is obvious for Propositions 5(ii), 6, and 7, also Proposition 5(i) remains valid, as the argument proving it does not rely on the specification of off-path beliefs (see Appendix A.5).

Proposition 5 characterizes the set of proposals that can be accepted by the principal in equilibrium. By definition of $D^*(Y)$, the set $A(Y)$ does not include any actions that belong to Y other than those already in the delegation set $D^*(Y)$. Hence, in equilibrium we have $\rho^*(x) = 0$ for all $x \in Y \setminus D^*(Y)$. This means that, given delegation set $D^*(Y)$, the agent can only gain from renegotiation if he discloses new actions of which the principal was previously unaware. Consider then an equilibrium where proposal $x \notin Y$ is accepted in the renegotiation phase. By the consistency condition, the principal's beliefs after this proposal have support

$$\Theta^*(x) = \{\theta \in [0, 1] : U^A(\theta, x) \geq \max_{y \in D^*(Y)} U^A(\theta, y)\}. \quad (15)$$

According to the principal's awareness, x is the only new action that the agent can propose. The principal thus believes that the agent proposes x whenever he prefers it over his best alternative in $D^*(Y)$. The question is then whether conditional on the agent preferring x over the actions belonging to $D^*(Y)$, the principal prefers x as well. The answer to this question is yes if and only if the principal would have preferred to add x to the delegation set $D^*(Y)$, i.e., if and only if $x \in A(Y)$. This is because adding an action to the delegation set changes the outcome only in those states where the agent prefers the action to the alternatives in the delegation set. In the renegotiation phase, the same consideration applies.

The benefit of partial commitment. The previous result demonstrates that by adding a renegotiation option to the optimal delegation set the principal can keep some flexibility to implement additional actions should her awareness grow while generating the same outcome as under full commitment in case her awareness remains unchanged. We show next that, if Assumption 1 holds, partial commitment indeed dominates full commitment.

Proposition 6. *Let Assumption 1 be satisfied. Fix an awareness set $Y \in \mathcal{Y}$, a delegation set $D = D^*(Y)$ and consider a renegotiation equilibrium $(x^*, \rho^*, (\mu^*(\cdot|x), x_x^*)_{x \in X})$ where $\rho^*(x) = 1$ for all $x \in A(Y)$ and $\rho^*(x) = 0$ otherwise. The principal's expected payoff in*

this equilibrium is $V^P(A(Y))$ and satisfies

$$V^P(A(Y)) \geq V^P(D^*(Y)). \quad (16)$$

Proposition 6 shows that, focusing on the equilibrium where the set of accepted proposals is maximal, the principal benefits from partially forgoing her commitment if Assumption 1 is satisfied. In equilibrium, the set of implementable actions for the agent is $A(Y)$ and the principal's equilibrium payoff (as viewed from the perspective of a fully aware outside observer) is given by $V^P(D^*(Y) \cup A(Y))$. By Assumption 1, the set $A(Y)$ includes all actions in $[y^A(0), \max D^*(Y)]$. Intuitively, this means that the agent can “close potential gaps” of the original delegation set $D^*(Y)$ through renegotiation. By the same assumption, convexifying the set not only benefits the agent but also the principal. The set $A(Y)$ may also include actions strictly higher than $\max D^*(Y)$. In the proof, we show that their inclusion benefits the principal as well.¹⁴

Information transmission. A striking feature of the described equilibrium is that the implemented action is strictly increasing in the state for all θ such that $y^A(\theta) \leq \max D^*(Y)$, even when the initial delegation set $D^*(Y)$ has gaps. This would not be possible under full awareness: in any candidate equilibrium where types perfectly separate themselves through their announcement, the fully aware principal learns the payoff state and has incentives to deviate to a strictly lower action, at least in some states. In the case of limited awareness, however, the principal cannot contemplate the agent's moves of which she remains unaware and this limits the extent to which she infers information from the agent's recommendation.

In particular, if the realized value is θ and the agent proposes an action $y^A(\theta) \notin Y$ such that $y^A(\theta) \in [y^A(0), \max D^*(Y)]$, the subjective game tree that represents the principal's frame of mind after updating does not include moves of the agent involving a proposal just below or above $y^A(\theta)$. As a consequence, the principal cannot conceive of the fact that she would have permitted such actions if the agent had proposed them instead. In

¹⁴To see why Assumption 1 is needed, suppose it is not satisfied. Then we cannot rule out the possibility to have an awareness Y , a delegation set $D^*(Y)$, and two actions $y, y' \notin D^*(Y)$ such that $V^P(D^*(Y) \cup \{y\}), V^P(D^*(Y) \cup \{y'\}) \geq V^P(D^*(Y))$ and $V^P(D^*(Y) \cup \{y, y'\}) < V^P(D^*(Y))$. The principal may thus be worse off after renegotiation.

the principal's subjective game following proposal $y^A(\theta)$, there is an equilibrium where the agent reveals $y^A(\theta)$ in all states where the agent prefers $y^A(\theta)$ over the actions in the initial delegation set. Each of the agent's equilibrium proposals is thus perceived to be consistent with an interval of states and these intervals overlap; that is, the principal's information can no longer be represented by a partition of the state space into pairwise disjoint sets. The discrepancy between the agent's true strategy and the principal's perception of it is exactly what allows for a continuum of on-path proposals. Sometimes the principal's coarse inference leads her to accept proposals that she should reject, e.g. when the agent proposes an action close to the lower boundary of any potential gap in $D^*(Y)$. In expectation, however, she gains from the additional flexibility that she grants in equilibrium.

Disclosure in the contracting phase. While for a fixed awareness set renegotiation unambiguously benefits the principal, the prospect of being able to renegotiate after the arrival of information affects the agent's disclosure incentives in the contracting phase. The agent's initial disclosure determines the principal's delegation set and with that the set of proposals the principal is willing to accept in the renegotiation phase. His goal is to maximize the final set of permitted actions, whether permission is given in the contracting phase or in the renegotiation phase. Focussing again on the case where, upon inducing awareness Y , the agent expects the principal to delegate $D^*(Y)$ and accept any additional proposal in $A(Y)$, the agent's optimal disclosure in the contracting phase solves the problem

$$\max_Y \int_0^1 U^A(\theta, y^*(\theta, A(Y))) dF(\theta) \quad (17)$$

subject to $Y^P \subseteq Y \subseteq Y^A$. When Assumption 1 is satisfied, this problem has a simple solution.

Proposition 7. *Let Assumption 1 be satisfied. A solution to problem (17) exists and is given by the awareness set Y that solves*

$$\max_{Y^P \subseteq Y \subseteq Y^A} \max D^*(Y).$$

The proposition shows that the agent's highest equilibrium payoff is attained by having the agent disclose a set of actions in the contracting phase that maximizes the upper threshold of the principal's resulting delegation set. Since through renegotiation, the agent is able to implement all actions below the upper threshold, this maximizes the agent's flexibility in equilibrium and thus his equilibrium payoff.

For concreteness, consider case (ii) of Proposition 2, where the principal's value $V^P([y^A(0), y])$ is single-peaked in y and the agent's optimal disclosure policy introduces a single gap around \hat{y} . Let $(\hat{y} - \bar{\Delta}_1, \hat{y} + \bar{\Delta}_2)$ be the largest feasible awareness gap such that the corresponding delegation set is $D^*(Y) = [y_{min}, \hat{y} - \bar{\Delta}_1] \cup \{\hat{y} + \bar{\Delta}_2\}$. The highest action the principal is willing to delegate in the contracting phase is thus

$$\max_{Y^P \subseteq Y \subseteq Y^A} \max D^*(Y) = \hat{y} + \bar{\Delta}_2.$$

Given awareness set $Y = [y_{min}, \hat{y} - \bar{\Delta}_1] \cap [\hat{y} + \bar{\Delta}_2, y_{max}]$ and delegation set $D^*(Y) = [y_{min}, \hat{y} - \bar{\Delta}_1] \cup \{\hat{y} + \bar{\Delta}_2\}$, there is an equilibrium in the renegotiation phase where the agent can implement any action in the interval $[y^A(0), \hat{y} + \bar{\Delta}_2]$. If the realized state θ is such that $y^A(\theta) \in D^*(Y)$, the agent does not renegotiate and takes his preferred action $y^A(\theta)$. If the realized θ is such that $y^A(\theta) > \hat{y} + \bar{\Delta}_2$, the agent does not renegotiate either, because, conditioning on the event that the agent prefers some action $x > \hat{y} + \bar{\Delta}_2$ over $\hat{y} + \bar{\Delta}_2$, the principal strictly prefers $\hat{y} + \bar{\Delta}_2$. If instead the state θ is such that $y^A(\theta) \in (\hat{y} - \bar{\Delta}_1, \hat{y} + \bar{\Delta}_2)$, the agent renegotiates and proposes his preferred action $y^A(\theta)$. The principal is unaware of other actions in the interval $(\hat{y} - \bar{\Delta}_1, \hat{y} + \bar{\Delta}_2)$ and thus only infers that the agent prefers the proposed action to all other actions in $D^*(Y)$. Conditioning on this information, she prefers x as well.

The model with renegotiation highlights an important aspect concerning the dynamics of unawareness. Much like information, awareness 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 the collection of the opponent's 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 the later stages of the game. In our

environment, this principle is reflected in the fact that the agent reveals fewer actions in the contracting phase when renegotiating is possible than when it is not. In the case discussed above, the optimal awareness gap in the contracting phase is maximal when renegotiating is possible. Notice that, even without renegotiation, the agent could implement *any single action below* $\hat{y} + \bar{\Delta}_2$ by revealing the ‘right’ set of actions. He cannot, however, implement *all actions below* $\hat{y} + \bar{\Delta}_2$ because some actions crowd out others. The agent has to make a choice based on the expected value of the feasible awareness sets and the resulting delegation sets. When renegotiation is possible, on the other hand, the agent can condition the principal’s awareness on the realization of θ .

Ex-ante welfare. Since the possibility to renegotiate limits the agent’s disclosure incentives in the contracting phase, the ranking of the principal’s expected payoff between the two cases, with and without renegotiation, depends on the principal’s initial level of awareness. To illustrate this, consider the following example based on the model analyzed in Section 3.1.

Example. Assume $U^P(y, \theta) = -(y - (\theta - \beta))^2$, $U^A = -(y - \theta)^2$. Let θ be uniformly distributed on $[0, 1]$ and assume $\beta < 1/2$, so that delegation is valuable. The optimal full awareness cap in this example is $\hat{y} = 1 - 2\beta$, while the parameter characterizing the agent’s unconstrained solution of the disclosure problem is $\Delta^* = 2(\sqrt{2} - 1)\beta$. Denote by $\bar{\Delta}(Y^P)$ the action in the principal’s initial awareness set closest to \hat{y} , where, for ease of notation, we will suppress the argument Y^P .

A calculation shows that under this specification, the principal’s expected payoff in the case with renegotiation is higher than in the case without it if and only if

$$\frac{\bar{\Delta}}{\beta} \leq \sqrt{2} - \frac{1}{2} \left(\sqrt{3(4\sqrt{2} - 5)} - 1 \right) \approx 1.21. \quad (18)$$

In words, accounting for the agent’s initial disclosure incentives, the principal is better off with renegotiation if and only if $\bar{\Delta}$ is sufficiently small with respect to the agent’s bias β . To understand this property, note first that if $\bar{\Delta} \leq \Delta^*$, the agent’s optimal disclosure in the contracting phase is not affected by the possibility to renegotiate later: in either case, the optimal awareness gap for the agent is $(\hat{y} - \bar{\Delta}, \hat{y} + \bar{\Delta})$ and the resulting delegation

set is $D^*(Y^*) = [0, \hat{y} - \bar{\Delta}] \cup \{\hat{y} + \bar{\Delta}\}$. Given that the principal prefers to have the gap in $D^*(Y^*)$ closed, she is strictly better off when renegotiation is allowed. Moreover, the larger the bias β , the larger Δ^* and hence the gain from renegotiation. If instead $\bar{\Delta} > \Delta^*$, the possibility to renegotiate gives the agent incentives to leave the principal unaware of more actions in the contracting phase than in the case of full commitment, thus a tradeoff arises. The principal's expected payoff in the case of pure delegation is now $V^P([0, \hat{y} - \Delta^*] \cup \{\hat{y} + \Delta^*\})$ (independent of $\bar{\Delta}$), while her expected payoff in the case of renegotiation is $V^P([0, \hat{y} + \bar{\Delta}])$.¹⁵ The ranking of these two payoffs depends on how high $\hat{y} + \bar{\Delta}$ is, i.e., how much additional flexibility the agent gains in the case of renegotiation. Condition (18) is equivalent to $V^P([0, \hat{y} + \bar{\Delta}]) \geq V^P([0, \hat{y} - \Delta^*] \cup \{\hat{y} + \Delta^*\})$.

5 Conclusion

This paper formulates a flexible delegation model with limited awareness and derives several properties of the optimal solution. The solution shows that by leaving the principal unaware of moderate options, the agent makes it optimal for the principal to permit actions closer to his own preferences. As argued in the Introduction, our framework has interesting implications for applications. We however believe that a key component of the contribution is to provide at least three general insights that apply to games with a principal-agent structure where the agent has superior awareness over feasible actions.

First, the paper illustrates that limited awareness can impose natural constraints on the language of contracts and that such limits may be exploited by the contracting party with superior awareness. This principle is not restricted to delegation problems but applies to other contracting problems. A principal facing a privately informed agent must resolve a tradeoff between exploiting the agent's private information and limiting the agent's information rents. The distortions solving this tradeoff are optimal for the principal but not for the agent. By manipulating the principal's awareness set and hence

¹⁵The closed form expressions are:

$$\begin{aligned} V^P([0, \hat{y} - \Delta^*] \cup \{\hat{y} + \Delta^*\}) &= -\beta^2 + \frac{8}{3}(3 - 2\sqrt{2})\beta^3 \\ V^P([0, \hat{y} + \bar{\Delta}]) &= -(1 - 2\beta + \bar{\Delta})\beta^2 + \frac{1}{3}[(\bar{\Delta} - \beta)^3 - \beta^3]. \end{aligned}$$

the set of feasible contracts, the agent can increase the principal's cost of such distortions, thereby increasing the principal's willingness to grant the agent higher information rents. The unconstrained solution to the agent's disclosure problem determines the maximal information rents he can get by modifying the set of feasible actions.

Second, the paper shows how in contracting situations with limited awareness, the option of renegotiation may be used as a tool to implement outcomes that are not describable at the contracting stage, without giving the other party blanket approval for unknown actions. While the existing literature largely focuses on the costs caused by the *impossibility of avoiding ex-post renegotiations* in the presence of ex-ante specific investments, we thus see *renegotiation as an opportunity* for the principal to improve the outcome. The downside of ex-post renegotiation for the principal in our setting is the reduction in the agent's incentives to disclose actions ex ante, giving rise to an interesting tradeoff. A clever design of the renegotiation process may shift this tradeoff further in favor of renegotiation (see Aghion et al. (1994) and Hart and Moore (2004)). An intriguing general question is indeed what a designer with limited awareness can achieve with mechanisms that are expressible in her language. The current paper may be viewed as a step in that direction.

Third, our modification of the game with renegotiation exemplifies how unawareness changes the ways in which agents infer information. If a player is unaware of the set of possible signals and only becomes aware of the signal s/he observes, the player cannot infer information from the fact that a different signal did not realize. This asymmetry gives rise to non-standard information structures and hence to rather different equilibrium outcomes with respect to the full awareness benchmark.

A Appendix

A.1 Proof of Proposition 1

Let $t : (Y^A)^2 \rightarrow [0, 1]$ be a symmetric function, indicating *the state at which the agent is indifferent between any two actions y and y'* . It is specified as follows. For $y = y'$, set $t(y, y') = s(y)$ (recall that $s(\cdot)$ is the inverse of $y^A(\cdot)$). For $y < y'$, $t(y, y')$ is defined by:

- if $U^A(\theta, y) < U^A(\theta, y')$ for all $\theta \in [0, 1]$, then $t(y, y') = 0$;
- if $U^A(\theta, y) > U^A(\theta, y')$ for all $\theta \in [0, 1]$, then $t(y, y') = 1$;
- otherwise $t(y, y')$ is such that

$$U^A(t(y, y'), y) = U^A(t(y, y'), y'). \quad (19)$$

Due to the single-crossing condition, the solution of (19) is unique. For $y > y'$, $t(y, y')$ is pinned down by the symmetry condition $t(y, y') = t(y', y)$. The following lemma links the slope of s with a partial derivative of t .

Lemma 8. *Consider y_0 such that $s(y_0) \in (0, 1)$, then*

$$\lim_{y \rightarrow y_0} \frac{dt(y, y_0)}{dy} = \frac{1}{2} s'(y_0)$$

Proof. For the case where t is determined by (19), we apply the implicit function theorem to derive

$$\frac{dt(y, y_0)}{dy} = \frac{U_y^A(t(y, y_0), y)}{U_\theta^A(t(y, y_0), y_0) - U_\theta^A(t(y, y_0), y)}$$

Taking the limit, we have

$$\begin{aligned} \lim_{y \rightarrow y_0} \frac{dt(y, y_0)}{dy} &= \lim_{y \rightarrow y_0} \frac{U_y^A(t(y, y_0), y)}{U_\theta^A(t(y, y_0), y_0) - U_\theta^A(t(y, y_0), y)} \\ &= \lim_{y \rightarrow y_0} \frac{U_{\theta y}^A(t(y, y_0), y) \frac{dt(y, y_0)}{dy} + U_{yy}^A(t(y, y_0), y)}{(U_{\theta\theta}^A(t(y, y_0), y_0) - U_{\theta\theta}^A(t(y, y_0), y)) \frac{dt(y, y_0)}{dy} - U_{\theta y}^A(t(y, y_0), y)} \\ &= \frac{U_{\theta y}^A(s(y_0), y_0) \lim_{y \rightarrow y_0} \frac{dt(y, y_0)}{dy} + U_{yy}^A(s(y_0), y_0)}{-U_{\theta y}^A(s(y_0), y_0)} \end{aligned}$$

where the second equality follows from L'Hôpital's rule. Also recall $t(y_0, y_0) = s(y_0)$.

We can solve the above equality for $\lim_{y \rightarrow y_0} \frac{dt(y, y_0)}{dy}$ and obtain

$$\lim_{y \rightarrow y_0} \frac{dt(y, y_0)}{dy} = -\frac{1}{2} \cdot \frac{U_{yy}^A(s(y_0), y_0)}{U_{y\theta}^A(s(y_0), y_0)}$$

From $U_y^A(s(y), y) = 0$, we derive via the implicit function theorem:

$$s'(y) = -\frac{U_{yy}^A(s(y), y)}{U_{\theta y}^A(s(y), y)}$$

which is well defined given our assumption that $U^A(\theta, y)$ is in C^2 and $U_{\theta y}^A > 0$. The two results together establish the claim. \square

Define $\bar{D}(y) := D^*(Y^A) \cap [y^A(0), y]$ as the set obtained by capping the optimal delegation set under full awareness at y .

Lemma 9. *If (6) holds as a strict inequality, there exists some $\underline{y} < \hat{y}$ such that for all $y \in (\underline{y}, \hat{y}) \cap D^*(Y^A)$,*

$$V^P(\bar{D}(y)) < V^P(\bar{D}(y) \cup \{\hat{y}\}).$$

Proof. We define the difference between the principal's expected payoffs when adding action \hat{y} to a delegation set whose highest action is $y < \hat{y}$. Given the strict monotonicity of the agent's preferred action in θ , adding action \hat{y} to a delegation set D with $\max D = y < \hat{y}$ only changes the outcome in the states where the agent optimally switches from y to \hat{y} . The set of states where this happens is $(t(y, \hat{y}), 1]$, so the payoff difference is:

$$\Delta \widehat{W}(y) := \int_{t(y, \hat{y})}^1 U^P(\theta, \hat{y}) dF(\theta) - \int_{t(y, \hat{y})}^1 U^P(\theta, y) dF(\theta)$$

Note that for all $y \in (\underline{y}, \hat{y}) \cap D^*(Y^A)$, we have

$$V^P(\bar{D}(y) \cup \{\hat{y}\}) - V^P(\bar{D}(y)) = \Delta \widehat{W}(y)$$

We calculate the first derivative of $\Delta \widehat{W}(\cdot)$ and evaluate it at \hat{y} :

$$\begin{aligned} \Delta \widehat{W}'(y) &= - \int_{t(y, \hat{y})}^1 U_y^P(\theta, y) dF(\theta) - (U^P(t(y, \hat{y}), \hat{y}) - U^P(t(y, \hat{y}), y)) f(t(y, \hat{y})) \frac{dt(y, \hat{y})}{dy} \\ \Delta \widehat{W}'(\hat{y}) &= - \int_{s(\hat{y})}^1 U_y^P(\theta, \hat{y}) dF(\theta) \end{aligned}$$

By (5), the above term is equal to zero. We must therefore consider the second derivative:

$$\begin{aligned}
\Delta\widehat{W}''(y) &= - \int_{t(y,\hat{y})}^1 U_{yy}^P(\theta, y) dF(\theta) + 2U_y^P(t(y,\hat{y}), y) f(t(y,\hat{y})) \frac{dt(y,\hat{y})}{dy} \\
&\quad - (U_\theta^P(t(y,\hat{y}), \hat{y}) - U_\theta^P(t(y,\hat{y}), y)) f(t(y,\hat{y})) \left(\frac{dt(y,\hat{y})}{dy} \right)^2 \\
&\quad - (U^P(t(y,\hat{y}), \hat{y}) - U^P(t(y,\hat{y}), y)) \left(f'(t(y,\hat{y})) \frac{dt(y,\hat{y})}{dy} + f(t(y,\hat{y})) \frac{d^2 t(y,\hat{y})}{dy^2} \right) \\
\Delta\widehat{W}''(\hat{y}) &= - \int_{s(\hat{y})}^1 U_{yy}^P(\theta, \hat{y}) dF(\theta) + 2U_y^P(s(\hat{y}), \hat{y}) f(s(\hat{y})) \frac{dt(y,\hat{y})}{dy} \Big|_{y=\hat{y}}
\end{aligned}$$

Since, by Lemma 8, we have $\frac{dt(y,\hat{y})}{dy} \Big|_{y=\hat{y}} = \frac{1}{2} s'(\hat{y})$, condition (6) holding as a strict inequality implies $\Delta\widehat{W}''(\hat{y}) > 0$. Remembering $\Delta\widehat{W}'(\hat{y}) = 0$, there is then an interval for y to the left of \hat{y} , where $\Delta\widehat{W}'(y) < 0$. With $\Delta\widehat{W}(\hat{y}) = 0$, this property implies, in turn, that there is some $\underline{y} < \hat{y}$ such that $\Delta\widehat{W}(y) > 0$ for all $y \in (\underline{y}, \hat{y})$. Hence, for all $y \in (\underline{y}, \hat{y}) \cap D^*(Y^A)$, we have $V^P(\bar{D}(y) \cup \{\hat{y}\}) - V^P(\bar{D}(y)) = \Delta\widehat{W}(y) > 0$. \square

Perturbation. We now construct a sequence of perturbations of Y^A as awareness levels for the principal, which are candidates for generating an improvement for the agent compared to full awareness.

Lemma 9 shows that there exists some $\underline{y} < \hat{y}$, such that for all $y \in (\underline{y}, \hat{y}) \cap D^*(Y^A)$, $V^P(\bar{D}(y)) < V^P(\bar{D}(y) \cup \{\hat{y}\})$ is satisfied. Hence, for each $y \in (\underline{y}, \hat{y}) \cap D^*(Y^A)$ there exists some $\delta_y > 0$ such that for all $\delta \in [0, \delta_y)$, $V^P(\bar{D}(y)) \leq V^P(\bar{D}(y) \cup \{\hat{y} + \delta\})$ holds.¹⁶ Setting $\delta_{\hat{y}} = 0$, we can then find a continuous, strictly decreasing function $\bar{\delta}$ which maps each $y \in (\underline{y}, \hat{y}) \cap D^*(Y)$ to a value of δ satisfying $V^P(\bar{D}(y)) \leq V^P(\bar{D}(y) \cup \{\hat{y} + \delta\})$ with image $\Phi := \{\delta \geq 0 : \bar{\delta}(y) = \delta \text{ for some } y \in (\underline{y}, \hat{y}) \cap D^*(Y^A)\}$. Letting $\bar{y}(\cdot)$ denote the inverse of $\bar{\delta}(\cdot)$, we define for each $\delta \in \Phi$ the associated awareness set by

$$Y(\delta) := Y^A \setminus (\bar{y}(\delta), \hat{y} + \delta).$$

Principal Optimality. In the last stage of the game, the agent's best response is described by $y^*(\theta, D)$. Given that the agent chooses according to y^* , the principal with awareness $Y \in \mathcal{Y}$ optimally selects a delegation set $D^*(Y) \subseteq Y$ to solve (3). In Proposition 12 we showed that the principal's value $V^P(\cdot)$ is continuous in D , where distances

¹⁶Since the agent simply chooses the better of the two actions closest to him, it is immediate to see that $V^P(\bar{D}(y) \cup \{\hat{y}\})$ is continuous in \hat{y} .

in D are defined according to the Hausdorff-metric d_H .

Let $y_1(\delta) = \max\{\tilde{y} \in D^*(Y(\delta)) : \tilde{y} \leq \hat{y}\}$ with $y_1(0) = \hat{y}$. We want to show that $y_1(\cdot)$ is continuous in δ on a right neighborhood of zero of its domain Φ . Suppose this is not true. Then, since $y_1(\cdot)$ is bounded above by \hat{y} , there exists a (sub)sequence $\{\delta_n\}$ with $\lim_{n \rightarrow +\infty} \delta_n = 0$ such that $\lim_{n \rightarrow +\infty} y_1(\delta_n) = y_0 \leq \hat{y}$, with y_0 possibly depending on the sequence. To violate continuity, one of them must satisfy $y_0 < \hat{y}$. Denote this sequence by $\{\tilde{\delta}_n\}$. Then,

$$d_H(D \setminus (y_0, \hat{y}), D^*(Y^A)) \geq \frac{\hat{y} - y_0}{2} \quad \forall D \in \mathcal{D}.$$

By continuity of V^P and uniqueness of the solution of (3), for all D and n sufficiently large, $V^P(\bar{D}(y(\tilde{\delta}_n))) > V^P(D \setminus (y_0, \hat{y}))$ is satisfied, which is a contradiction to $y_1(\tilde{\delta}_n) \in D^*(Y(\tilde{\delta}_n))$. Hence, $y_1(\cdot)$ is continuous on a right neighborhood of 0.

The principal's optimization regarding the inclusion of actions below $y_1(\delta)$ is equivalent to that under full awareness, as their potential inclusion only affects the agent's choice in states below $s(y_1(\delta))$. Indeed, given delegation set D and state θ , the agent has to consider at most two actions, which are the points in D on the left and right from his preferred action $y^A(\theta)$. Conditional on $y_1(\delta)$ belonging to the delegation set, the principal's design problem for actions below $y_1(\delta)$ can thus be separated from that for actions above. Hence, the optimal delegation set under awareness $Y(\delta)$ satisfies $\bar{D}(y_1(\delta)) \subseteq D^*(Y(\delta))$.

Next, we show that the principal permits at least one action above \hat{y} when having awareness $Y(\delta)$ for all $\delta \in \Phi$ sufficiently close to zero. To see this notice that for all $\delta \in \Phi$ sufficiently close to zero and all D with $\max D \leq \bar{y}(\delta)$, the following inequalities hold:

$$V^P(D) \leq V(\bar{D}(\bar{y}(\delta))) \leq V^P(\bar{D}(y(\delta)) \cup \{\hat{y} + \delta\}).$$

The last inequality has been established above. The first inequality follows from the facts that: i) by continuity of V^P and the fact that $D^*(Y^A)$ is a unique maximizer, generating a payoff close to $V^P(D^*(Y^A))$ requires that $\max D$ is close to \hat{y} , ii) conditional on $\max D$ being close to \hat{y} , we have $V^P(D) \leq V^P(\bar{D}(\max D))$, and iii) $V^P(\bar{D}(y))$ is increasing on a left neighborhood of \hat{y} (and $\bar{y}(\delta) \geq \max D$). We thus established $\max D^*(Y(\delta)) > \hat{y}$ for $\delta \in \Phi$ sufficiently close to zero.

With this observation, we can define $y_2(\delta) := \max\{\tilde{y} \in D^*(Y(\delta)) : \tilde{y} > \hat{y}\}$. By an

analogous argument to the one above, y_2 is continuous in $\delta \in \Phi$ on a right neighborhood of zero. Since $y_2(\delta)$ is bounded below by $\hat{y} + \delta$ and satisfies $y_2(0) = \hat{y}$, it must be increasing on a right neighborhood of 0. For all $\delta \in \Phi$ which are sufficiently close to zero, $D^*(Y(\delta))$ satisfies

$$\bar{D}(y_1(\delta)) \cup \{y_2(\delta)\} \subseteq D^*(Y(\delta)). \quad (20)$$

Agent optimality. Let \hat{V}^A denote the agent's value as a function of the delegation set. Property (20) implies that $\hat{V}^A(\bar{D}(y_1(\delta)) \cup \{y_2(\delta)\})$ constitutes a lower bound for the payoff the agent obtains when the principal's awareness set is $Y(\delta)$: additional actions in $D^*(Y(\delta))$ can only benefit the agent. For ease of notation, we change variables and write $y^+ = y_2(\delta)$ and $y^-(y^+) = y_1(y_2^{-1}(y^+))$. The agent's expected payoff for the delegation set $\bar{D}^+(y^+) := \bar{D}(y^-(y^+)) \cup \{y^+\}$ can then be written as:

$$\begin{aligned} \hat{V}^A(\bar{D}^+(y^+)) &= \int_0^{s(y^-(y^+))} U^A(\theta, y^*(\theta, D^*(Y^A))) dF(\theta) + \int_{s(y^-(y^+))}^{t(y^-(y^+), y^+)} U^A(\theta, y^-(y^+)) dF(\theta) \\ &\quad + \int_{t(y^-(y^+), y^+)}^1 U^A(\theta, y^+) dF(\theta). \end{aligned}$$

The first derivative of this payoff with respect to y^+ is

$$\frac{d \hat{V}^A(\bar{D}^+(y^+))}{d y^+} = \int_{s(y^-(y^+))}^{t(y^-(y^+), y^+)} U_y^A(\theta, y^-(y^+)) \frac{d y^-(y^+)}{d y^+} dF(\theta) + \int_{t(y^-(y^+), y^+)}^1 U_y^A(\theta, y^+) dF(\theta).$$

Evaluated at $y^+ = \hat{y}$, this derivative is equal to:

$$\left. \frac{d \hat{V}^A(\bar{D}^+(y^+))}{d y^+} \right|_{y^+ = \hat{y}} = \int_{s(\hat{y})}^1 U_y^A(\theta, \hat{y}) dF(\theta).$$

Since $U_y^A(s(\hat{y}), \hat{y}) = 0$ and $U_{\theta y}^A > 0$, we have $U_y^A(\theta, y) > 0$ for all $\theta > s(\hat{y})$. The derivative of the agent's value at $y^+ = \hat{y}$ is thus strictly positive. Hence, we can find a positive $\delta \in \Phi$ sufficiently close to zero, and an associated $y^+ = y_2(\delta) > \hat{y}$, such that $\hat{V}^A(\bar{D}(y_1(\delta)) \cup \{y_2(\delta)\}) = \hat{V}^A(\bar{D}^+(y^+)) > \hat{V}^A(\bar{D}^+(\hat{y})) = \hat{V}^A(\bar{D}(\hat{y})) = V^A(Y^A)$. Revealing all actions in Y^A is thus strictly dominated for the agent. \square

Changing bias. The result can be extended to the case where the agent is downward biased for a set of low states. Suppose there is state $\bar{\theta} \in (0, 1)$ such that $y^A(\theta) < y^P(\theta)$ for all $\theta < \bar{\theta}$ and $y^A(\theta) > y^P(\theta)$ for all $\theta > \bar{\theta}$. In this case, the optimal full awareness delegation set $D^*(Y^A)$ restricts the agent's choice from below: $\check{y} := \min D^*(Y^A) > y^A(0)$. Assume \check{y} is a limit point of $D^*(Y^A)$ and the analogue of condition (6) holds, that is:

$$U_y^P(s(\check{y}), \check{y})f(s(\check{y}))s'(\check{y}) - \int_0^{s(\check{y})} U_{yy}^P(\theta, \check{y})dF(\theta) < 0.$$

Following steps analogous to those above, we can then find a sequence of perturbations of Y^A —this time introducing a gap around \check{y} —inducing delegation sets that include an action $y^- < \check{y}$ as well as all elements of $D^*(y)$ above some threshold $y^+(y^-)$. Letting $\bar{D}^-(y^-)$ denote such delegation set, the agent's expected payoff as a function of y^- is

$$\begin{aligned} \hat{V}^A(\bar{D}^-(y^-)) &= \int_0^{t(y^-, y^+(y^-))} U^A(\theta, y^-)dF(\theta) + \int_{t(y^-, y^+(y^-))}^{s(y^+(y^-))} U^A(\theta, y^+(y^-))dF(\theta) \\ &\quad + \int_{s(y^+(y^-))}^1 U^A(\theta, y^*(\theta, D^*(Y^A)))dF(\theta). \end{aligned}$$

Taking the derivative with respect to y^- and evaluating it at $y^- = \check{y}$ yields

$$\left. \frac{d\hat{V}^A(\bar{D}^-(y^-))}{dy^-} \right|_{y^- = \check{y}} = \int_0^{s(\check{y})} U_y^A(\theta, \check{y})dF(\theta).$$

Since $U_y^A(s(\check{y}), \check{y}) = 0$ and $U_{\theta\check{y}}^A > 0$, we have $U_y^A(\theta, \check{y}) < 0$ for all $\theta < s(\check{y})$, so the derivative is strictly negative. This means that perturbing the full awareness set by introducing a gap around \check{y} is strictly optimal for the agent.

A.2 Proof of Proposition 2

First, we prove statement (i). Noticing that under Assumption 1 the upper threshold \hat{y} is a limit point of $D^*(Y^A)$, the claim that $\hat{y} \notin Y^P$ implies $Y^* \neq Y^A$ follows from Proposition 1. We thus want to prove the converse of this claim, namely that $\hat{y} \in Y^P$ implies $Y^* = Y^A$. Towards a contradiction, suppose this is not true. Then there exists an awareness set Y with $\hat{y} \in Y$ such that the agent strictly prefers $D^*(Y)$ over $[y^A(0), \hat{y}]$. This implies that $D^*(Y)$ contains a non-empty set of actions \tilde{X} such that $x > \hat{y}$ for all

$x \in \tilde{X}$. Since $D^*(Y^A)$ is the largest optimal delegation set in $\mathcal{D}(Y^A)$,

$$V^P(D^*(Y^A)) > V^P(D^*(Y^A) \cup \tilde{X})$$

holds. Monotonicity of the agent's policy then implies that conditional on permitting action \hat{y} , the principal is strictly better off by removing all actions in \tilde{X} . Hence, \hat{y} cannot belong to $D^*(Y)$. Since restricting the agent's choice from below is never optimal, we have $\hat{y} > \min D^*(Y)$ and hence $\hat{y} \in \text{Conv}(D^*(Y))$. Assumption 1 then implies $V^P(D^*(Y) \cup \{\hat{y}\}) \geq V^P(D^*(Y))$, a contradiction. Disclosing all actions is thus optimal for the agent.

To prove statement (ii), notice first that if $\hat{y} \in Y^P$, then $\Delta_1 = \Delta_2 = 0$. Consider then the case $\hat{y} \notin Y^P$. We start by showing that the principal permits at most one action weakly greater than \hat{y} . Suppose instead there is an awareness set Y given which the principal optimally delegates a set $D^*(Y)$ which contains two distinct actions weakly greater than \hat{y} and let \bar{y} be the largest action in $D^*(Y)$. Given that (7) is single-peaked in y , we know that for any $y \in (\hat{y}, \bar{y})$ we have $V^P([y^A(0), y]) > V^P([y^A(0), \bar{y}])$ and hence:

$$\int_{s(y)}^1 U^P(\theta, y) dF(\theta) > \int_{s(y)}^{s(\bar{y})} U^P(\theta, y^A(\theta)) dF(\theta) + \int_{s(\bar{y})}^1 U^P(\theta, \bar{y}) dF(\theta). \quad (21)$$

This inequality, together with the assumption that permitting \bar{y} is optimal for the principal, implies that \bar{y} cannot be a limit point of $D^*(Y)$. Consider then action $y > \hat{y}$ such that $y < \bar{y}$ and $[y, \bar{y}] \cap D^*(Y) = \{y, \bar{y}\}$. We can show the following:

$$\begin{aligned} & V^P(D^*(Y)) \\ &= \int_0^{s(y)} U^P(\theta, y^*(\theta, D^*(Y))) dF(\theta) + \int_{s(y)}^{t(y, \bar{y})} U^P(\theta, y) dF(\theta) + \int_{t(y, \bar{y})}^1 U^P(\theta, \bar{y}) dF(\theta) \\ &\leq \int_0^{s(y)} U^P(\theta, y^*(\theta, D^*(Y))) dF(\theta) + \int_{s(y)}^{s(\bar{y})} U^P(\theta, y^A(\theta)) dF(\theta) + \int_{s(\bar{y})}^1 U^P(\theta, \bar{y}) dF(\theta) \\ &< \int_0^{s(y)} U^P(\theta, y^*(\theta, D^*(Y))) dF(\theta) + \int_{s(y)}^1 U^P(\theta, y) dF(\theta) \\ &= V^P(D^*(Y) \setminus \{\bar{y}\}) \end{aligned}$$

where the weak inequality follows from Assumption 1 and the strict inequality follows from (21). Taken together, these inequalities imply that the principal can strictly improve

her payoff by removing action \bar{y} from the delegation set, which yields the contradiction.

Consider now an optimal awareness set Y^* . The set Y^* must clearly satisfy $\max D^*(Y^*) > \hat{y}$, as any delegation set with an upper bound weakly smaller than \hat{y} is dominated by the full awareness delegation set and $D^*(Y^*) \neq D^*(Y^A)$. We set $\Delta_2 = \max D^*(Y^*) - \hat{y} > 0$. We know $[\hat{y}, \hat{y} + \Delta_2) \cap D^*(Y^*) = \emptyset$, because we have shown that the principal allows at most one action above \hat{y} . Assumption 1 and the fact that the agent is upward biased imply that the principal permits all actions in Y^* that are weakly smaller than $\hat{y} + \Delta_2$. Define $\Delta_1 \geq 0$ as the smallest value of Δ that satisfies the inequality

$$\begin{aligned} & \int_{s(\hat{y}-\Delta)}^1 U^P(\theta, \hat{y} - \Delta) dF(\theta) \\ \leq & \int_{s(\hat{y}-\Delta)}^{t(\hat{y}-\Delta, \hat{y}+\Delta_2)} U^P(\theta, \hat{y} - \Delta) dF(\theta) + \int_{t(\hat{y}-\Delta, \hat{y}+\Delta_2)}^1 U^P(\theta, \hat{y} + \Delta_2) dF(\theta) \end{aligned}$$

By continuity of $U^P(\cdot, \cdot)$, $t(\cdot, \cdot)$ and $s(\cdot)$ (due to the continuity of $U^A(\cdot, \cdot)$ and recalling our tie-breaking rule), Δ_1 is well-defined. Since $\Delta_2 > 0$, also Δ_1 is positive.

It is easy to see that Y^* must satisfy $Y^* \cap (\hat{y} - \Delta_1, \hat{y}) = \emptyset$. If instead there is an action $y \in (\hat{y} - \Delta_1, \hat{y})$ such that $y \in Y^*$, then Assumption 1 implies $y \in D^*(Y^*)$, since y belongs to the convex hull $D^*(Y^*)$. However, by definition of Δ_1 , given that the principal permits y , she strictly prefers not to permit $\hat{y} + \Delta_2$, a contradiction. Assumption 1, together with the fact that the agent is upward biased, further implies that, conditional on $\hat{y} - \Delta_1 \in D^*(Y^*)$, revealing any action below $\hat{y} - \Delta_1$ and weakly above $y^A(0)$ results in the inclusion of that action in the delegation set and thus strictly benefits the agent. Optimality of Y^* thus requires $Y^* \cap [y^A(0), \hat{y} + \Delta_2] = [y^A(0), \hat{y} - \Delta_1] \cup \{\hat{y} + \Delta_2\}$. Given such awareness set, the principal optimally chooses the delegation set $D^*(Y^*) = [y^A(0), \hat{y} - \Delta_1] \cup \{\hat{y} + \Delta_2\}$. \square

A.3 Proof of Proposition 3

We first want to verify that the conditions of statement (ii) in Proposition 2 are satisfied. It is well known that condition (9) implies Assumption 1 (see Martimort and Semenov (2006)). What remains to show is that (7) is strictly quasiconcave in y . For the case

considered here, we can write the principal's payoff in (7) as

$$V^P([0, y]) = -F(y)\beta^2 - \int_y^1 (y - (\theta - \beta))^2 dF(\theta), \quad (22)$$

with the first and second derivative respectively given by

$$\frac{dV^P([0, y])}{dy} = -2 \int_y^1 (y - (\theta - \beta)) dF(\theta)$$

and

$$\frac{d^2V^P([0, y])}{dy^2} = 2(\beta f(y) - (1 - F(y))).$$

Note that by condition (9) the second derivative is strictly increasing in y , which in turn implies that the first derivative is strictly convex. At the boundary points $y = 0$ and $y = 1$, the first derivative, respectively, takes the values

$$\left. \frac{dV^P([0, y])}{dy} \right|_{y=0} = 2\mathbb{E}[\theta - \beta] > 0 \quad \text{and} \quad \left. \frac{dV^P([0, y])}{dy} \right|_{y=1} = 0.$$

Together with strict convexity, this implies that $\frac{dV^P([0, y])}{dy}$ has exactly one intersection with zero on $[0, 1]$ at \hat{y} . That is, for all $y < \hat{y}$, $\frac{dV^P([0, y])}{dy} > 0$, for all $y \in (\hat{y}, 1)$ $\frac{dV^P([0, y])}{dy} < 0$. Hence, the function $V^P([0, y])$ is strictly quasi-concave.

Given quasi-concavity, we know that the equilibrium disclosure and delegation sets are described by two parameters Δ_1 and Δ_2 , as described in Proposition 2. Next, we want to show that $\Delta_1 = \Delta_2 = \Delta$. In equilibrium, the principal must be indifferent between delegation sets $[y_{min}, \hat{y} - \Delta_1]$ and $[y_{min}, \hat{y} - \Delta_1] \cup \{\hat{y} + \Delta_2\}$, as otherwise the agent could reduce Δ_1 without destroying the principal's incentives to permit $\hat{y} + \Delta_2$ and thereby obtain a strictly higher payoff. Let $t := \hat{y} + \frac{\Delta_2 - \Delta_1}{2}$ denote the state at which the agent is indifferent between the two actions. The change in the principal's payoff when adding action $\hat{y} + \Delta_2$ to the set $[y_{min}, \hat{y} - \Delta_1]$ is given by

$$\begin{aligned} & - \int_t^1 (\hat{y} + \Delta_2 - \theta + \beta)^2 dF(\theta) + \int_t^1 (\hat{y} - \Delta_1 - \theta + \beta)^2 dF(\theta) \\ &= -2(\Delta_1 + \Delta_2) \int_t^1 (t - (\theta - \beta)) dF(\theta) \\ &= (\Delta_1 + \Delta_2) \left. \frac{dV^P([0, y])}{dy} \right|_{y=t} \end{aligned}$$

which, as we just argued, is weakly positive if and only if $t \leq \hat{y}$. For the principal to be indifferent, we must have $t = \hat{y}$ or equivalently $\Delta_1 = \Delta_2 = \Delta$.

Having shown this property, the agent's optimization problem boils down to the choice of $\Delta \leq \bar{\Delta}(Y^P)$. The agent's payoff as a function of Δ is specified in (13). Taking the first derivative with respect to Δ yields the first-order condition (11). The second derivative of (13) is

$$-2(1 - F(\hat{y} - \Delta)) < 0,$$

so the agent's expected payoff is strictly concave in Δ . It follows that the unconstrained maximizer Δ^* is described by (11) and that the constrained maximizer is $\max\{\Delta^*, \bar{\Delta}(Y^P)\}$.

□

A.4 Proof of Proposition 4

Let $\hat{y}(\beta)$ be the optimal cap under full awareness when the bias is β . Hence, $\hat{y}(\beta)$ is the maximizer of (22) and is implicitly defined by the first-order condition

$$\int_{\hat{y}}^1 (\theta - \beta) dF(\theta) - \hat{y}(1 - F(\hat{y})) = 0.$$

As we argued in the proof of Proposition 3, the first derivative of $V^P([0, y])$ with respect to y is strictly convex and crosses \hat{y} from above. The second-order condition thus holds as a strict inequality:

$$\beta f(\hat{y}(\beta)) - (1 - F(\hat{y}(\beta))) < 0. \quad (23)$$

We can then use the implicit function theorem (condition (9) implies that the cumulate F is \mathcal{C}^1) to show that $\hat{y}(\beta)$ admits a derivative at each β , which equals:

$$\hat{y}'(\beta) = -\frac{1 - F(\hat{y}(\beta))}{\beta f(\hat{y}(\beta)) - (1 - F(\hat{y}(\beta)))} < 0 \quad (24)$$

Continuous differentiability is guaranteed by the implicit function theorem and can be checked directly in the above expression

Next, consider the agent's payoff as a function of Δ and the parameter β , as described in (13), and denote it by $U(\Delta; \beta)$. As we showed in the proof of Proposition 3, this payoff is strictly concave with the interior solution characterized by (11). The conditions for

applying the implicit function theorem are again satisfied, hence there is a function $\Delta^*(\beta)$ describing the unconstrained 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:

$$\Delta^{*\prime}(\beta) = -\frac{U_{\Delta\beta}(\Delta^*(\beta); \beta)}{U_{\Delta\Delta}(\Delta^*(\beta); \beta)}$$

To prove the statement of the proposition, we must then show $U_{\Delta\beta}(\Delta^*(\beta); \beta) > 0$. Differentiating the expression of the first-order condition (11) with respect to β keeping Δ^* fixed, after some rearrangement, delivers:

$$U_{\Delta\beta}(\Delta^*(\beta); \beta) = -\hat{y}'(\beta) [1 + F(\hat{y}(\beta) - \Delta^*(\beta)) - 2F(\hat{y}(\beta))].$$

Given $\hat{y}'(\beta) < 0$ (see (24)), we are done if $1 + F(\hat{y}(\beta) - \Delta^*(\beta)) - 2F(\hat{y}(\beta)) > 0$, or equivalently:

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

Using $\hat{y}(\beta) = \mathbb{E}[\theta - \beta | \theta \geq \hat{y}(\beta)]$, the first-order condition (11) can be written as:

$$[1 - F(\hat{y}(\beta) - \Delta^*(\beta))] [\mathbb{E}[\theta | \theta \geq \hat{y}(\beta) - \Delta^*(\beta)] - (\hat{y}(\beta) - \Delta^*(\beta))] = 2[1 - F(\hat{y}(\beta))]\beta. \quad (25)$$

Since $\hat{y}(\beta) - \Delta^*(\beta)$ is strictly smaller than $\hat{y}(\beta)$, the following condition holds:

$$\mathbb{E}[\theta - \beta | \theta \geq \hat{y}(\beta) - \Delta^*(\beta)] - (\hat{y}(\beta) - \Delta^*(\beta)) > 0.$$

Equivalently we can write

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

Given this inequality, (25) requires $2(1 - F(\hat{y}(\beta))) > 1 - F(\hat{y}(\beta) - \Delta^*(\beta))$, as desired. \square

A.5 Proof of Proposition 5

Proof. We begin with part (i). Consider a renegotiation equilibrium and let $\hat{X} := \{x \in Y : \rho^*(x) = 1\}$ be the set of actions in the principal's initial awareness set which are allowed in equilibrium. By upper semicontinuity of ρ^* , this set \hat{X} is closed. Since also

$D^*(Y)$ is closed, the set $D^*(Y) \cup \hat{X}$ is closed as well. We first want to show that

$$V^P(D^*(Y) \cup \hat{X}) \geq V^P(D^*(Y)).$$

Since the agent can always guarantee choices in $D^*(Y)$ by proposing N , agent's optimality in the eyes of the principal implies $x_x^*(\theta) = y^*(\theta, D^*(Y) \cup \hat{X})$ for all $x \in \hat{X}$ and $\theta \in \bigcup_{x \in \hat{X}} \Theta^*(x)$, with the exclusion of points where the agent is indifferent between two actions in $D^*(Y) \cup \hat{X}$. For each $x \in D^*(Y) \cup \hat{X}$, we thus have:

$$cl(\Theta^*(x)) = \{\theta \in [0, 1] : U^A(\theta, x) \geq \max_{y \in D^*(Y) \cup \hat{X}} U^A(\theta, y)\}$$

Using the fact that for all $x \in \hat{X}$ and all $\theta \in \Theta^*(x)$, we have $x = y^*(\theta, D^*(Y) \cup \hat{X})$ and that for each $B \subseteq \Theta^*(x)$ we have $\int_{\Theta^*(x)} \mu^*(B|x) f(\theta') d\theta' = \int_B dF(\theta)$,¹⁷ the principal's optimality condition yields

$$\begin{aligned} \int_{\Theta^*(x)} \int_{\Theta^*(x)} U^P(\theta, x) d\mu^*(\theta|x) dF(\theta') &\geq \int_{\Theta^*(x)} \int_{\Theta^*(x)} U^P(\theta, y^*(\theta, D^*(Y))) d\mu^*(\theta|x) dF(\theta') \\ \iff \int_{\Theta^*(x)} U^P(\theta, y^*(\theta|D^*(Y) \cup X)) dF(\theta) &\geq \int_{\Theta^*(x)} U^P(\theta, y^*(\theta, D^*(Y))) dF(\theta) \end{aligned}$$

for all $x \in \hat{X}$.

Next, for $\theta \notin \bigcup_{x \in \hat{X}} \Theta^*(x)$, we have $x_x^*(\theta) = N$ or $\rho^*(x_x^*(\theta)) = 0$. In either case, the agent will take $y^*(\theta, D^*(Y))$. We further have $y^*(\theta, D^*(Y)) = y^*(\theta, D^*(Y) \cup \hat{X})$ for all $\theta \notin \bigcup_{x \in \hat{X}} \Theta^*(x)$, since the agent is free to take any action in $D^*(Y) \cup \hat{X}$.

Setting $\Theta^C(\hat{X}) := [0, 1] \setminus (\bigcup_{x \in \hat{X}} \Theta^*(x))$, we can now integrate over \hat{X} and obtain

$$\begin{aligned} &\int_{\Theta^C(\hat{X})} U^P(\theta, y^*(\theta, D^*(Y) \cup \hat{X})) dF(\theta) + \int_{\hat{X}} \int_{\Theta^*(x)} U^P(\theta, y^*(\theta, D^*(Y) \cup \hat{X})) dF(\theta) dx \\ &\geq \int_{\Theta^C(\hat{X})} U^P(\theta, y^*(\theta, D^*(Y))) dF(\theta) + \int_{\hat{X}} \int_{\Theta^*(x)} U^P(\theta, y^*(\theta, D^*(Y))) dF(\theta) dx. \end{aligned}$$

or equivalently

$$V^P(D^*(Y) \cup \hat{X}) \geq V^P(D^*(Y)).$$

Recall that $D^*(Y) \cup \hat{X} \subseteq Y$ is a closed set. Since $D^*(Y)$ is the largest closed optimal

¹⁷Note that monotonicity of the agent's optimal policy y^* in θ implies that $\Theta^*(x)$ is either of positive measure or a singleton.

awareness set with respect to Y that includes actions that will actually be taken by the agent under some contingency, this inequality yields a contradiction unless $D^*(Y) \cup \hat{X} = D^*(Y)$.

Having shown that the principal only accepts additional actions in the renegotiation phase if they do not belong to Y , consider proposal $x \in Y^A \setminus Y$ such that $\rho^*(x) = 1$. Perceived agent optimality then requires $x_x^*(\theta) = x$ for all θ such that $U^A(\theta, x) > \max_{y \in D^*(Y)} U^A(\theta, y)$. Principal optimality in turn requires that conditioning on the event $U^A(\theta, x) > U^A(\theta, y^*(\theta, D^*(Y)))$, the principal prefers x over $y^*(\theta, D^*(Y))$ in expectation. This is the case only if $x \in A(Y)$.

To show part (ii), we need to construct an equilibrium of the renegotiation game where proposal x is accepted by the principal whenever $x \in A(Y)$. To this end, we set $\rho^*(x) = 1$ for all $x \in A(Y)$ and $\rho^*(x) = 0$ otherwise. Recall that $\rho^*(N) = 0$. For the agent we set $x^*(\theta) = \arg \max_{x \in A(Y)} U^A(\theta, x)$ if $\max_{x \in A(Y)} U^A(\theta, x) \geq \max_{y \in D^*(Y)} U^A(\theta, y)$ and $x^*(\theta) = N$ otherwise.¹⁸ Similarly, for the agent's strategy as perceived by the principal when receiving proposal x , we set $x_x^*(\theta) = x$ if $U(\theta, x) \geq \max_{y \in D^*(Y)} U^A(\theta, y)$ and $x_x^*(\theta) = N$ otherwise. The principal beliefs system μ^* is defined as follows. For all $x \in X$:

$$\mu^*(B|x) = \frac{\int_B dF(\theta)}{\int_{\Theta^*(x)} dF(\theta)} \quad \forall B \subseteq \Theta^*(x),$$

whenever $\Theta^*(x) = (x_x^*)^{-1}(x)$ is of positive measure, and $\mu^*(\{\theta\}|x) = 1$ if $\Theta^*(x) = \{\theta\}$. It can be checked directly that this strategy and belief profile satisfy principal optimality, agent optimality, and consistency of beliefs and thus constitute a PBE of the renegotiation game, as specified in Definition 1. \square

A.6 Proof of Proposition 6

Proof. Note that, under Assumption 1, for any $x \in (\min D^*(Y), \max D^*(Y))$, we have $V^P(D^*(Y) \cup \{x\}) \geq V^P(D^*(Y))$, so x belongs to the set of acceptable proposals. The same is true for all $x < \min D^*(Y)$, since, conditioning on the fact that the (upward biased) agent prefers x over $\min D^*(Y)$, the principal prefers x as well. The set of implementable action $A(Y)$ thus includes all action in $[y^A(0), \max D^*(Y)]$. Under Assumption 1 we

¹⁸Recall, in case of indifference the agent breaks ties in favour of the principal.

clearly have $V^P([y^A(0), \max D^*(Y)]) \geq V^P(D^*(Y))$. If $\max D^*(Y) = \max A(Y)$, this concludes the argument. For the other case, let $\bar{y} := \max A(Y)$ assume $\bar{y} > \max D^*(Y)$. By definition of A , we have $V^P(D^*(Y) \cup \{\bar{y}\}) \geq V^P(D^*(Y))$. Monotonicity of the agent's action in θ then implies

$$V^P([y^A(0), \max D^*(Y)] \cup \{\bar{y}\}) \geq V^P([y^A(0), \max D^*(Y)])$$

But given that permitting action \bar{y} weakly increases the principal's expected payoff, permitting any additional action in $(\max D^*(Y), \bar{y})$ benefits the principal as well (again by Assumption 1), so (16) is satisfied. \square

A.7 Proof of Proposition 7

Proof. We start by showing the existence of a solution of $\max_Y D^*(Y)$. Recall that $BR^P(\cdot)$ denotes the principal's solution correspondence for problem (3). Let us then define:

$$\check{y}(Y) := \sup_{D \in BR^P(Y)} \max D.$$

Since any $D \in BR^P(Y)$ is compact, the last max is well-defined. From the proof of Proposition 12 we know that BR^P is upper hemicontinuous. If we show that the function $m(D) := \max D$ is continuous in D , the generalized version of the maximum theorem (e.g., Theorem 17.30 in Aliprantis and Border (2006)) implies that the max exists for each Y and $\check{y}(Y)$ is upper semicontinuous in Y . This in turn implies that the following object is well-defined:

$$\bar{y}^* := \max_{Y^P \subseteq Y \subseteq Y^A} \check{y}(Y).$$

Lemma 10. *The function $m(D) := \max D$ is continuous with respect to the Hausdorff metric.*

Proof. Recall again that we are working with metric spaces. Take a sequence $D_n \rightarrow^H D$. Consider now the sequence of real numbers $d_n := \max D_n \forall n$. We need to show that the sequence converges to $d := \max D \in \mathbb{R}$. Since D_n converges, it is Cauchy. We want to show that also d_n is Cauchy. For any δ , let N_δ be such that $d_H(D_n, D_m) \leq \delta \forall n, m \geq N_\delta$. Now, if $d_n = d_m$ there is nothing to prove. Suppose $d_n \neq d_m$. Then, w.l.o.g assume

$d_n > d_m$. We have

$$|d_n - d_m| = |\max D_n - \max D_m| = \inf_{y \in D_m} |d_n - y| \leq d_H(D_n, D_m) \leq \delta.$$

Given that δ is generic, d_n is Cauchy and since \mathbb{R} is complete, the sequence d_n must converge. Let d^* be the converging point of the sequence. Again, w.l.o.g. assume $d^* > d$. But then, by the definition of convergence, it must be that for N large enough, $d_n > d$ for all $n \geq N$. Delivering

$$d_H(D, \tilde{D}_n) \geq \inf_{y \in D} |d_n - y| = |d_n - d| > 0 \quad \forall n \geq N,$$

which contradicts the fact that D_n converges to D . Hence, it must be that $d^* = d$. \square

Next, we want to show that for each Y , $\tilde{y}(Y)$ is equal to $\max D^*(Y)$. Suppose this is not true. Then there exists an awareness set \tilde{Y} such that $\tilde{y}(\tilde{Y}) > \max D^*(\tilde{Y})$ and a delegation set $\tilde{D} \in BR^P(\tilde{Y})$ such that $\max \tilde{D} > \max D^*(\tilde{Y})$. Since it is never optimal to curtail the agent's flexibility from below and Assumption 1 holds, we have

$$\min \tilde{D} = \min D^*(\tilde{Y}) = y^*(0, \tilde{Y})$$

Given $\min \tilde{D} = \min D^*(\tilde{Y}) \leq \max D^*(\tilde{Y}) < \max \tilde{D}$, Assumption 1 implies $V^P(\tilde{D} \cup D^*(\tilde{Y})) \geq V^P(\tilde{D})$ and hence $(\tilde{D} \cup D^*(\tilde{Y})) \in BR^P(\tilde{Y})$. But since D^* selects the agent-preferred delegation set from $BR^P(\tilde{Y})$ and since $D^*(\tilde{Y}) \subseteq (\tilde{D} \cup D^*(\tilde{Y}))$, we must have $D^*(\tilde{Y}) = \tilde{D}$ and hence $\max D^*(\tilde{Y}) = \max \tilde{D}$, a contradiction. Combining these results, we have shown that

$$\sup_{D \in BR^P(Y)} \max D = \max D^*(Y)$$

and hence $\bar{y}^* = \max_{Y^P \subseteq Y \subseteq Y^A} \max D^*(Y)$.

Let \bar{Y}^* be an awareness set that maximizes $\max D^*(Y)$ over Y subject to $Y^P \subseteq Y \subseteq Y^A$. We now want to show that \bar{Y}^* solves (17). Suppose not. Since

$$A(\bar{Y}^*) = [y^A(0), \bar{y}^*]$$

there must then exist an awareness set Y such that $\max(A(Y)) > \bar{y}^*$. By definition of

\bar{y}^* , there is no awareness set Y such that $Y^P \subseteq Y \subseteq Y^A$ and $\bar{y}^* \in D^*(Y)$. Hence, there must be a proposal $x > \bar{y}^*$ that the principal accepts in the renegotiation phase. By Proposition 5, this requires $x \in A(Y)$, or equivalently

$$V^P(D^*(Y) \cup \{x\}) \geq V^P(D^*(Y)).$$

This in turn implies $x \in D^*(Y \cup \{x\})$ and hence $\bar{y}^* \geq x$, a contradiction. \square

A.8 Existence results

Consider the following properties of our setup.

- (a) The set $Y \subseteq Y^A = [y_{min}, y_{max}]$ is a compact subset of the complete and separable metric space $(\mathbb{R}, |\cdot|)$.
- (b) Recall, \mathcal{Y} denotes the set of closed subsets of $[y_{min}, y_{max}]$, and

$$\hat{\mathcal{D}}(Y) := \{D \in \mathcal{Y} : D \subseteq Y\}.$$

Then $\hat{\mathcal{D}}(Y)$ is a closed subset of $2^{[Y]}$ with respect to the Hausdorff-metric

$$d_H(D, D') = \max \left\{ \sup_{y \in D} \inf_{y' \in D'} |y - y'|, \sup_{y' \in D'} \inf_{y \in D} |y - y'| \right\},$$

and hence compact in the topology generated by the Hausdorff metric d_H (see point 3 of Theorem 3.85 in Aliprantis and Border (2006)).

- (c) Recall: U^A and U^P are continuous and uniformly bounded on their domains $[0, 1] \times [y_{min}, y_{max}]$, and F admits a density.
- (d) Recall: $U_{\theta y}^A > 0, U_y^A(\theta, y) > 0, U_{yy}^A(\theta, y) < 0$. Hence, if we fix a closed set $D \subseteq Y^A$ and an open interval $O \subseteq Y^A \setminus D$, there is at most one value of θ such that $y^A(\theta) \in O$ and $BR^A(\theta, D)$ is not single-valued.
- (e) The set $(y_{min}, y_{max}) \setminus D$ is an open set, and hence it can be uniquely defined as a countable union of disjoint open intervals (e.g., Theorem 6, page 51, in Kolmogorov and Fomin (1975)).

Recall the agent chooses according to $BR^A(\theta, D) := \arg \max_{y \in D} U^A(y, \theta)$. Note that by continuity $BR^A(\theta, D)$ is non-empty for each θ , since $D \subseteq Y^A$ is compact from (a) and U^A is continuous from (c). In addition, since the feasibility set D changes continuously with D in the Hausdorff norm (and U^A is continuous in (D, y, θ)), the maximum theorem implies that BR^A is an upper hemicontinuous correspondence when seen as a function of (D, θ) . In addition, combining (d) with (e), and noticing that for all θ such that $y^A(\theta) \in D$ we have $BR^A(\theta, D) = \{y^A(\theta)\}$, we conclude that for each D the set of θ 's for which $BR^A(\theta, D)$ is not single-valued is countable. Since F admits a density, the set of values for which the agent is indifferent is then of F -measure zero. In summary:

Lemma 11. *BR^A is a non-empty upper hemicontinuous correspondence in (θ, D) . Moreover, for each D , the set $\Theta^A(D) := \{\theta \in [0, 1] \mid BR^A(\theta, D) \text{ is not a singleton}\}$ has measure zero according to F .*

Recall as well, we denoted with y^* the selection that resolves ties in favour of the principal. Due to Lemma 11, for each awareness set $Y \in \mathcal{Y}$, the principal optimally selects a delegation set $D \subseteq Y$ to solve

$$\max_{D \in \mathcal{D}(Y)} V^P(D) \quad \text{where} \quad V^P(D) = \int_0^1 U^P(\theta, y^*(\theta, D)) dF(\theta). \quad (26)$$

From Lemma 4 in Holmström (1980), V^P is upper semicontinuous in D for each closed $Y \subseteq Y^A$ (where distances in D are defined according to the Hausdorff-metric). Since according to this metric the feasibility set $\hat{\mathcal{D}}(Y)$ is compact, we have:

Proposition 12. *An optimal solution to the principal's problem in $\hat{\mathcal{D}}(Y)$ —and hence in $\mathcal{D}(Y)$ —exists. In addition, V^P is continuous in D .*

The existence of an optimal solution in $\mathcal{D}(Y)$ is guaranteed by the fact that for any solution to (26) in $\hat{\mathcal{D}}(Y)$ we obtain a solution in $\mathcal{D}(Y)$ by eliminating actions that the agent does not take in equilibrium. The continuity of V is implied by the fact that whenever an u.h.c. correspondence is single-valued it is continuous. Hence the second part of Lemma 11 implies that any selection h from BR^A will have discontinuities at a

set of points that have probability zero according to F . That is,

$$V^P(D) = \int_0^1 U^P(\theta, y^*(\theta, D))dF(\theta) = \int_0^1 U^P(\theta, h(\theta, D))dF(\theta),$$

and the latter varies continuously with D by the continuity property of the selection h and the continuity of U^P summarized in (c).

Denote by $BR^P(Y)$ the solution correspondence for the principal's problem. If we can show that the correspondence from Y to $\mathcal{D}(Y)$ is continuous, thanks to properties (a)-(e), we can apply the theorem of the maximum to show that $BR^P(Y)$ is upper hemicontinuous.

Now recall, we indicate with $D^*(Y)$ the selection from $BR^P(Y)$ that resolves ties in favor of the agent. The problem of the agent at the initial disclosure stage solves:

$$\max_{Y \in \mathcal{Y}} V^A(Y) \quad \text{where} \quad V^A(Y) = \int_0^1 U^A(\theta, y^*(\theta, D^*(Y)))dF(\theta) \quad \text{s.t.} \quad Y^P \subseteq Y \subseteq Y^A.$$

If we can show that the value $V^A(Y)$ is upper semicontinuous by the property of the selection D^* , we have a solution. Let us hence show the following result:

Proposition 13. *$BR^P(Y)$ is upper hemicontinuous and the problem of the agent has at least one solution.*

Proof. As argued above, there are two crucial steps. First, the upper semicontinuity of V^A and then the continuity of the correspondence $\mathcal{D}(Y)$.

Lemma 14. *The function V^A is upper semicontinuous in Y under the Hausdorff metric.*

Proof. Recall that D^* has the following property:

$$D^*(Y) = \arg \max_{D \in BR^P(Y)} \int_0^1 U^A(\theta, y^*(\theta, D))dF(\theta).$$

To show upper semicontinuity, take a converging sequence $Y_n \rightarrow^H \bar{Y}$ and suppose there is a sequence (in the real numbers) V_n converging to \bar{V} (in the $|\cdot|$ metric), where for each n

$$V_n = V^A(Y_n) = \int_0^1 U^A(\theta, y^*(\theta, D^*(Y_n)))dF(\theta).$$

We need to show that $\bar{V} \leq V^A(\bar{Y})$. To simplify notation, let $\hat{D}_n = D^*(Y_n)$. We hence have a sequence \hat{D}_n such that $\hat{V}_n = \int_0^1 U^A(\theta, y^*(\theta, \hat{D}_n))dF(\theta) \rightarrow \bar{V}$. Since \mathcal{Y} is compact,

there is a converging subsequence, $\hat{D}_{n'} \rightarrow^H \hat{D}$, and recall that $\hat{V}_{n'} \rightarrow \bar{V}$. Note now that the function $T(D) := \int_0^1 U^A(\theta, y^*(\theta, D))dF(\theta)$ is continuous in D based on the same arguments as in the proof of Proposition 12. It must hence be the case that $T(\hat{D}_{n'}) \rightarrow T(\hat{D})$. This implies that $\bar{V} = \int_0^1 U^A(\theta, y^*(\theta, \hat{D}))dF(\theta)$. Now, recall that BR^P is upper hemicontinuous, that is, it has a closed graph Gr . We have shown that \hat{D} is the limit of a sequence D_n such that $\hat{D}_n \in BR^P(Y_n)$ for all n , that is, $(\hat{D}_n, Y_n) \in Gr$ for all n . The limit point must be in the graph as well: $(\hat{D}, \bar{V}) \in Gr$. This is equivalent to saying that $\hat{D} \in BR^P(\bar{V})$. And hence, from the definition of D^* we have the desired inequality:

$$\bar{V} = \int_0^1 U^A(\theta, y^*(\theta, \hat{D}))dF(\theta) \leq \int_0^1 U^A(\theta, y^*(\theta, D^*(\bar{V})))dF(\theta).$$

□

Lemma 15. *The correspondence mapping each set Y from the metric space (\mathcal{Y}, d_H) to $\hat{D}(Y)$ is both u.h.c. and l.h.c.*

Proof. First of all, note that from $\hat{D}(Y)$, we have

$$D \in \hat{D}(Y) \iff D \subseteq Y.$$

Since we are working with metric spaces (and hence first countable topological spaces), we can prove our statement using sequences. (i) u.h.c.: take any $Y \in \mathcal{Y}$ and a generic converging sequence $Y_n \rightarrow^H Y$. Now, take a sequence D_n such that $D_n \subseteq Y_n$ for all n . We want to show that there is a subsequence D_{n_s} converging to $D \subseteq Y$. The existence of a converging sequence is implied by the compactness of the space. So let D be such a point. We need to show that $D \subseteq Y$. This is implied by the convergence condition $d_H(D_n, D) \rightarrow 0$ in the Hausdorff metric. Suppose $D \supset Y$. It must hence be that $d_H(D, Y) = \varepsilon > 0$, that is, the distance between the two sets is positive. Now, since both sequences converge, for each δ there is a N_δ such that for all $n \geq N_\delta$ we have both $d_H(D_n, D) \leq \delta$ and $d_H(Y_n, Y) \leq \delta$. This indicates that D_n cannot be smaller than the maximal reduction of D compatible with the distance and Y cannot be larger than the maximal extension of Y compatible with the distance. Such reductions and extensions can be made arbitrarily small, so if we have $D \supset Y$, we must have for N large enough

$D_N \supset Y_N$, which is a contradiction.

(ii) l.h.c.: take any $D \subseteq Y$ and a converging sequence $Y_n \rightarrow^H Y$. We need to show that there is a sequence $D_n \rightarrow^H D$ such that $D_n \subseteq Y_n$ for all n . If $D = Y$, we can take the sequence $D_n = Y_n$. Alternatively, suppose $D \subset Y$. Since Y_n converges to Y , for N sufficiently large, we have $D \subset Y_n$ for all $n \geq N$. Hence consider the following sequence: $D_n = Y_n$ for $n < N$ and $D_n = D$ for $n \geq N$. □

This concludes the proof. □

References

- P. Aghion, M. Dewatripont, and P. Rey. Renegotiation design with unverifiable information. *Econometrica*, pages 257–282, 1994.
- C. D. Aliprantis and K. C. Border. *Infinite Dimensional Analysis: a Hitchhiker’s Guide*. Springer, 2006.
- R. Alonso and N. Matouschek. Optimal delegation. *The Review of Economic Studies*, 75(1):259–293, 2008.
- M. Amador and K. Bagwell. The theory of optimal delegation with an application to tariff caps. *Econometrica*, 81(4):1541–1599, 2013.
- M. Armstrong and J. Vickers. A model of delegated project choice. *Econometrica*, 78(1):213–244, 2010.
- S. Auster. Asymmetric awareness and moral hazard. *Games and Economic Behavior*, 82:503–521, 2013.
- S. Auster and N. Pavoni. Limited awareness and financial intermediation. mimeo, 2022.
- D. P. Baron and B. Holmström. The investment banking contract for new issues under asymmetric information: Delegation and the incentive problem. *The Journal of Finance*, 35(5):1115–1138, 1980.
- Y. Feinberg. Games with unawareness. *The BE Journal of Theoretical Economics*, 21(2):433–488, 2021.
- E. Filiz-Ozbay. Incorporating unawareness into contract theory. *Games and Economic Behavior*, 76(1):181–194, 2012.
- A. Francetich and B. C. Schipper. Rationalizable screening and disclosure under unawareness. *University of California, Davis*, 2020.

- S. J. Grossman and O. D. Hart. The costs and benefits of ownership: A theory of vertical and lateral integration. *Journal of Political Economy*, 94(4):691–719, 1986.
- L. Guiso and T. Jappelli. Awareness and stock market participation. *Review of Finance*, 9(4):537–567, 2005.
- Y. Guo and E. Shmaya. Project choice from a verifiable proposal. Technical report, Working paper, 2021.
- M. Halac and P. Yared. Commitment vs. flexibility with costly verification. *Journal of Political Economy*, 128(12), 2020.
- J. Y. Halpern and L. C. Rêgo. Extensive games with possibly unaware players. *Mathematical Social Sciences*, 70:42–58, 2014.
- O. Hart and J. Moore. Incomplete contracts and renegotiation. *Econometrica*, pages 755–785, 1988.
- O. Hart and J. Moore. Agreeing now to agree later: Contracts that rule out but do not rule in. mimeo, 2004.
- A. Heifetz, M. Meier, and B. C. Schipper. Dynamic unawareness and rationalizable behavior. *Games and Economic Behavior*, 81:50–68, 2013.
- A. Heifetz, M. Meier, and B. C. Schipper. Prudent rationalizability in generalized extensive-form games with unawareness. *The BE Journal of Theoretical Economics*, 21(2):525–556, 2021.
- B. Holmström. On the theory of delegation. mimeo, 1980.
- E. Karni and M.-L. Vierø. “reverse bayesianism”: A choice-based theory of growing awareness. *American Economic Review*, 103(7):2790–2810, 2013.
- E. Karni and M.-L. Vierø. Awareness of unawareness: A theory of decision making in the face of ignorance. *Journal of Economic Theory*, 168:301–328, 2017.
- A. N. Kolmogorov and S. V. Fomin. *Introductory real analysis*. Courier Corporation, 1975.
- E. Kováč and T. Mylovanov. Stochastic mechanisms in settings without monetary transfers: The regular case. *Journal of Economic Theory*, 144(4):1373–1395, 2009.
- H. Lei and X. Zhao. Delegation and information disclosure with unforeseen contingencies. *The BE Journal of Theoretical Economics*, 2021.
- D. Martimort and A. Semenov. Continuity in mechanism design without transfers. *Economics Letters*, 93(2):182–189, 2006.

- N. D. Melumad and T. Shibano. Communication in settings with no transfers. *The RAND Journal of Economics*, pages 173–198, 1991.
- R. C. Merton. A simple model of capital market equilibrium with incomplete information. *The Journal of Finance*, XLII(3):483–510, 1987.
- E. Piermont. Introspective unawareness and observable choice. *Games and Economic Behavior*, 106:134–152, 2017.
- A.-K. Roesler and B. Szentes. Buyer-optimal learning and monopoly pricing. *American Economic Review*, 107(7):2072–80, 2017.
- W. P. Rogerson. Economic incentives and the defense procurement process. *Journal of Economic Perspectives*, 8(4):65–90, 1994.
- D. Szalay. The economics of clear advice and extreme options. *The Review of Economic Studies*, 72(4):1173–1198, 2005.
- J. Tirole. Cognition and incomplete contracts. *American Economic Review*, 99(1):265–94, 2009.
- E.-L. Von Thadden and X. Zhao. Incentives for unaware agents. *The Review of Economic Studies*, 79(3):1151–1174, 2012.
- X. Zhao. Framing contingencies in contracts. *Mathematical Social Sciences*, 61(1):31–40, 2011.