Discussion Paper Series – CRC TR 224

Discussion Paper No. 256
Project B 02

Optimal Delegation and Information Transmission Under Limited Awareness

Sarah Auster [1]
Nicola Pavoni [2]

January 2021

[1] Department of Economics, University of Bonn: auster.sarah@gmail.com
[2] Department of Economics, Bocconi University: pavoni.nicola@gmail.com

Funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 is gratefully acknowledged.

Optimal Delegation and Information Transmission under Limited Awareness*

Sarah Auster‡ and Nicola Pavoni§

January 2021

Abstract

We study the delegation problem between a principal and an agent, who not only has better information about the performance of the available actions but also has superior awareness of the set of actions that are actually feasible. The agent decides which of the available actions to reveal and which ones to hide. We 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 thereby increases the principal's willingness to grant him higher information rents. We also consider communication between the principal and the agent after the contract is signed and the agent receives information. We show that limited awareness of actions improves communication in such signalling games: the principal makes a coarser inference from the recommendations of the privately informed agent and accepts a larger number of his proposals.

JEL Codes: D82, D83, D86

*The authors thank Sandeep Baliga, Pierpaolo Battigalli, Sylvain Chassang, Wouter Dessein, Antonio Guarino, Marina Halac, Johannes Hörner, Navin Kartik, Alessandro Pavan, and Yuval Salant for very helpful comments. The present draft benefited from several comments received by the seminar participants at the Toulouse School of Economics, Collegio Carlo Alberto, University College London, EIEF in Rome, Bonn University and at the Max Plank Institute, at the University of Bristol, the University of Cardiff, at the Bocconi and Cattolica universities in Milan, at the University of Verona, University of Malaga, Oxford University, Stockholm School of Economics and the Cowles Conference on Economic Theory. Nicola Pavoni acknowledges financial support from the MIUR-PRIN grant 20157NH5TP. 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. Headquarters rely on division managers who have superior information about the profitability of available projects but also a desire to attract additional resources to their own division, voters entrust decisions to politicians whose preferences may reflect a political bias or the interest of certain lobbies, financial investors seek advice from non-neutral financial professionals with a better understanding of the risks and returns of the available portfolios. The tension underlying these situations has been formalised in the delegation model—first introduced by Holmström (1977)—where an uninformed principal specifies a set of permissible actions to the informed agent and contingent transfers are infeasible.

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

[1]For a recent application of the classical delegation model in political competition see Kartik et al. (2017). Somin (2016) and Carpini and Keeter (1996) present the results of a number of surveys on US voters over various decades and document the lack of knowledge of basic institutional rules and of the set of policies available to local governments. For example, Somin (2016) documents that 34% of US voters cannot name the three branches of the federal government, a similar percentage do not know which government officials are responsible for which issues; Carpini and Keeter (1996) document that less than 50% of US voters know whether the local governors have to approve the decisions made by their higher state court; 25% do not know whether states can pass a law prohibiting abortion.

[2]Following Baron and Holmström (1980), the standard delegation model has been adopted by several authors to study financial advice, investment banking and delegated portfolio management. The presence of partial awareness of financial products by investors has been recognised since at least Merton (1987). More recently, Guiso and Jappelli (2005) document the lack of awareness of financial assets among the 1995 and 1998 waves of the survey of Italian households (SHIW). Only 65% of potential investors were aware of stocks and only 30% of investment accounts; mutual funds and corporate bonds were known by only 50% of the sample. The share of wealth in the hand of unaware agents was also substantial. The share of wealth owned

This paper studies the implications of such asymmetry by incorporating unawareness into the canonical delegation model. We consider the problem of a principal (she) who needs to take an action and delegates the task to an agent (he). The agent receives private information about the payoffs of each of the available actions and the principal's problem is to determine a set of actions from which the agent can choose (see for example Alonso and Matouschek, 2008). We depart from the traditional framework of optimal delegation by considering a situation where the principal is *unaware of some feasible actions*. Our key assumption is that the principal's unawareness limits her language to 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 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 rents. We address this question in an environment with a continuum of payoff states, a continuum of feasible actions and an agent who in each state prefers a higher action than the principal. Given her awareness, the principal's optimal delegation set solves the usual tradeoff between minimising 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 cap above which no action is permitted. How high this cap is depends on the principal's awareness set. We show under minimal restrictions that the agent optimally leaves the principal unaware of an interval of actions around the optimal upper cap 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. We derive the agent's optimal disclosure policy and the resulting delegation set explicitly for the case of quadratic utility functions and a uniform bias.

by households that were not aware of corporate bonds was approximately 20%, and so was the share owned by those unaware of mutual funds.

The baseline model assumes that the disclosure of feasible actions takes place before the agent receives private information. We relax this assumption in the second part of the paper and allow the agent to propose additional actions after the contract is signed and the agent observes the state. Principal and agent thus play a signalling game in the last stage. The study of this problem reveals interesting implications of asymmetric awareness on strategic information acquisition. We characterise the outcome of the agent's best equilibrium and show that the agent can fill all potential gaps of the delegation set below its upper cap. From a modeller's viewpoint 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. Due to this asymmetry, each of the agent's equilibrium proposals is perceived to be consistent with an interval of states and these intervals overlap. Thus, in contrast to the case of full awareness, the principal's information cannot be represented by a partition of the state space into pairwise disjoint sets. Translated to the canonical problem of cheap talk (Crawford and Sobel, 1982), the result shows that asymmetric awareness allows for substantially finer communication and improved outcomes for both parties compared to the case of full awareness.

We discuss some extensions of the baseline model (without renegotiation). In particular, we consider the case where the set of feasible actions is an arbitrary subset of the reals, where actions are multidimensional, and where the agent does not know the principal's initial awareness. Finally, we discuss the role of the principal's sophistication.

After the literature review, the paper is organised as follows. Section 2 presents the delegation model with limited awareness. In Section 3 we analyse the agent's optimal disclosure and the resulting delegation set. Section 4 analyses the game with renegotiation after the agent receives private information. Section 5 discusses extensions and Section 6 concludes.

1.1 Related Literature

The paper makes both applied and theoretical contributions. It introduces unawareness to the canonical delegation problem and shows that a biased agent has incentives to hide moderate options from the principal in order to implement more extreme ones. The identified distortion may have significant effects on economic outcomes in a range of situations, as argued above. The paper thus relates to the literature on optimal delegation and on applications of unawareness to games and contracting problems.

Holmström (1977) 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.

There are only few papers that apply 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 unaware (Von Thadden and Zhao 2012, 2014; Zhao 2011; Filiz-Ozbay 2012; Auster 2013). One exception is Francetich and Schipper (2020) who consider 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. Lei and Zhao (2020) consider a particular case of our model (quadratic utility, uniform bias, uniform distribution) to study unawareness of contingencies (nature's moves) rather than players' actions.[3]

On the theoretical side, the study of the disclosure problem reveals how the agent's 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, e.g. Roesler and Szentes (2017), Garrett et al. (2020).

[3]Conceptually, there is also a connection to the literature on incomplete contracts and unforeseen contingencies (Grossman and Hart, 1986). In contrast to that literature, our model entails that the principal is limited in her choice of contracts but correctly foresees the outcome of any contract she can specify.

Finally, the paper introduces a new class of communication games with non-verifiable information, where the receiver's awareness of the possible signals depends on the realised state. We show that the discrepancy between the sender's signalling strategy and the receiver's perception of it can substantially change the outcomes in such games. The equilibrium we describe also provides a neat illustration for how unawareness differs from the standard model and, in particular, 'zero probability' beliefs. We hence contribute to the literature on signalling and strategic information transmission (Crawford and Sobel, 1982). To the best of our knowledge, we are the first to study the equilibrium implications of limited awareness on strategic information transmission in a setting where information is not verifiable. Heifetz et al. (2020) study the strategic disclosure of hard information and find that if the information is multidimensional and the receiver is unaware of some dimension, unraveling is not a necessary outcome of the game.

In a companion paper, Auster and Pavoni (2020), we apply our delegation model (without signalling) to financial intermediation, considering a market with multiple fully aware brokers (agents) and a continuum of partially aware investors (principals). Self-reported data from customers in the Italian retail investment sector support the key predictions of the model: the menus offered to less knowledgable investors contain fewer products, which are perceived to be more extreme.

2 Environment

There is a principal and an agent. The agent has access to an interval of actions $Y^A = [y_{min}, y_{max}]$.[4] 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 expected

[4] We discuss in Section 5 the extension to general subsets of $\mathbb{R}$, for instance, finite collections of points, as well as the case of multi-dimensional actions.

utility functions with *continuous* Bernoulli components given by[5]

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

Fixing θ, $U^i, i = P, A$ 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 $U^i_{y\theta} > 0$, 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, $y^P(\theta) < y^A(\theta)$.[6]

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.[7] 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.

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.[8] Our

[5]Note 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.

[6]To ease exposition we assume that the bias is always positive. However, it would be easy to extend the analysis where the bias is negative and, slightly adjusting our assumptions, to the case where the bias changes the sign.

[7]See Karni and Vierø (2013, 2017) for a decision theoretic model capturing this type of unawareness.

[8]The 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,

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 is, the richer is 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 which the principal does not anticipate. We will discuss this case in Remark 1 in the following section and concentrate now on the case where the principal specifies the actions which she permits. Since the principal cannot specify actions of which she is unaware, the principal's delegation set is then a subset of her awareness set. The timing of the game can be summarised as follows:

1. The principal's initial awareness Y^P is realised and observed by all parties.
2. The agent reveals a closed set of actions $X \subseteq Y^A$ and the principal updates her awareness to $Y = Y^P \cup X$.
3. Given awareness set Y, the principal chooses a delegation set $D \subseteq Y$.
4. The agent observes θ and chooses an action from set D.
5. Payoffs are realised.

Notice that we have not made any explicit assumption on whether or not the principal is aware of her unawareness. The principal might take the world at face value or she might understand that there exist actions outside her awareness. Since she cannot include such actions in the delegation set, awareness of their possible existence neither affects her expected payoff nor optimisation problem.[9] Furthermore, within the constraints of her awareness, the

the principal commits to a mechanism that specifies an action as a function of the agent's message. Their argument continues to hold in our setting. There are thus two possible interpretations: after observing the payoff parameter, the agent might directly choose an action or he might make a recommendation which the principal has committed to follow.

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

principal is perfectly rational: she anticipates correctly the expected payoff associated to each feasible delegation set and will not be surprised ex-post.

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 modeller'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. In Appendix B.1, we provide a more extensive description of the family of game trees representing the generalised game with unawareness associated to our delegation model according to the approach proposed by Heifetz et al. (2013). Figure 3 in the Appendix reports a graphical example.

As a solution concept, we use a strong version of Perfect Bayesian Nash Equilibrium (PBE) which implies subgame perfection, adapted to generalised extensive-form games with unawareness (e.g., see Halpern and Rêgo (2014) and Feinberg (2020)). In Appendix B.2, 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 generalised game. Rationalizability assumes that players have common knowledge of their rationality and their preferences but it does not assume, for example, that a player is automatically certain of a ready-made convention of play upon becoming aware.[10] Despite the clear appeal of this notion for games with unawareness, our focus on Bayesian Nash equilibrium as a solution concept in the main body of the paper facilitates considerably the comparison to existing results in the literature, especially in Section 4, where signalling arises as part of the game.

3 Equilibrium Analysis

We will work backwards 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

[10] See Guarino (2020) and Heifetz et al. (2020).

stage of the game is defined by

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

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 favour of the principal).[11]

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^*:

$$V^P(D) := \int_0^1 U^P(\theta, y^*(\theta, D))\mathrm{d}F(\theta). \tag{2}$$

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 $y^*(\theta, D) = y$.[12] 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

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

Again, if problem (3) has multiple solutions, we assume that the principal choses the agent-preferred set and denote by $D^*(\cdot)$ such selection from the set of maximisers for each Y. 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 that the bias is not too large and implies that the principal prefers the delegation set $[y^A(0), y_0^*]$ to the singleton $\{y_0^*\}$ (see also Alonso and Matouschek (2008), Corollary 2).

[11] Such selection is well defined as it is easy to show - from the joint continuity of U^A - that the BR^A correspondence is upper hemicontinuous.

[12] This restriction reduces multiplicities by eliminating 'redundant' delegations sets that contain actions that will not be chosen under any contingency.

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 which the principal already knows, the induced awareness set must contain the principal's initial awareness set Y^P. The smaller Y^P is, the larger is the collection of awareness sets from which the agent can choose. An optimal awareness set Y^* solves the problem

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

Since different awareness sets might induce the same delegation set, the solution to problem (4) is 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.

Equilibrium disclosure. The central question of this paper is whether the agent distorts the principal's delegation choice in his favour 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 realisations of θ to take an action closer to the principal's conditionally preferred action.

Let $\hat{y} := \max D^*(Y^A) < y^A(1)$ denote the *upper cap* which the principal imposes under the optimal delegation set in the full awareness benchmark. The following proposition provides conditions under which unawareness of $\hat{y}$ is sufficient to ensure that the agent benefits from the principal's limited awareness.

Proposition 1 (Optimality of Limited Awareness)**.** *Assume that problem (3) has a unique maximiser $D^*(Y^A)$ and that the upper cap $\hat{y}$ is not an isolated point of $D^*(Y^A)$. If $\hat{y} \notin Y^P$, then generically the agent strictly prefers not to disclose all actions in Y^A.*

Proposition 1 shows that 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 cap $\hat{y}$. 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. Hence, by leaving the principal unaware of actions around $\hat{y}$, the agent can implement an action $y^+ > \hat{y}$ that is not permitted under full awareness.

The agent's gain in flexibility 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. To see this, define $s(\cdot)$ as the inverse of $y^A(\cdot)$ and consider the state $s(\hat{y})$, where the agent's preferred action is $\hat{y}$. By the assumption that $\hat{y}$ is a limit point of $D^*(Y^A)$, there is an interval of actions to the left of $\hat{y}$ that are permitted under $D^*(Y^A)$. This implies that there is an interval of states to the left of $s(\hat{y})$ such that for all states belonging to the interval, the agent gets to take his preferred action under $D^*(Y^A)$. The perturbation forces the agent to move away from his bliss point in these states. 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 to the right of $s(\hat{y})$.

We should emphasise that when the optimal delegation set under full awareness is not an interval, the agent profits from perturbations around other pooling points as well. For example, if $D^*(Y^A)$ has an intermediate gap $(\underline{y}, \overline{y})$, the agent benefits from moving up the lower bound $\underline{y}$ at the cost of losing some flexibility below $\underline{y}$.[13] The main complication here is that such perturbation may affect the principal's optimal choice of $\overline{y}$. If the optimal value for $\overline{y}$ 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 which need to be compared. To guarantee the

[13]Similarly, the agent benefits from lowering $\overline{y}$ at the cost of losing some actions above $\overline{y}$.

profitability of the perturbation in this case, more stringent assumptions on the principal's initial awareness set are 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 $\overline{y}$.

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).

The same principle applies to other screening problems. When deciding on the contract, a principal facing a privately informed agent solves a tradeoff between eliciting the agent's information to take a suitable action and limiting the agent's information rent. *By making specific actions unavailable, the agent can increase the cost for the principal not to use the agent's information and thereby increase her willingness to grant the agent higher information rents.*

Interval delegation. The literature on optimal delegation 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 $Conv(D)$. Then, for all $A \subseteq Conv(D)$:*

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

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 $(\underline{y}, \overline{y})$. The removal of actions in $(\underline{y}, \overline{y})$ means that there is an interval of states where the agent switches to the lower action $\underline{y}$ and an interval of states where the agent switches to the higher action $\overline{y}$ with respect to the original delegation set. Since the principal is downward biased with respect to the agent, the switch to the lower action benefits 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 favours 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 $y < y^A(1)$.[14] The optimal cap is the value of y that maximises $W : Y \to \mathbb{R}$, where

$$W(y) := \int_0^{s(y)} U^P(\theta, y^A(\theta))\,\mathrm{d}F(\theta) + \int_{s(y)}^1 U^P(\theta, y)\,\mathrm{d}F(\theta). \tag{5}$$

We now show that unawareness of the optimal cap, which we denote again by $\hat{y}$, is not only a sufficient condition for less-than-full revelation to be strictly optimal but also a necessary one. We further show that the resulting delegation set has a single gap around $\hat{y}$.

[14] Recall, $y^A(\cdot)$ is the unrestricted choice function for the agent.

Proposition 2 (Optimal Delegation under Limited Awareness)**.** *Let Assumptions 1 be satisfied. (i) The agent optimally reveals all feasible actions to the principal if and only if the principal is aware of the action $\hat{y}$. (ii) If in addition W in (5) is single peaked,*[15] *there exist two parameters $\Delta_1, \Delta_2 \geq 0$ such that:*

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

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 maximises his discretion by revealing all actions below $\hat{y}$. Hence, the agent 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 $\hat{y}$. The resulting delegation set includes all relevant actions below the interval and one action above it.

Remark: *Our model captures situations where the principal is unaware of actions, such as economic policies, financial and non-financial products, procurement tasks, etc. We assume these actions are one-dimensional. In contrast, Heifetz et al. (2020) consider a situation where actions are multidimensional and a receiver is unaware of certain dimensions (interpreted as attributes of the action). We could consider this type of unawareness in the context of our problem. If the language of the contract is such that the agent is free to choose on dimensions which are left unspecified, it is clear that the agent would have no incentives to reveal additional attributes of actions to the principal. Since disclosing new attributes can only reduce the agent's flexibility, disclosure cannot be profitable. The same is true in our one-dimensional setting if instead of specifying the actions that are permitted, the principal would specify those actions that are not permitted. Hence, in contrast to the benchmark case of full awareness where both types of specifications lead to the same outcomes, the language in which a contract is written matters here.*[16]

[15] Over the unidimensional space, single-peakedness is equivalent to strict quasi-concavity.

[16] One could generalise our model allowing for a more complex relationship between principal awareness

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

$$U^P(y,\theta) = -(y-(\theta-\beta))^2, \quad U^A = -(y-\theta)^2. \tag{6}$$

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 guarantying interval delegation to be optimal is the following regularity condition on the distribution function (see Martimort and Semenov, 2006):

$$f'(\theta)\beta + f(\theta) > 0 \text{ for all } \theta \in (0,1). \tag{7}$$

Delegation is valuable for the principal if $\mathbb{E}[\theta - \beta] > 0$. When these two 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, and Alonso and Matouschek, 2008):[17]

$$\hat{y} = \mathbb{E}[\theta - \beta | \theta \geq \hat{y}]. \tag{8}$$

The agent chooses his preferred action $y = \theta$ for all $\theta < \hat{y}$ and the action $\hat{y}$ in all remaining states.

Delegation choice. We start by deriving the optimal delegation set for an arbitrary awareness set $Y \in \mathcal{Y}$. As the following proposition shows, this delegation set has a close relation to the optimal delegation set under full awareness.

Proposition 3 (Generalised Interval Delegation)**.** *Let $Y \in \mathcal{Y}$ and define $\hat{y}_Y := \arg\min_{y \in Y} |y - \hat{y}|$. If condition (7) is satisfied, the optimal delegation set with respect to awareness set Y is*

$$D^*(Y) = \{y \geq 0 : y \in Y \text{ and } y \leq \hat{y}_Y\}.$$

and her available actions. The key trade off faced by the agent in our model will be preserved as long as a more aware principal has both a non-trivial possibility of enlarging the set of actions permitted to the agent and the possibility of excluding some actions.

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

Since the agent is upward biased and Assumption 1 is satisfied, the optimal delegation set for awareness set Y is fully described by its upper bound. Proposition 3 shows that the optimal upper bound is given by the element of Y that is closest to $\hat{y}$. This property is a consequence of the symmetry of the principal's payoff function. Note that $\hat{y}$ is the best action for the principal if such action must be constant on $\theta \in [\hat{y}, 1]$. Now, consider $\bar{y} > \hat{y}$ and $\underline{y} < \hat{y}$. If the delegation set includes $\underline{y}$ and $\bar{y}$ and no actions in between, the agent will take $\bar{y}$ in all states $\theta \geq \frac{\underline{y}+\bar{y}}{2}$. If $\bar{y}$ is removed, the agent switches from $\bar{y}$ to $\underline{y}$ in all those states. Now, suppose $|\bar{y} - \hat{y}| > |\underline{y} - \hat{y}|$. Then, by symmetry of the payoff, the principal strictly prefers action $\underline{y}$ over $\bar{y}$ for $\theta \in [\hat{y}, 1]$, as $\underline{y}$ is closer to the principal's ideal action $\hat{y}$. For $|\bar{y} - \hat{y}| - |\underline{y} - \hat{y}|$ small, we have $\frac{\underline{y}+\bar{y}}{2} \approx \hat{y}$, hence the principal is better off by not permitting $\bar{y}$. The same applies, a fortiori, for larger values of $|\bar{y} - \hat{y}| - |\underline{y} - \hat{y}|$, because a unit increase in $\bar{y}$ increases the range of θ for which $\bar{y}$ is chosen by the agent only by one half.

The property $\max D^*(Y) = \hat{y}_Y$ has two important implications: first, the optimal delegation set includes all actions belonging to Y that are weakly smaller than $\hat{y}$; second, it includes at most one action strictly greater than $\hat{y}$. The optimal delegation set under partial awareness can thus be seen as the closest approximation of the optimal interval under full awareness, $[0, \hat{y}]$, which is available to the principal given her restricted awareness. This approximation includes an element $y > \hat{y}$ if and only if y is closer to $\hat{y}$ than any element of Y smaller than $\hat{y}$. For a graphical illustration see Figure 1.

Optimal Disclosure. With the characterisation of D^*, we can turn our attention to the agent's optimal strategy of expanding the principal's awareness. From Proposition 2 we know that if the principal is aware of the threshold action $\hat{y}$, the agent optimally reveals all other actions. Since there is no action closer to $\hat{y}$ than $\hat{y}$ itself, the upper bound of the optimal delegation set will always be $\hat{y}$. Disclosing actions above $\hat{y}$ is thereby irrelevant; the principal will never allow the agent to implement any of them. On the other hand, revealing actions below the threshold $\hat{y}$ is strictly optimal since they will be included in the optimal delegation set, therefore expanding the agent's choice.

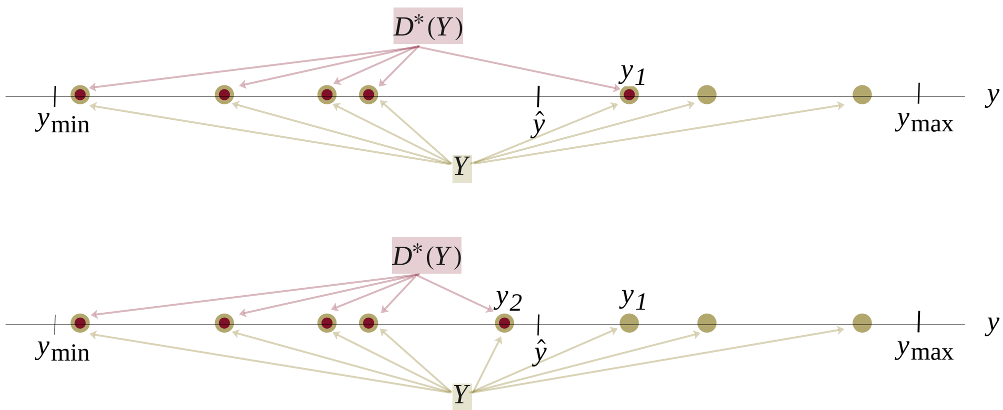

Figure 1: Optimal delegation set $D^*(Y)$. The figures represent two examples of the principal's awareness set Y. In both figures, the yellow bullets represent the set Y while the red bullets represent the resulting optimal delegation set $D^*(Y)$. In the upper figure, the principal includes action y_1 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 $D^*(Y)$. Differently put, the awareness of action y_2 by the principal 'crowds out' action y_1 from the resulting delegation set.

Starting now from an arbitrary set Y^P, the above argument implies that the optimal awareness set Y^* is such that the upper bound of the corresponding delegation set $D^*(Y^*)$ is at least $\hat{y}$. Moreover, the only reason for the agent to leave the principal unaware of certain actions is to induce the principal to permit some action strictly greater than $\hat{y}$. By Proposition 3 this is optimal for the principal if and only if the principal is not aware of any action closer to $\hat{y}$. Hence, the gap around $\hat{y}$ is symmetric, so the optimal awareness gap characterised in Proposition 2 is determined by a single parameter $\Delta_1 = \Delta_2 = \Delta$. The corresponding delegation set is then given by:

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

Given such delegation set, the agent's optimal policy is as follows. In states below $\hat{y} - \Delta$ the agent uses his flexibility and implements his preferred action $y = \theta$. In states above $\hat{y} - \Delta$ the preferred action is not available, so the agent chooses the one closest to his bliss point. For states in the interval $(\hat{y} - \Delta, \hat{y}]$ this is action $\{\hat{y} - \Delta\}$, for the remaining states it is $\{\hat{y} + \Delta\}$.

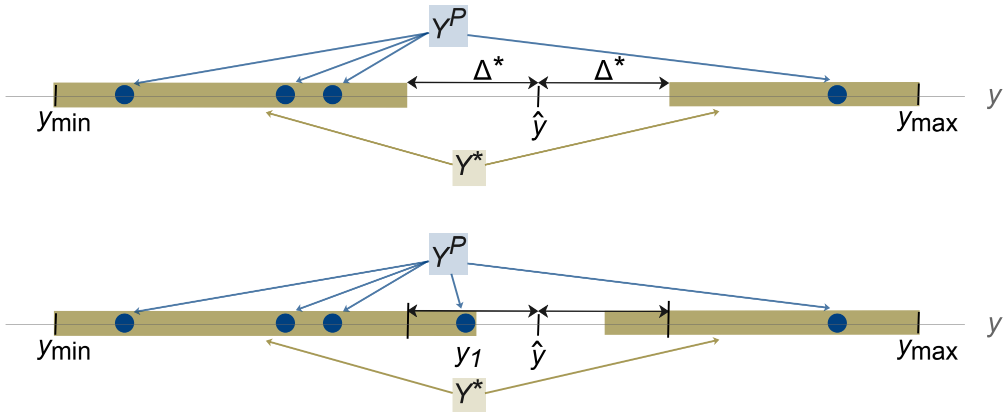

Figure 2: Optimal awareness set Y^*. The figures represent two examples of the principal's initial awareness set Y^P and associated awareness sets $Y^* = Y^P \cup X^*$ after including disclosed actions X^*. In both figures, the blue bullets represent the set Y^P, while the yellow set represents the resulting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

Using this policy, we can write the agent's expected payoff as a function of Δ. The agent chooses the value of Δ that maximises his expected payoff. This choice is restricted by the principal's initial awareness set. Letting $\bar{\Delta}(Y^P) := \arg\min_{y \in Y^P} |y - \hat{y}|$ indicate the maximum feasible awareness gap, the agent's optimisation problem amounts to:

$$\max_{\Delta \geq 0} \quad -\int_{\hat{y}-\Delta}^{\hat{y}} (\hat{y} - \Delta - \theta)^2 dF(\theta) - \int_{\hat{y}}^{1} (\hat{y} + \Delta - \theta)^2 dF(\theta) \quad \text{s.t.} \quad \Delta \leq \bar{\Delta}(Y^P). \tag{9}$$

The solution to problem (9) is given by $\min\{\bar{\Delta}(Y^P), \Delta^*\}$, where Δ^* solves the first-order condition

$$\int_{\hat{y}-\Delta^*}^{\hat{y}} [\theta - (y - \Delta^*)] \, dF(\theta) = \int_{\hat{y}}^{1} [\theta - (y + \Delta^*)] \, dF(\theta). \tag{10}$$

A gap parametrised by Δ^* is implemented whenever the principal's initial awareness does not constrain the agent in his choice. 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 shown in Figure 2.

Proposition 4 (Comparative Statics)**.** *Let* $\Delta^*(\beta)$ *be the unconstrained solution to problem* (9) *when the principal's preferences parameter is* $\beta \in (0, \mathbb{E}[\theta])$ *and condition (7) is satisfied. Then* $\Delta^*(\cdot)$ *is an increasing function.*

Proposition 4 shows an intuitive result: the larger is 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.[18] 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 (8). Considering the agent's tradeoff when choosing Δ, notice that when F is uniform, the cost associated to the loss of flexibility for a given gap Δ 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 $\hat{y}(\beta)$ is, the larger is the range of values for θ above $\hat{y}(\beta) + \Delta$ and hence the set of types to whom this rent accrues.

4 Renegotiation

Thus far, we have assumed that the agent can only reveal actions to the principal before he learns the payoff state θ. This entails that even in states where the agent knows of an action that makes both parties better off, additional communication is not possible. An interesting question is how the outcome changes if after learning the payoff parameter θ, the agent can reveal additional actions to the principal, who then decides whether to permit a new action or to maintain the original contract. We thus consider a model where the agent can renegotiate with the principal. If the agent proposes to replace the original delegation set with a new action, the principal understands that the agent's choice signals something about the state. In particular, the principal can infer that the agent only reveals an action if that benefits him. However—due to the principal's limited awareness—she cannot conceive of alternative

[18]It turns out that our argument holds true for all distributions satisfying condition (7).

actions the agent could have disclosed. This implies that the principal cannot learn from particular actions not being proposed, a key difference to the case of full awareness.

Modified game. The possibility of renegotiation considerably changes the nature of the game. While the probability distribution over θ is common knowledge, renegotiation occurs under asymmetric information, implying that the principal and the agent play a signalling game. We will distinguish two phases of the game: the contracting phase and the renegotiation phase. To simplify the analysis, we assume that the principal is initially unaware of the possibility of renegotiation. The contracting phase is then the same as before: the agent reveals a set of actions to the principal and the principal chooses the optimal delegation set with respect to her updated awareness set. Note, however, that the equilibrium we derive continues to exist when this restriction is removed. According to the principal's perception, the allocation of any pure-strategy equilibrium involving renegotiation can be replicated by directly adding the actions that are permitted on path to the delegation set. Since the principal cannot foresee the possibility of renegotiating over actions of which she is unaware in the contracting phase, she sees no value in renegotiating the optimal contract.

In the renegotiation phase, the agent learns the realised value of θ and can propose additional actions to the principal. The principal updates her awareness and her beliefs about the payoff state. We will restrict attention to a game where the agent proposes single actions to the principal in the renegotiation phase.[19] The principal's decision is then between maintaining the original delegation set and permitting the proposed action. A detailed description of the strategies of the two players and the definition of Perfect Bayesian Equilibrium (generalised to allow for differential awareness) is reported in Appendix A.5. In Appendix B.3 we describe the generalised extensive-form game à la Heifetz et al. (2013).

Disclosure in the renegotiation phase. As usual, we move backwards and start with the analysis of the renegotiation stage. The following proposition provides a key property of

[19] It can be easily shown that the agent cannot improve on the agent-best equilibrium we describe in Proposition 6 by revealing more than one action at a time.

our model: in any equilibrium, the principal permits an agent's proposal only if she whould have included the action in the delegation set at the contracting phase.

Proposition 5 (Ex-ante like Reasoning)**.** *Let Y denote the principal's awareness set at the beginning of the renegotiation phase. (i) Fix an equilibrium, and let A be the (measurable) set of proposals which the principal accepts in the renegotiation phase. Then, for all $x \in A$,*

$$V^P(D^*(Y) \cup \{x\}) \geq V^P(D^*(Y)). \tag{11}$$

(ii) Conversely, for any set A constituted of points satisfying condition (11)*, there is an equilibrium such that A is the set of actions accepted in the renegotiation phase.*

Recall that V^P is defined as the expected utility of the principal under 'pure' delegation, that is, the case where the principal is not expecting to have any renegotiation (that is why the payoff only depends on the delegation set). Proposition 5 characterises the set of 'implementable' proposals: for each possible awareness level Y, it defines the set of proposals that can be accepted by the principal in equilibrium. Condition (11) vacuously holds for $x \in D^*(Y)$. Accepting the proposal or keeping the original delegation set results in the same action. On the other hand, for actions not belonging to $D^*(Y)$, the inequality can be satisfied only if x does not belong to Y. Hence, the agent can only gain from renegotiation if he discloses new actions with respect to the contracting phase.

To prove Proposition 5, we first show that there is no equilibrium in which the principal permits an action in Y which does not belong to $D^*(Y)$. Building on this result, we argue that in an equilibrium where proposal $x \notin Y$ is accepted in the renegotiation phase, the principal's beliefs after proposal x are described by the set

$$\{\theta \in [0,1] : U^A(\theta, x) \geq \max_{y \in D^*(Y)} U^A(\theta, y)\}. \tag{12}$$

According to the principal's awareness, x is the only new action which 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 (11) is satisfied. Indeed, 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.

Disclosure in the contracting phase. The disclosure of actions in the contracting phase determines a delegation set, which in turn determines a set of actions that can be implemented through further revelation in the renegotiation phase. For each initial disclosure, the set of actions that are ultimately implementable is characterised by (11). Due to the signalling nature of the game, there are typically multiple equilibria. For instance, there is always a trivial equilibrium, where the agent believes that no proposal will be accepted and hence does not renegotiate. We are interested in the question of *what is possible* to achieve through strategic disclosure of actions and focus on the outcome of the best equilibrium for the agent.

Letting $Z : \mathcal{Y} \to \mathcal{Y}$ describe the mapping from awareness sets to the *maximal set* of implementable actions, the agent solves the problem

$$\max_{Y \supseteq Y^P} \int_0^1 \max_{y \in Z(Y)} U^A(\theta, y) dF(\theta),$$

$$\text{where} \qquad Z(Y) := \{z \in Y^A : V^P(D^*(Y) \cup \{z\}) \geq V^P(D^*(Y))\}.$$

Without further restrictions, the analysis of this problem can be intricate. For an illustration, consider the following example with three actions $y_1 < y_2 < y_3$. Suppose the principal's initial awareness set is $Y^P = \{y_2\}$ and she has the following preferences:

$$V^P(\{y_1, y_3\}) > V^P(\{y_1\}) > V^P(\{y_1, y_2\}) > V^P(\{y_1, y_2, y_3\}).$$

Conditional on action y_2 being in the delegation set, the principal does not want to include action y_3. In order for the agent to take action y_3, he must make the principal aware of y_1

in the contracting stage. Action y_1 crowds out y_2 and thereby opens up the possibility for the principal to permit action y_3. The agent could also reveal y_1 and y_3 in the contracting phase so that the principal delegates the set $\{y_1, y_3\}$. To see why it can be better to hold back with action y_3, suppose there is a fourth action y_4 such that

$$V^P(\{y_1, y_3\}) > V^P(\{y_1, y_4\}), V^P(\{y_1, y_3, y_4\}) > V^P(\{y_1\}).$$

Revealing y_3 in the contracting phase implies that the agent will not be allowed to take action y_4. If instead the agent initially discloses only y_1, he can implement both y_3 and y_4, depending on which action he prefers after observing the realisation of θ.

In order to simplify the analysis, we impose some more structure on the principal's preference over delegation sets and require Assumption 1 to be satisfied. To recall, Assumption 1 assures that the principal always benefits from closing gaps in the delegation set. In order to state the result, we define the highest action that the principal is willing to delegate in the contracting phase for some awareness set that is consistent with the principal's initial awareness:

$$y^{max} := \max\{y \in Y^A : y \in D^*(Y) \text{ for some } Y \supseteq Y^P\}$$

For concreteness, consider the example of quadratic utility functions and a uniform bias, as specified in Section 3.1. Under this specification, the highest implementable action is given by $\hat{y} + \bar{\Delta}(Y^P)$, where $\bar{\Delta}(Y^P)$ was defined as the minimal distance between an action belonging to Y^P and the upper bound of the optimal delegation set under full awareness, $\hat{y}$.

The last-stage outcome of a pure-strategy equilibrium is described by a function $y^E : \Theta \to Y^A$, which maps each state θ to the implemented action.

Proposition 6 (Agent Best Equilibrium)**.** *Let Assumption 1 be satisfied. In the agent's best equilibrium of the renegotiation game, the outcome is described by*

$$y^E(\theta) = \begin{cases} y^A(\theta) & \text{if} \quad \theta \leq s(y^{max}) \\ y^{max} & \text{if} \quad \theta > s(y^{max}). \end{cases}$$

Recall that $s(y) = (y^A)^{-1}(y)$ is the state in which y is the agent's preferred action. Proposition 6 shows that when the agent can renegotiate after observing the realisation of θ, he is able to implement all actions below y^{max}. Hence, by disclosing actions in sequence the agent cannot only increase the highest permitted action (as in the case without renegotiation), but he can do it without any loss of flexibility below the cap.

In Appendix A.7, we provide an example of equilibrium strategies and beliefs delivering the equilibrium outcome described by y^E. For a simple illustration, let us consider the case where the principal's optimisation problem under full awareness has a unique local maximum (W in (5) is single peaked). Under this assumption, we can restrict attention to initial disclosures that induce an awareness set with a single gap.[20] Suppose that in the contracting phase, the agent discloses a set $Y_{\bar{\Delta}}$ to generate the largest possible gap in the delegation set:

$$D^*\left(Y_{\bar{\Delta}}\right) = [y^A(0), \hat{y} - \bar{\Delta}] \cup \{y^{max}\}.$$

Moving to the renegotiation phase, we distinguish three different regions. If $\theta \leq s(\hat{y} - \Delta)$, there is no need for the agent to renegotiate, since the agent's preferred action already belongs to the delegation set. In states belonging the interval $(s(\hat{y} - \Delta), s(y^{max}))$, the agent proposes his preferred action $y^A(\theta)$. By Assumption 1, the principal would have preferred to add this action to the initial delegation set. As we have shown in Proposition 11, this implies that the principal is willing to accept $y^A(\theta)$ in the renegotiation phase. Hence, the principal permits all actions in the interval $(\hat{y} - \bar{\Delta}, y^{max})$. Finally, if $\theta > s(y^{max})$, the agent would like to take an action that is higher than y^{max}. However, by definition of y^{max},

$$V^P(D^*(Y_{\bar{\Delta}})) > V^P(D^*(Y_{\bar{\Delta}}) \cup \{x\})$$

holds for all $x > y^{max}$, which, by Proposition 5, implies that the principal rejects all such proposals. The best alternative for the agent is thus to take y^{max}. The argument further implies that there is no equilibrium of the renegotiation game where the principal permits a higher action than y^{max}. Hence, the outcome described in Proposition 6 maximises the

[20] See the proof of Proposition 2.

agent's expected payoff. Notice also that y^{max} is equal to $\hat{y}$ (the optimal cap under full awareness) if and only if $\hat{y} \in Y^P$.

Dynamic awareness. The model with renegotiation highlights two important aspects of games with limited awareness. The first concerns the dynamics of unawareness. Much like information, 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 is 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 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 renegotiation is possible than when it is not. In the case of quadratic utility functions, the optimal value of Δ parametrising the delegation set in the contracting phase is maximal when renegotiation is possible. Notice that even without renegotiation, the agent could implement *any single action below* y^{max} by revealing the 'right' set of actions. However, he cannot not implement *all actions below* y^{max} 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 realisation of θ.

Information transmission. In the renegotiation stage, the principal and the agent play a signalling game. Unless $y^{max} = \hat{y}$, the delegation set from the contracting phase has a gap below y^{max}. One striking feature of the equilibrium described in Proposition 6 is that the agent's implemented action is strictly increasing in θ for all θ such that $y^A(\theta) \leq y^{max}$. In other words, there is no pooling of types below y^{max}. This would not be possible under full awareness: in any candidate equilibrium where types 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, the principal cannot contemplate moves of the agent of which she is unaware and this limits the extent to which she infers information from the agent's recommendation. In particular, if the realised value is θ and the agent proposes $y^A(\theta) \notin Y$, the subjective game tree that represents the principal's frame of mind after updating does not include moves of the agent involving an action 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 the principal's subjective game, 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. Conditional on that information, the principal indeed prefers action $y^A(\theta)$ to the initial delegation set.

4.1 Communication without commitment.

A general takeaway from this analysis is the insight that *reduced information inference due to limited awareness can foster communication in situations with a conflict of interest.* The effect is not restricted to our renegotiation game but also arises in other communication settings. To see this, consider Crawford and Sobel's (1982) canonical cheap talk model. Their famous result shows that in the absence of commitment, any conflict of interest severely limits the information transmission that can be sustained in equilibrium. Communication is necessarily coarse. Our renegotiation game corresponds to the Crawford and Sobel (1982) setting if we remove the commitment assumption of the principal. Indeed, without commitment the initial delegation set plays no role in the renegotiation stage. Hence, for a fully aware principal the renegotiation game is identical to the standard cheap talk game.

What are the consequences of removing the principal's commitment and how does the set of equilibria differ from that of the standard cheap talk game? As under full awareness, removing the principal's commitment changes profoundly the nature of the game and Proposition 5 no longer applies. First of all, when becoming aware of a new action, the principal's equilibrium inference is no longer based on the initial delegation set as in (12) but on an analogous condition, where the 'reference set' $D^*(Y)$ is replaced by the set S of actions proposed by the agent and accepted by the principal in all contingencies θ where the principal's awareness does not increase.[21] Still, after correcting for the principal's inference, one might hope to characterise the set of implementable actions by the analogue of (11) with S replacing $D^*(Y)$. This is however not possible, since in the absence of commitment, the principal (*not the agent*) gets to pick her (interim) most preferred action in $Y \cup x$.

Intuitively, the best case scenario in terms of implementable actions and the agent's payoff is when S is small—so that inference is coarse—and Y is small—so that the principal's choices at the interim stage are limited. Consider an extreme case where the principal is initially aware of a single action y. If the agent reveals no further actions to the principal, the principal is forced to take y. The situation is thus analogous to one with a contracting phase and an initial delegation set $\{y\}$. Following our arguments above, it is then easy to see that the communication game without contracting has an equilibrium where the principal accepts all recommendations x such that $V^P(\{y, x\}) \geq V^P(\{y\})$. Since the agent is upward biased, this condition is always satisfied for $x < y$. It will also holds for certain $x > y$, as long as y is sufficiently low. Hence, in stark contrast to the case of full awareness, communication without commitment can result in full disclosure of information (from a modeller's point of view) on an interval of states.

Corollary 7. *(Cheap Talk with Unaware Receiver) Suppose the principal cannot commit and communication takes place after the agent observes θ. If $Y = \{y\}$, then there is a PBE with equilibrium outcome*

$$\forall\, \theta \in [0,1], \qquad y^E(\theta) = \max_{x \in \hat{Z}(\{y\})} U^A(\theta, x),$$

$$\textit{where} \qquad \hat{Z}(\{y\}) = \left\{ z \in Y^A : V^P(\{y, z\}) \geq V^P(\{y\}) \right\}.$$

Finally, it is interesting to point out that improvements in the information transmission due to limited awareness not only benefit the agent but typically also the principal gains. To see this, let Assumption 1 be satisfied and consider an initial awareness set Y and a

[21] Since the cheap talk game with constant awareness has typically multiple equilibria, S is not uniquely determined by Y but depends on the equilibrium.

communication equilibrium with A as the set of actions the principal takes on path. If Y is sufficiently rich, then equilibrium information transmission is necessarily coarse, which means that A is not an interval. By shrinking the set Y, e.g. to a singleton, gaps between actions belonging to A can be filled. Under Assumption 1, the principal benefits from closing such gaps ex-post. Hence, in a situation where the principal has no commitment, the principal may actually benefit from having limited awareness.

5 Discussion

Although we hope that our framework is able to bring important insights about general principal-agent problems where the agent has superior awareness of actions, we focus on a relatively simple delegation framework to identify the main effects. In this section, we discuss a few modifications of our baseline model (without renegotiation) where the results directly extend.

Set of feasible actions. So far we assumed that the set of available actions is an interval. A possible concern is that the principal, despite being unaware of certain elements of the set of feasible actions, might understand that this set is an interval and try to implement the optimal delegation set under full awareness by describing it indirectly. In particular, the principal could attempt to include actions outside her awareness, maybe through a description of the properties of such actions. Although Proposition 1 relies on a variational argument and hence uses the fact that Y^A is a continuum, the general point we are making also applies to the case with discrete Y^A, so that *a priori* there is no specific structure of the set of available actions—or simply the awareness set of the agent—that might be commonly known. In both cases, the agent has incentives to leave the principal unaware of actions around the optimal cap under full awareness.

For an illustration consider the specification with quadratic utilities and a uniform bias as in Section 3.1 and assume that Y^A is an arbitrary closed subset of $\mathbb{R}$. The analysis of the optimal delegation set for a given awareness set $Y \subseteq Y^A$ remains valid, so we have

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

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

Multidimensional Actions One might further wonder how our analysis extends to situations where choices are multidimensional. First of all, in several applications, although choice sets are multidimensional, the choice is often restricted either by budget/resource or other feasibility constraints so that the relevant choice becomes unidimensional (e.g., Amador et al. (2006) and Amador and Bagwell (2013)).

In addition, the results of this paper directly apply to the case where y represents an index obtained by aggregating the different actions. For $n = 1, 2, \dots N$, let $x_n \in X^A_n$, with $X^A := \times_{n=1}^N X^A_n$ an arbitrary subset of $\mathbb{R}^N$. Then we might assume there is a function $g : X \to \mathbb{R}$ such that $y := g(x_1, \dots, x_N)$, with $Y^A = G(X^A)$ and $U^i(\cdot, \theta)$ for $i = A, P$ as before. Obviously, not knowing y amounts to not knowing any $x \in X^A$ such that $g(x) = y$. In a finance application, x might be a vector of attributes of the asset (return, volatility, liquidity, maturity, etc.) and the index y might be interpreted as 'the asset', which must be confronted to 'the state of the market' θ.

Another multidimensional case where our results apply is when θ is unidimensional and the objective function is additively separable and identical across the different choice dimensions. In this case, the optimal contract under full awareness replicates the same delegation set in each dimension (Koessler and Martimort, 2012).

Allowing for more general frameworks, such as those with a multidimensional state θ and multidimensional and/or asymmetric aggregators significantly complicates the model. Such multi-dimensional setups, among other things, roughly imply the possibility of imperfect transfers between the principal and the agent. This reduces the possibility of pooling and induces agent's choices that are generically different from the agent's bliss points for a given dimension (see Koessler and Martimort 2012). Our results use the pooling point of the unidimensional case, so an extension to this framework would in general require a separate analysis. The main principle, however, should apply to this case as well. Take for example the bi-dimensional version of the quadratic case considered above, and assume that $\beta_1 = -\beta_2 > 0$. This is a particularly simple case, since under full awareness, the principal is able to obtain her first-best allocation: $y_i(\theta) = \theta - \beta_i$, $i = 1, 2$. It is immediate to see that the agent can increase his information rents by hiding actions $[-\beta_1 - \Delta_1, 0]$ in the first dimension and actions $[1, \beta_2 + \Delta_2]$ in the second dimension for $\Delta_i \geq 0$ sufficiently large (of course, assuming these actions are outside the awareness set of the principal).

Private awareness. There are many situation where the principal's initial awareness may not be known to the agent. For instance, a financial expert may be uncertain about the investment options an investor has encountered before their meeting, a division manager may not know which of the feasible projects are known to the headquarter, etc. The following proposition shows that uncertainty over the principal's initial awareness does not fundamentally change the solution of the problem. For simplicity, we restrict again attention to the specification of Section 3.1.

Proposition 8. *Let utility functions be specified by (6) and condition (7) be satisfied. Furthermore, let the agent's belief about the principal's awareness set be described by a probability distribution over the set of closed subsets of Y^A with finite support. The set of actions that the agent optimally reveals takes the form*

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

for some $\Delta \geq 0$. If the agent assigns a positive probability to the event that the principal's *awareness set does not include $\hat{y}$, then $\Delta > 0$.*

Proof. See Appendix A.8. □

In the proof of Proposition 8, we show that the agent can improve on an arbitrary awareness set by disclosing all actions that have a weakly greater distance to $\hat{y}$ than the closest action in the set. No matter what the realised awareness set of the principal is, actions that are further away from $\hat{y}$ than the closest one do not crowd out any additional actions. The optimal size of the awareness gap for the agent is determined by his beliefs about the principal's initial awareness. Proposition 8 shows that the agent leaves the principal unaware of some actions whenever he assigns a strictly positive probability to the event that the principal's awareness is bounded away from $\hat{y}$. The agent's cost of not disclosing actions around $\hat{y}$ when facing a principal who is aware of $\hat{y}$ is the (potential) loss of flexibility below $\hat{y}$. However, as we argue in Section 3, at $\Delta = 0$ the agent's marginal utility loss associated to the reduced flexibility equals zero, since in states below $\hat{y}$ the agent is at his bliss point. This implies that as long as the agent assigns a positive probability to the event that the principal is not aware of $\hat{y}$ and, hence, that there is a strict gain of introducing a gap, the net effect of marginally increasing Δ at $\Delta = 0$ is positive.

Principal's sophistication and awareness of unawareness We assumed that when additional options get revealed to the principal, she updates her awareness to the union of what she knew initially and what the agent reveals. This might suggest that the principal is naive, as she does not contemplate the possibility of other actions of which she is not aware. With regard to this point, we should note that the agent's equilibrium announcement is justifiable for the principal in the sense that it is consistent with the principal believing that the agent acts rationally. In our game, the requirement of justifiability, introduced by Ozbay (2007), states that given the principal's updated awareness Y, the principal cannot conceive of any announcement strategy which yields a higher expected payoff for the agent than Y. In other words, in equilibrium the principal should not believe that if the agent had revealed fewer actions, he would have been better off. Since the delegation set resulting from the optimal awareness set Y^* yields a weakly greater payoff for the agent than the payoff associated to any other subset $Y \subseteq Y^*$, this requirement is always satisfied in our setting.

It would however be interesting to consider a dynamic environment—much richer then ours—where the principal has ways to expand her awareness set (for example, by using a costly technology or by sampling other agents). In this case, the initial awareness of being partially aware and, perhaps more importantly, the assessment of the value of discovering new options would be key (see, for instance, Karni and Vierø (2017)). Moreover, it would be natural to assume that becoming aware can change the principal's perception about the possibility of being unaware of further actions. We leave this interesting issue for future research.

6 Conclusion

This paper formulates a flexible model of delegation with limited awareness and derives a number of 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 allows for a number of useful applications that span from financial intermediation, human resources, and political economy. We however believe that a key component of the contribution is to provide at least two general insights that apply to games with principal-agent or sender-receiver structure where the agent/sender 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 contracting parties with superior awareness. This principle is not restricted to problems of delegation but applies to other contracting problems as well. 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 the set of feasible contracts, the agent can increase the principal's cost of such distortions and thereby increase 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, which may provide an interesting new angle to look at principal-agent relationships.[22]

Second, 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 he observes, the player cannot infer information from the fact that a different signal did not realise. This asymmetry gives rise to non-standard information structures and hence to rather different equilibrium outcomes with respect to the full awareness benchmark. In the context of our game, we show that limited awareness can foster communication considerably and improve equilibrium outcomes in situations with conflicts of interest.

References

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. Amador, I. Werning, and G.-M. Angeletos. Commitment vs. flexibility. *Econometrica*, 74(2):365–396, 2006.

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.

[22]This is related to a recent literature looking at determinants of agency rents, which we mention in the Introduction. The focus of this literature is on the information structure as the modifiable primitive.

S. Auster and N. Pavoni. Limited awareness and financial intermediation. mimeo, 2020.

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.

M. X. D. Carpini and S. Keeter. *What Americans know about politics and why it matters.* Yale University Press, 1996.

J. Copic and A. Galeotti. Awareness as an equilibrium notion: Normal-form games. mimeo, 2006.

V. P. Crawford and J. Sobel. Strategic information transmission. *Econometrica: Journal of the Econometric Society*, pages 1431–1451, 1982.

Y. Feinberg. Games with unawareness. *The BE Journal of Theoretical Economics*, 1(ahead-of-print), 2020.

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.

D. Garrett, G. Georgiadis, A. Smolin, and B. Szentes. Optimal project design. mimeo, 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.

P. Guarino. An epistemic analysis of dynamic games with unawareness. *Games and Economic Behavior*, 120:257–288, 2020.

L. Guiso and T. Jappelli. Awareness and stock market participation. *Review of Finance*, 9 (4):537–567, 2005.

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.

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*, 1(ahead-of-print), 2020.

B. Holmström. On incentives and control in organisations, chapter ii: Delegation. *unpublished PhD dissertation*, 1977.

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.

N. Kartik, R. Van Weelden, and S. Wolton. Electoral ambiguity and political representation. *American Journal of Political Science*, 61(4):958–970, 2017.

F. Koessler and D. Martimort. Optimal delegation with multi-dimensional decisions. *Journal of Economic Theory*, 147(5):1850–1881, 2012.

E. Kováč and T. Mylovanov. Stochastic mechanisms in settings without monetary transfers: The regular case. *Journal of Economic Theory*, 144(4):1373–1395, 2009.

A. Lei and X. Zhao. Delegation and unawareness. mimeo, 2020.

D. Martimort and A. Semenov. Continuity in mechanism design without transfers. *Economics Letters*, 93(2):182–189, 2006.

M. Meier and B. C. Schipper. Bayesian games with unawareness and unawareness perfection. *Economic Theory*, 56(2):219–249, 2014.

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. Y. Ozbay. Unawareness and strategic announcements in games with uncertainty. In *Proceedings of the 11th conference on Theoretical aspects of rationality and knowledge*, pages 231–238, 2007.

A.-K. Roesler and B. Szentes. Buyer-optimal learning and monopoly pricing. *American Economic Review*, 107(7):2072–80, 2017.

I. Somin. *Democracy and political ignorance: Why smaller government is smarter*. Stanford University Press, 2016.

D. Szalay. The economics of clear advice and extreme options. *The Review of Economic Studies*, 72(4):1173–1198, 2005.

E.-L. Von Thadden and X. Zhao. Incentives for unaware agents. *The Review of Economic Studies*, 79(3):1151–1174, 2012.

E.-L. Von Thadden and X. Zhao. Multi-task agency with unawareness. *Theory and decision*, 77(2):197–222, 2014.

X. Zhao. Framing contingencies in contracts. *Mathematical Social Sciences*, 61(1):31–40, 2011.

Appendix

A Proofs

A.1 Proof of Proposition 1

Preliminaries. We start by introducing two functions. First, define $s := (y^A)^{-1}$ as the inverse of the agent's bliss point function y^A; that is, $s(y)$ *represents the state in which* y *is the most preferred action for the agent.* Second, let $t : (Y^A)^2 \to [0,1]$ be a symmetric function, indicating *the state at which the agent is indifferent between any two action* y *and* y'. It is specified as follows. For $y = y'$, set $t(y,y') = s(y)$. 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

$$U^A(t(y,y'),y) = U^A(t(y,y'),y'). \tag{13}$$

Due to the single-crossing condition, the solution of (13) 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 9. *Consider* y_0 *such that* $s(y_0) \in (0,1)$*, then*

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

Proof. For the case where t is determined by (13), we apply the implicit function theorem to derive

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

Taking the limit, we have

$$\begin{aligned}
\lim_{y\to y_0} \frac{\mathrm{d}t(y,y_0)}{\mathrm{d}y} &= \lim_{y\to y_0} \frac{U_y^A(t(y,y_0),y)}{U_\theta^A(t(y,y_0),y_0) - U_\theta^A(t(y,y_0),y)} \\
&= \lim_{y\to y_0} \frac{U_{\theta y}^A(t(y,y_0),y)\frac{\mathrm{d}t(y,y_0)}{\mathrm{d}y} + U_{yy}^A(t(y,y_0),y)}{\left(U_{\theta\theta}^A(t(y,y_0),y_0) - U_{\theta\theta}^A(t(y,y_0),y)\right)\frac{\mathrm{d}t(y,y_0)}{\mathrm{d}y} - U_{\theta y}^A(t(y,y_0),y)} \\
&= \frac{U_{\theta y}^A(s(y_0),y_0)\lim_{y\to y_0}\frac{\mathrm{d}t(y,y_0)}{\mathrm{d}y} + U_{yy}^A(s(y_0),y_0)}{-U_{\theta y}^A(s(y_0),y_0)}
\end{aligned}$$

where the second equality follows from L'Hôspital's rule. Also recall $t(y_0, y_0) = s(y_0)$. We can solve the above equality for $\lim_{y\to y_0} \frac{\mathrm{d}t(y,y_0)}{\mathrm{d}y}$ and obtain

$$\lim_{y\to y_0} \frac{\mathrm{d}t(y, y_0)}{\mathrm{d}y} = -\frac{1}{2} \cdot \frac{U^A_{yy}(s(y_0), y_0)}{U^A_{y_0}(s(y_0), y_0)}$$

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

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

The two results together establish the claim. □

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 10. *Generically, there exists some $\underline{y} < \hat{y}$ such that for all $y \in (\underline{y}, \hat{y})$,*

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

Proof. Under full awareness, the upper bound of the principal's optimal delegation is $\hat{y}$. Under the assumption that $\hat{y}$ is a limit point $D^*(Y^A)$, $\hat{y}$ must maximise

$$\int_{\underline{y}}^{s(y)} U^P(\theta, y^A(\theta))dF(\theta) + \int_{s(y)}^{1} U^P(\theta, y)dF(\theta) \tag{14}$$

over y for some $\underline{y} < \hat{y}$. The first-order condition is

$$\int_{s(\hat{y})}^{1} U^P_y(\theta, \hat{y})dF(\theta) = 0 \tag{15}$$

and the second-order condition is

$$\int_{s(\hat{y})}^{1} U^P_{yy}(\theta, \hat{y})dF(\theta) - U^P_y(s(\hat{y}), \hat{y})f(s(\hat{y}))s'(\hat{y}) \leq 0. \tag{16}$$

Since s is fully determined by U^A and thus exogenous and since $W'(y) := \int_{s(y)}^1 U^P_y(\theta, y)dF(\theta)$ is strictly decreasing on a neighbourhood around $\hat{y}$ (recall that $V^P(D^*(Y^A))$ is a strict maximum), *generically* this condition holds as a strict inequality.

Define the difference between the principal's expected payoffs under delegation sets $\bar{D}(y) \cup \{\hat{y}\}$

and $\bar{D}(y)$:

$$\Delta W(y) := V^P\left(\bar{D}(y) \cup \{\hat{y}\}\right) - V^P\left(\bar{D}(y)\right)$$
$$= \int_{t(y,\hat{y})}^1 U^P(\theta,\hat{y})dF(\theta) - \int_{t(y,\hat{y})}^1 U^P(\theta,y)dF(\theta).$$

The first derivative is given by

$$\Delta W'(y) = -\int_{t(y,\hat{y})}^1 U_y^P(\theta,y)dF(\theta) - \left(U^P(t(y,\hat{y}),\hat{y}) - U^P(t(y,\hat{y}),y)\right) f(t(y,\hat{y})\frac{d\,t(y,\hat{y})}{d\,y}$$

with

$$\Delta W'(\hat{y}) = -\int_{s(\hat{y})}^1 U_y^P(\theta,\hat{y})dF(\theta).$$

By (15), the above term is equal to zero. We must therefore consider the second derivative:

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

with

$$\Delta W''(\hat{y}) \quad = \quad -\int_{s(\hat{y})}^1 U_{yy}^P(\theta,\hat{y})dF(\theta) + 2U_y^P(s(\hat{y}),\hat{y})f(s(\hat{y}))\left.\frac{d\,t(y,\hat{y})}{d\,y}\right|_{y=\hat{y}}$$

Since, by Lemma 9, we have $-\left.\frac{d\,t(y,\hat{y})}{d\,y}\right|_{\varepsilon=0} = \frac{1}{2}s'(\hat{y})$, condition (16) holding as a strict inequality implies $\Delta W''(\hat{y}) > 0$. Remembering $\Delta W'(\hat{y}) = 0$, there is then an interval for y to the left of $\hat{y}$, where $\Delta W'(y) < 0$. With $\Delta W(\hat{y}) = 0$, this property implies, in turn, that there is an interval for y to the left of $\hat{y}$ such that $\Delta W(y) > 0$. Hence, for y sufficiently close to $\hat{y}$, the inequality $V^P\left(\bar{D}(y)\right) < V^P\left(\bar{D}(y) \cup \{\hat{y}\}\right)$ holds. □

Perturbation. Lemma 10 shows that, generically, there exists some $\underline{y} < \hat{y}$, such that for all $y \in (\underline{y}, \hat{y}]$, $V^P(\bar{D}(y)) < V^P(\bar{D}(\bar{y}) \cup \{\hat{y}\})$ is satisfied. Hence, for each $y \in (\underline{y}, \hat{y}]$ 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.[23] Due to this property, we can find a continuous, strictly decreasing function which maps each $y \in (\underline{y}, \hat{y}]$ to a value of δ

[23]Since the agent simply chooses the better of the two actions closest to him, it is immediate to see that $V^P(\bar{D}(\bar{y}) \cup \{\hat{y}\})$ is continuous in $\hat{y}$ (see expression (14)).

satisfying this condition. Letting $y(\cdot)$ denote the inverse of this function, we define

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

as the associated awareness set.

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 $D \subseteq Y$ to solve (3). Holmström (1977, Theorem 1) provides conditions under which the principal's value function V^P is upper semi-continuous. These conditions are satisfied in our framework. In particular:

- The set $Y^A = [y_{min}, y_{max}]$ is a compact subset of $\mathbb{R}$, a complete, separable metric space.
- $\mathcal{D}(Y^A)$ is a closed subset of $2^{[Y^A]}$ with respect to the Hausdorff-metric

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

- U^A and U^P are uniformly bounded on their domains $[0, 1] \times [y_{min}, y_{max}]$.

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 neighbourhood of zero. Suppose this is not true. Then, since $y_1(\cdot)$ is bounded above by $\hat{y}$, there exists a (sub)sequence $\{\delta_n\}$ with $\lim_{n \to +\infty} \delta_n = 0$ such that $\lim_{n \to +\infty} y_1(\delta_n) = y_0 \leq \hat{y}$, with y_0 possibly depending on the sequence. To violate continuity it must be that one of them satisfies $y_0 < \hat{y}$. Denote this sequence by $\{\bar{\delta}_n\}$. Then,

$$d_H(D \backslash (y_0, \hat{y}), D^*(Y^A)) \geq \frac{\hat{y} - y_0}{2}.$$

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

The principal's optimisation 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 $Y(\delta)$ satsifies $\bar{D}(y_1(\delta)) \subseteq D^*(Y(\delta))$.

Next, we show that the principal permits at least one action above $\hat{y}$. To see this notice that for all D with $\max D \leq y(\delta)$, the following inequalities hold:

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

The first inequality follows from the facts that 1) generating a payoff close to $V^P(D^*(Y^A))$ requires that $\max D$ is close to $\hat{y}$, 2) conditional on $\max D$ being close to $\hat{y}$, we have $V^P(D) \leq V^P(\bar{D}(\max D))$, and 3) $V^P(\bar{D}(y))$ is strictly increasing on a left neighbourhood of $\hat{y}$. We thus established $\max D^*(Y(\delta)) > \hat{y}$.

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 δ on a right neighbourhood of zero. Since the function y_2 is bounded below by $\hat{y} + \delta$ and satisfies $y_2(0) = \hat{y}$, it must be increasing on a right neighbourhood of 0.

$D^*(Y(\delta))$ thus satisfies

$$\bar{D}(y_1(\delta)) \cup \{y_2(\delta)\} \subseteq D^*(Y(\delta)). \tag{17}$$

Agent optimality. Let V^A denote the agent's value as a function of the delegation set. Property (17) implies that $V^A(\bar{D}(y_1(\delta)) \cup \{y_2(\delta)\})$ 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 $D(y^+) := \bar{D}(y^-(y^+)) \cup \{y^+\}$ can then be written as:

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

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

$$\begin{aligned} \frac{\mathrm{d}\, V^A(D(y^+))}{\mathrm{d}\, y^+} \;&=\; \int_{s(y^-(y^+))}^{t(y^-(y^+),y^+)} U_y^A(\theta, y^-(y^+))y'(\varepsilon)\mathrm{d}F(\theta) \\ &\quad + \int_{t(y^-(y^+),y^+)}^{1} U_y^A(\theta, y^+)\mathrm{d}F(\theta). \end{aligned}$$

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

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

Since $U_y^A(s(\hat{y}), \hat{y}) = 0$ and $U_{\theta y}^A > 0$, we have $U^A(\theta, y) > 0$ for all $\theta > s(\hat{y})$. The derivative of the agent's value at $y^+ = \hat{y}$ is thus strictly positive.

Taken toegther, there exists an $y^+ > \hat{y}$ and an associated awareness set that yields a delegation set which the agent strictly prefers to $D^*(Y^A)$. Hence, revealing all actions in Y^A is strictly dominated for the agent.

□

A.2 Proof of Proposition 2

Proof. First, we prove statement (i). Let $\hat{y} \notin Y^P$. Since Assumption 1 is satisfied, the optimal delegation set under full awareness takes the form $D^*(Y^A) = [y^A(0), \hat{y}]$. Hence, $\hat{y}$ is a limit point of $D^*(Y^A)$ and, by Proposition 1, the agent does not disclose all actions. Next, let $\hat{y} \in Y^P$. We want to show that full disclosure is optimal for the agent. 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 $\hat{y} \in Conv(D^*(Y))$. 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), 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 (5) is single-peaked, 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:

$$\int_{s(y)}^{1} U^P(\theta, y)\,\mathrm{d}F(\theta) > \int_{s(y)}^{s(\bar{y})} U^P(\theta, y^A(\theta))\,\mathrm{d}F(\theta) + \int_{s(\bar{y})}^{1} U^P(\theta, \bar{y})\,\mathrm{d}F(\theta). \tag{18}$$

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:

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

where the weak inequality follows from Assumption 1 and the strict inequality follows from (18). 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^* with $\bar{y} = \max D^*(Y^*)$. The set Y^* must clearly satisfy $\max D^*(Y^*) \geq \hat{y}$, as any delegation set with an upper bound smaller than $\hat{y}$ is dominated by the full awareness delegation set. We set $\Delta_2 = \max D^*(Y^*) - \hat{y}$ and define $\Delta_1 \geq 0$ as the smallest value that satisfies the equality

$$\begin{aligned}
&\int_{s(\hat{y}-\Delta_1)}^{1} U^P(\theta, \hat{y} - \Delta_1)\, \mathrm{d}F(\theta) \\
= \; &\int_{s(\hat{y}-\Delta_1)}^{t(\hat{y}-\Delta_1, \hat{y}+\Delta_2)} U^P(\theta, \hat{y} - \Delta_1)\, \mathrm{d}F(\theta) + \int_{t(\hat{y}-\Delta_1, \hat{y}+\Delta_2)}^{1} U^P(\theta, \hat{y} + \Delta_2)\, \mathrm{d}F(\theta).
\end{aligned}$$

Notice that $\Delta_2 = 0$ if and only if $\Delta_1 = 0$. 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 $\bar{y}$. All actions in the interval $(\hat{y} - \Delta_1, \hat{y} + \Delta_2)$ would crowd out $\hat{y} + \Delta_2$, hence $Y^* \cap (\hat{y} - \Delta_1, \hat{y} + \Delta_2) = \emptyset$. At the same time, making the principal aware of actions weakly below $\hat{y} - \Delta_1$ does not crowd out any actions. If these actions are weakly above $y^A(0)$, they strictly expand the agent's choice set, so the agent strictly prefers to reveal them. 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 $D^*(Y^*) = [y^A(0), \hat{y} - \Delta_1] \cup \{\hat{y} + \Delta_2\}$. □

A.3 Proof of Proposition 3

For the proof of this and the following proposition, it is useful to introduce the terms

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

and

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

in the literature referred to as, respectively, backward bias and forward bias (see Alonso and Matouschek 2008). By condition (7) we have

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

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

Having introduced these terms, we are ready to prove Proposition 3. This proof is presented via three lemmas.

Lemma 11. *Let condition (7) be satisfied and consider $y_1, y_2 \in Y$ with $y_1 < y_2$. If $y_1, y_2 \in D^*(Y)$,* *then all $y \in (Y \cap (y_1, y_2))$ belong to $D^*(Y)$.*

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

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

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

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

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

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

Lemma 12. *The optimal delegation set satisfies* $\min D^*(Y) = \min Y$.

Proof. Consider delegation set D with $\min D(Y) > \min Y$. Letting $y = \min Y$ and $\underline{y} = \min D(\widehat{y})$, the state at which the agent is indifferent between the two actions is given by $t = (y + \underline{y})/2$. If the principal includes y in the delegation set, the agent switches from $\underline{y}$ to y in all states $\theta \leq t$. The principal's change in expected payoff when including y is hence given by

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

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

Lemma 13. *Let condition (7) be satisfied. The optimal delegation set is such that*

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

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

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

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

The previous results together conclude the proof. □

A.4 Proof of Proposition 4

Let $\hat{y}(\beta)$ be the optimal cap under full awareness when the bias is β. Hence, $\hat{y}(\beta)$ describes the solution of the problem

$$\max_y -\int_{y_{min}}^{y} \beta^2 dF(\theta) - \int_y^1 \left(y - (\theta-\beta)\right)^2 dF(\theta). \tag{19}$$

and is implicitly defined by the first-order condition

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

The following second-order necessary condition must be satisfied at $\hat{y}(\beta)$:

$$\beta f(\hat{y}(\beta)) - (1 - F(\hat{y}(\beta))) \leq 0. \tag{20}$$

By condition (7), the third derivative of expression (19) with respect to $\bar{y}$, given by $\beta f'(\bar{y}) + f(\bar{y})$, is strictly positive. Hence, the first derivative is strictly convex. Strict convexity of the first derivative, together with the fact that it equals zero from below at $\bar{y} = 1$, implies that at $\hat{y}$ the first derivative crosses zero from above. Condition (24) must hence be satisfied with strict inequality. We can then use the implicit function theorem (condition (7) implies that the cumulate F is $\mathcal{C}^1$) to show that

$\hat{y}(\beta)$ admits a derivative at each β, which equals:

$$\hat{y}'(\beta) = -\frac{1 - F(\hat{y}(\beta)}{\beta f(\hat{y}(\beta)) - (1 - F(\hat{y}(\beta)))} < 0 \tag{21}$$

where we used the necessary second order condition (24) with strict inequality. Continuous differentiability is guaranteed by the implicit function theorem and can be checked directly in the above expression

Turning to the case of partial awareness, let us first write the agent's payoff as a function of Δ and the parameter β:

$$U(\Delta; \beta) = -\int_{\hat{y}(\beta)-\Delta}^{\hat{y}(\beta)} (\hat{y}(\beta) - \Delta - \theta)^2 \mathrm{d}F(\theta) - \int_{\hat{y}(\beta)}^{1} (\hat{y}(\beta) + \Delta - \theta)^2 \mathrm{d}F(\theta). \tag{22}$$

The first and second derivative of this function with respect to Δ are

$$U'_{\Delta}(\Delta; \beta) = 2\int_{\hat{y}(\beta)-\Delta}^{\hat{y}(\beta)} [\hat{y}(\beta) - \Delta - \theta] \,\mathrm{d}F(\theta) - 2\int_{\hat{y}(\beta)}^{1} [\hat{y}(\beta) + \Delta - \theta] \,\mathrm{d}F(\theta) \tag{23}$$

$$U''_{\Delta\Delta}(\Delta; \beta) = -2[1 - F(\hat{y}(\beta) - \Delta)] < 0. \tag{24}$$

Strict concavity implies that the agent's optimisation problem has a unique solution on $[0, \bar{\Delta}(Y)]$. The interior solution of this problem is characterised by the first-order condition

$$\int_{\hat{y}(\beta)-\Delta^*}^{\hat{y}(\beta)} [\hat{y}(\beta) - \Delta^* - \theta] \, f(\theta)\mathrm{d}\theta - \int_{\hat{y(\beta)}}^{1} [\hat{y}(\beta) + \Delta^* - \theta] \, f(\theta)\mathrm{d}\theta = 0. \tag{25}$$

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

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

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 (10) with respect to β keeping Δ^* as fixed, after some rearrangement, delivers:

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

Since $\hat{y}'(\beta) < 0$ (see (21)), we would be done if $1 + F\left(\hat{y}(\beta) - \Delta^*(\beta)\right) - 2F(\hat{y}(\beta)) > 0$. Note that this inequality can be equivalently written as:

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

Using $\hat{y}(\beta) = \mathbb{E}[\theta - \beta | \theta \geq \hat{y}(\beta)]$, the first order condition (25) can be written as:

$$[1 - F(\hat{y}(\beta) - \Delta^*(\beta))] \left[\mathbb{E}[\theta | \theta \geq \hat{y}(\beta) - \Delta^*(\beta)] - (\hat{y}(\beta) - \Delta^*(\beta))\right] = 2[1 - F(\hat{y}(\beta))]\beta. \tag{26}$$

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

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

Equivalently we can write

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

Given this inequality, (26) requires $(1 - F(\hat{y}(\beta) - \Delta^*)\beta))) < 2(1 - F(\hat{y}(\beta)))$, as desired.

A.5 The Game with Renegotiation

An equilibrium for the new extensive form game will be defined as an awareness choice Y at the root for the agent, a delegation set $D \subseteq Y$ in the contracting phase for the principal and a PBE at the renegotiation stage given (Y, D).

Suppose that after the contracting phase the principal's awareness is Y and the optimal delegation set is $D \subseteq Y$. The game played at the renegotiation phase is a signalling game. We now define the strategies of the principal and the agent at this stage. We will first of all concentrate on agent's moves constituted by either 'no new proposal' (let us call it N) or proposal of singletons y. Let $X := N \cup Y^A$ be the set of possible proposals of the agent. The N proposal will be interpreted as if the agent does not propose any action outside the original delegation set D, and hence, in particular, he does not increase the awareness of the principal. When the agent proposes a new action $x \in X$ the principal can decide whether to replace the original delegation set D with x.

It is clear that the agent will never benefit from proposals $x \in D$ as these are weakly dominated by the 'null' or 'no proposal' $x = N$. It will also be clear below that the agent will not be able to increase his payoff in the best equilibrium by augmenting the proposal/signal y with an additional payoff-irrelevant message. We hence - to simplify notation - do not allow for any further message at the renegotiation stage.

In Appendix B.3 we show that, given initial history (Y, D), that the component of agent's strategy at this stage can be described by a map from each realisation of $\theta \in [0, 1]$ into a recommendation $x \in X$ and then a map from the new delegation set D' into an action $y \in D'$.

At this stage, the agent's strategy is thus a pair of functions $x(Y, D, \cdot), y(Y, D, \cdot, \cdot)$, where $x(Y, D, \theta)$ is the proposal given (Y, D, θ), and $y(Y, D, \theta, D')$ is the final choice given (Y, D, θ, D'). The principal's strategy is a mapping $\rho(Y, D, \cdot) \in \{0, 1\}$, where $\rho(Y, D, x) = 0$ implies $D' = D$ while $\rho(Y, D, x) = 1$ implies $D' = \{x\}$.[24] With her action, the principal assigns to each proposal a delegation set which can be either the original D or the singleton $\{x\}$. If the agent proposes $x(Y, D, \cdot) = N$, the principal can only choose $\rho(Y, D, N) = 0$.

[24] Of course, we could have equivalently defined the feasibility set for ρ as the pair $\{D, x\}$.

We also concentrate on outcomes in pure strategies. We can hence define a beliefs' support function Θ for the principal mapping each triplet (Y, D, x) to a subset of values for θ. The set $\Theta(Y, D, x) \subseteq [0,1]$ describes the states which the principal considers possible when the agent proposes x when the principal has awareness level Y and decided D. For completeness, we set $\Theta(Y, D, N) = [0,1]$.

Definition 1. *Given, (Y, D), with $D \subseteq Y$, the strategy profile (x^*, y^*, ρ^*) and beliefs' support function Θ^* constitute a PBE of the renegotiation game if and only if the following conditions hold:*[25]

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

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

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

$$\begin{aligned} x^*(Y, D, \theta) &\in \arg\max_x \rho^*(Y, D, x) U^A(\theta, x) + (1 - \rho^*(Y, D, x)) U^A(\theta, y^*(Y, D, \theta, D)); \\ y^*(Y, D, \theta, D) &\in \arg\max_{y \in D} U^A(\theta, y); \end{aligned}$$

- *Consistency of beliefs: for all (Y, D, x) that are allowed on path,*

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

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

Finally, note that in the last stage, for each (Y, D), the function $y^*(Y, D, \cdot, \cdot)$ coincides with the function y^* we defined after condition (1), hence, its use in the main text (abusing notation). Recall we assumed that the principal is unaware of the renegotiation possibility.

A PBE equilibrium for the extensive-form game is constituted by an initial awareness choice Y^*, the delegation function $D^*(\cdot)$ as described in the 'pure' delegation game of the previous section together with a PBE equilibrium as defined in Definition 1 for each (Y, D), with $D \subseteq Y$. This notion is strong in the sense that it imposes PBE also for subtrees not reached along the equilibrium path.

Proposition 15 in Appendix B.3 shows the substantial equivalence between the equilibrium outcomes according to this definition and a more complete definition of equilibrium for extensive form games with unawareness we provide there.

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

A.6 Proof of Proposition 5

Fix an equilibrium and let Y denote the principal's awareness set after the contracting phase with associated delegation set $D^*(Y)$ (for a detailed description of strategies and beliefs see Section A.5). Consider the renegotiation phase and let S denote the set of actions that the agent will be able to take in the last stage for some contingency and that will actually be taken by the agent under some contingency. Since the agent always has the option not to renegotiate, we have $D^*(Y) \subseteq S$. We write $\bar{X} := S \setminus D^*(Y)$ and assume that $\bar{X}$ measurable.

We first want to show $\bar{X} \cap Y = \emptyset$. Towards a contradiction, suppose this intersection is nonempty and call the non-empty set $\hat{X} = \bar{X} \cap Y = (S \cap Y) \setminus D^*(Y)$. Consider the principal's subjective game after receiving a proposal $x' \in \hat{X}$. The principal is now aware that the agent can renegotiate and therefore of all proposals in Y. By definition of S, since $\hat{X}$ is disjoint from $D^*(Y)$, we have $\rho^*(Y, D^*(Y), x) = 1$ for all $x \in \hat{X}$ and $\rho^*(Y, D^*(Y), x) = 0$ for all $x \in Y \setminus S$. Consistency of beliefs implies

$$\Theta^*(Y, D^*(Y), x) = \{\theta \in [0,1] : U^A(\theta, x) \geq \sup_{x' \in S \cap Y} U^A(\theta, x')\} \text{ for all } x \in \hat{X}.$$

Given Θ^*, principal optimality requires for all $x \in \hat{X}$[26]

$$\int_{\Theta^*(Y,D^*(Y),x)} U^P(\theta, x)\, \mathrm{d}F(\theta) \geq \int_{\Theta^*(Y,D^*(Y),x)} U^P(\theta, y^*(\theta, D^*(Y)))\mathrm{d}F(\theta) \tag{27}$$

Summing (27) over all $x \in \hat{X}$, we obtain:

$$\int_{\hat{X}} \int_{\Theta^*(Y,D^*(Y),x)} U^P(\theta, x)\, \mathrm{d}F(\theta) dx \geq \int_{\hat{X}} \int_{\Theta^*(Y,D^*(y),x)} U^P(\theta, y^*(\theta, D^*(Y)))\mathrm{d}F(\theta) dx \tag{28}$$

It is easy to show that the properties of U^A imply $\{\Theta^*(Y, D^*(Y), x)\}_{x \in \hat{X}}$ is a collection sets with measure zero intersection. Letting $\Theta^C(Y, D^*(Y), \bar{X}) := [0,1] \setminus \left(\bigcup_{x \in \hat{X}} \Theta^*(Y, D^*(Y), x)\right)$, (28) can be written as

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

or equivalently

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

Since $\hat{X}$ is a subset of Y and D^* is the largest optimal 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. Hence, $S \cap Y = D^*(Y)$.

Having shown that the principal only accepts proposals in the renegotiation phase if they do

[26]Note that the set $\Theta^*(Y, D^*(Y), x)$ is measurable for all $x \in \hat{X}$, as both S and Y are measurable and U^A is continuous.

not belong to Y, consider proposal $x \in S \setminus Y$. Consistency of beliefs requires

$$\Theta^*(Y, D^*(Y), x) = \{\theta \in [0,1] : U^A(\theta, x) \geq \max_{y \in D^*(Y)} U^A(\theta, y)\}. \tag{29}$$

Principal optimality is then satisfied if and only if (27) holds (with Θ^* defined by (29)), or equivalently:

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

Finally, for each set A constituted of actions x satisfying (11) we can construct an equilibrium of the renegotiation game where proposal x is accepted by the principal whenever $x \in A$. To this end, we set $\rho^*(Y, D^*(Y), x) = 1$ for all $x \in A$ and $\rho^*(Y, D^*(Y), x) = 0$ otherwise, $x^*(Y, D^*(Y), \theta) = \arg\max_{x \in A} U^A(\theta, x)$ for all θ, and we define $\Theta^*(Y, D^*(Y), x)$ by (29) for all $x \in A$. Off-path beliefs can be set arbitrarily such that the principal rejects proposals outside A. This strategy and belief profile clearly satisfies principal optimality, agent optimality and consistency of beliefs and thus constitutes a PBE of the renegotiation game, as specified in Definition 1 of Section A.5. □

A.7 Proof of Proposition 6

We want to show that in the agent's best equilibrium the set of actions which the agent can implement with and without renegotiation is given by $S^* := [y^A(0), y^{max}]$. By definition of y^{max} and Proposition 5, there is no equilibrium where the principal accepts a proposal $y > y^{max}$. To see this, let S be the set of implementable actions with $\bar{y} = \max S$ and assume $\bar{y} > y^{max}$. By definition of y^{max}, there is no awareness set Y such that $Y^P \subseteq Y$ and $\bar{y} \in D^*(Y)$. Hence, $\bar{y}$ must be proposed in the renegotiation phase. By Proposition 5, $\bar{y} \in S$ implies that the awareness set Y after the contracting phase is such that

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

This inequality implies $\bar{y} \in D^*(Y \cup \{\tilde{y}\})$ and hence $Y^P \not\subseteq Y$, a contradiction. All we need to show is then that there exists an equilibrium where the set of the agent's implementable actions is S^*.

Recall the strategies we defined above and consider the following candidate equilibrium. In the contracting phase (the root of the game) the agent reveals a set Y such that $\max D^*(Y) = y^{max}$. Let this set be denoted by Y^{max}. The principal is unaware of the renegotiation stage and thus offers a delegation set that is optimal without renegotiation, $D^*(Y^{max})$. By Proposition 5, there is then an equilibrium where each x satisfying

$$V^P(D^*(Y^{max}), x) \geq V^P(D^*(Y^{max})) \tag{30}$$

is accepted in the renegotiation stage. For instance, we can specify, for all (Y, D), the agent's strategy by $x^*(Y, D, \theta) = y^A(\theta), \forall \theta$, the principal's strategy by $\rho^*(Y, D, x) = 1$ if $V^P(D \cup x) \geq$

$V^P(D)$ and $\rho^*(Y, D, x) = 0$ otherwise, and the beliefs by

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

for all x. We then want to show that (30) is satisfied if and only if $x \in S^*$. We distinguish three cases:

- If $x \in S^*$ and $x \leq \min D^*(Y^{max})$, then (30) follows from the fact that the agent is upward biased—conditional on the agent preferring x over $\min D^*(Y^{max})$, the principal prefers x as well.
- If $x \in S^*$ and $\min D^*(Y^{max}) < x \leq \max D^*(Y^{max})$, inequality (30) follows from Assumption 1.
- If $x \notin S$ and hence $x > y^{max}$, then (30) is violated by definition of y^{max}.

Taken together, the agent can implement any action belonging to S^*, hence the equilibrium outcome y^E is as specified in Proposition 6. □

A.8 Proof of Proposition 8

Proof. We want to show that revealing an awareness set of the form $[y^A(0), \hat{y} - \Delta] \cup \{\hat{y} + \Delta\}$ is optimal. Towards a contradiction, suppose this is not the case and let the optimal awareness set be denoted by Y. Define $\tilde{\Delta}$ to be the largest value of Δ such that $Y \subseteq (-\infty, \hat{y} - \Delta] \cup [\hat{y} + \Delta, +\infty)$ and $\tilde{Y} := [y^A(0), \hat{y} - \tilde{\Delta}] \cup \{\hat{y} + \tilde{\Delta}\}$. Suppose the principal's realised awareness set is Y^P. According to Proposition 3, the induced delegation sets from revealing, respectively, Y and $\tilde{Y}$ are

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

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

To prove the second part of the statement, let $EU(\Delta)$ denote the agent's expected payoff associated to the disclosure of a set of actions parametrised by Δ and let μ be the probability which the agent assigns to the event that the principal's awareness set does not include $\hat{y}$. Then, given our assumption that awareness sets are closed, there exists some $\varepsilon > 0$ such that no action in $(\hat{y}-\varepsilon, \hat{y}+\varepsilon)$ belongs to any of the realisations of the principal's awareness sets in the support that do not contain $\hat{y}$. For $\Delta \in [0, \varepsilon]$, the agent's expected payoff conditional on facing a principal who is not aware of $\hat{y}$ is then described by the function $U(\Delta)$, as defined in (22), with $U'(0) = -2\int_{\hat{y}}^{1} (\hat{y}-\theta)\,\mathrm{d}F(\theta) > 0$ (see (23))

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

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

with

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

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

□

B Generalised extensive-form games with unawareness

Heifetz et al. (2013, from now on HMS) define generalised extensive-form games that allow for evolving unawareness.

To introduce the generalised extensive-form game Γ, let N be a set of decision nodes, C be a set of chance nodes, and Z be a set of terminal nodes. The nodes $\bar{N} = N \cup C \cup Z$ constitute a tree. HMS capture limited awareness via the notion of subtrees, defined as subsets of nodes of $\bar{N}$. Letting $\mathbf{T}$ be a family of subtrees of $\bar{N}$, for $T, T' \in \mathbf{T}$ the relation $T' \preccurlyeq T$ signifies that the nodes of T' constitute a subset of the nodes of T. One element of $\mathbf{T}$ represents the modeler's view of the paths of play that are objectively feasible. The other elements of $\mathbf{T}$ represent feasible paths of play as subjectively viewed by some player, or as the frame of mind attributed to a player by other players or by the same player at a later stage of the game.

HMS propose a number of natural properties for generalised extensive-form games. These include basic extensions of standard requirements of extensive form games but also new properties that are specific to unawareness. All of these features are satisfied in our generalised game, which we describe in more detail now.

B.1 'Pure' Delegation Generalised Extensive-form Game

Consider first the model of the initial part of the paper where renegotiation is not possible. There are are three players: the principal, the agent, and nature/chance. Each game tree in $\mathbf{T}$ is very similar and all have virtually the same structure. An example of trees in $\mathbf{T}$ for the 'pure' delegation model is reported in Figure 3. At the root, the agent moves with $Y \supseteq Y^p$, then the principal proposes $D \subseteq Y$, then nature picks θ which is revealed to the agent,x and finally the agent picks $y \in D$ knowing θ. The agent's view is the richest one and coincides with the modeller's view and with the principal's view in case of full awareness. Let the associated game tree be denoted by T_{Y^A}. The other elements of $\mathbf{T}$ represent feasible paths of play as subjectively viewed by the principal. They also coincide with the way the agent see the principal's view. We denote by T_Y the subtree associated to principal's awareness level Y. The tree T_Y can be depicted starting from T_{Y^A} and deleting all moves Y' such that $Y \cap Y' = \emptyset$ or $Y \subset Y'$ (or both) and all nodes following those moves.

At the outset of the game, the principal's awareness is Y^P and her subjective view of the game is described by a subtree T_{Y^P}. Within the confined view of T_{Y^P} the principal can only envisage the agent announcing Y^P. If the agent reveals additional projects to the principal—i.e., he announces $Y \supset Y^P$—the principal updates her awareness and therefore her subjective view of the game. Given updated awareness Y, the principal's subjective game tree T_Y includes additional nodes. Once the principal becomes aware of the additional nodes, she can also contemplate less expressive game trees $T_{Y'}, Y^P \subseteq Y' \subseteq Y$. That is, the principal can envisage how the game would have unfolded had the agent revealed fewer actions. However, the principal cannot contemplate the paths of play that would have been feasible if the agent would have revealed more. After the initial stage, awareness no longer changes, which means that, given the constraints unawareness imposes on the players' succeeding moves, agent and principal play a standard game. Finally, it might be useful to note that, within the HMS structure, while the agent knows that the 'true' game tree is T_{Y^A}, the tree T_Y represents how the agent's thinks the principal sees the game when her awareness level is Y. Furthermore, the tree T_Y represents both how the principal actually sees the game when her awareness level is Y (so the agent is right) and what the principal thinks the agent's view of the game is (so the principal is wrong). The tree T_Y also represents all the higher order beliefs about each other views of both the agent and the principal.

Information Sets and Strategies. The 'forest' of trees $\mathbf{T}$ is hence constituted by $\mathbf{T} = \left(\{T_Y\}_{Y \in \mathcal{Y}(Y^P)}\right)$, where T_{Y^A} is the objective tree and, recall, $\mathcal{Y}(Y^P)$ is the set of all closed sets in $\mathbb{R}$ that are subset of Y^P.

A complete description of the strategic interactions among the players however requires considerations across the trees. The delegation game is of perfect information so for each player each information set contains a sigle node. The same node however, might appear in several trees. When a node is a decision node we treat the corresponding node in a different tree as a different node—we thereby generate copies of it—while keeping the obvious order of the nodes among all copies. Some nodes, however, do not constitute information sets of any player. So there is a redundancy in this description.

Information sets represent both information and awareness. In each tree T_Y, *the principal has exactly one information set*, where she decides the delegation set given his awareness. However, as explained above, given her awareness Y, the principal can now contemplate what she could have done under lower awareness. These views are described by nodes in T_Y that lead to information sets in less expressive trees: $T_{Y'}$ for $Y^P \subseteq Y' \subseteq Y$. The nodes representing lower awareness than Y in T_Y are reported in T_Y but they do not constitute principal's information sets as they do not represent the 'state of mind' the principal would have at that node. The *agent's information sets* are the root in each T_Y and the nodes (Y', D, θ) for $Y' \subseteq Y$, $D \subseteq Y$ an $\theta \in [0, 1]$.

A T_Y*-partial game* in our framework is constituted by the partially ordered set of trees including T_Y and all $T_{Y'}$ with $Y' \subseteq Y$, with all the information sets as specified above. The T_{Y^A} together with the linked less expressive trees and the information sets of the two players specifying subsets of the nodes in $\mathbf{T}$ constitutes the *generalised extensive-form game*.

The *agent's pure strategy* restricted to the T_Y-partial game can be described by a collection of pairs $\sigma_{Y'}$ and $y_{Y'}$, one pair for each $Y^P \subseteq Y' \subseteq Y$, where $\sigma_{Y'}$ represents the move at the root of the tree $T_{Y'}$ with feasibility $Y^P \subseteq \sigma_{Y'} \subseteq Y'$ and the function $y_{Y'}(\cdot)$ maps one $y \in D$ for each $\sigma_{Y'} \subseteq Y', \theta \in [0, 1], D \subseteq \sigma_{Y'}$; that is, $y_{Y'}(\sigma_{Y'}, \theta, D) \in D$. The subscript in the functions indicates the tree where the nodes constituting agent's information sets are located.

Recall that to each node $Y' \subset Y$ in T_Y we associated an information set (and hence a decision node) for the principal in tree $T_{Y'}$ (i.e., to a tree outside T_Y). The structure of $\mathbf{T}$ is relatively simple in that the less expressive trees are linked to more expressive ones in a simple way. This simplicity allows us to describe *a pure strategy for the principal* as a function $D(\cdot)$ that associate to each $Y \in \mathcal{Y}(Y^P)$ a set in $\mathcal{D}(\mathcal{Y})$.[27] Actions taken by the principal at a given informations set

[27]Since the principal has one information set in each tree, we could have equivalently written her strategy as an action D_Y in the sole information sets for the tree T_Y.

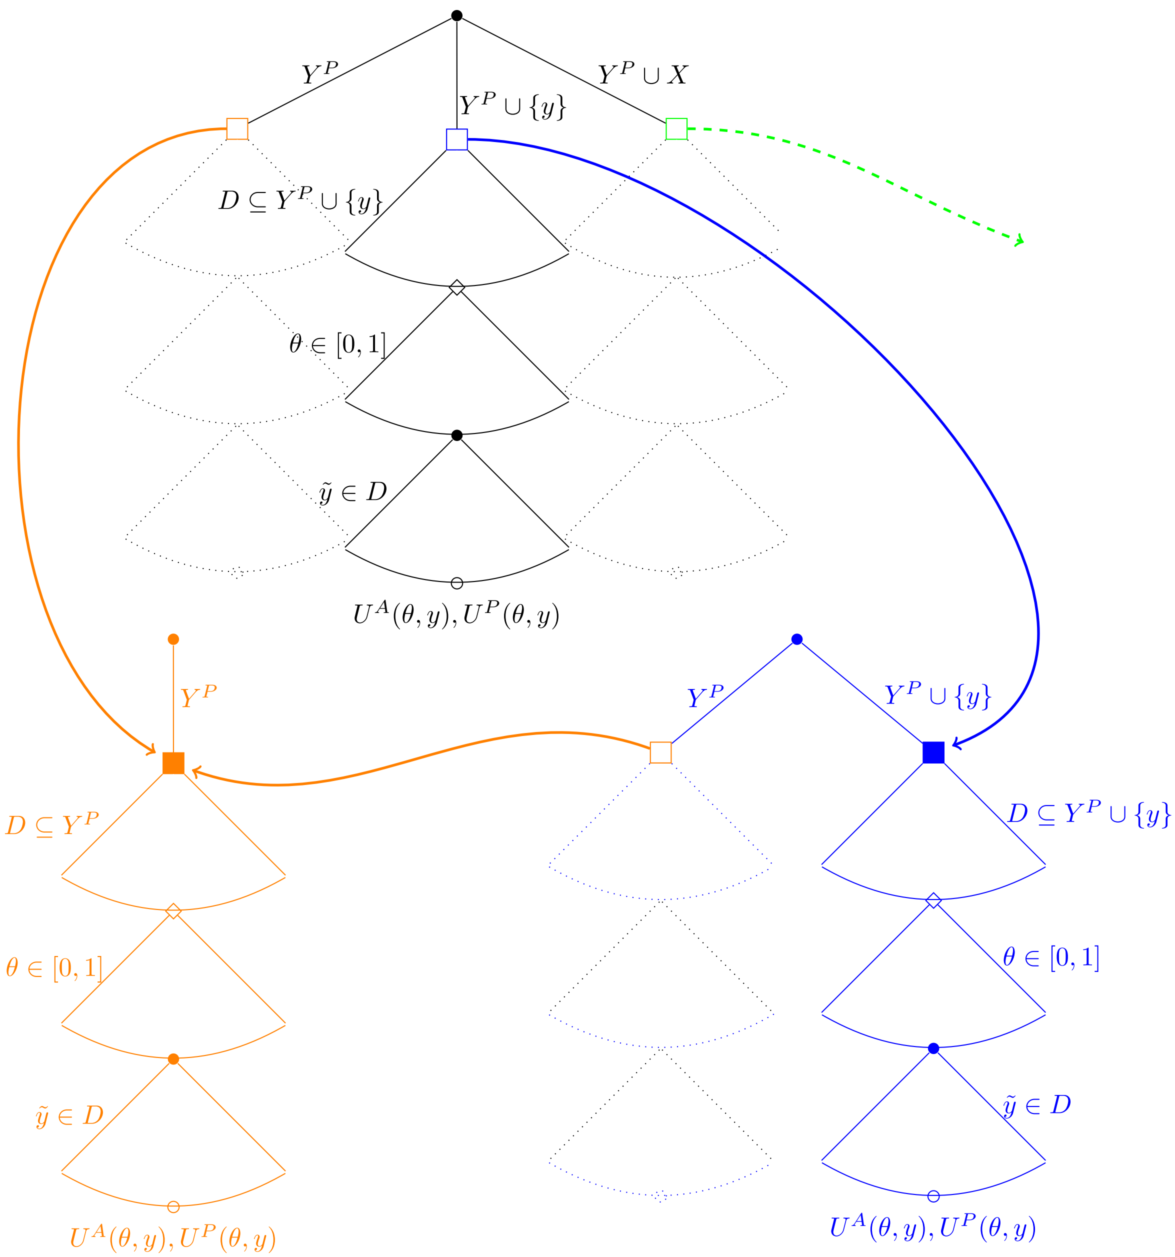

Figure 3: Objective and subjective game trees: circles are decision nodes of the agent, diamonds those of nature, and squares those of the principal. The tree on the top depicts the objective game tree. The tree on the bottom left illustrates the principal's subjective view of the game before the agent discloses additional actions. The tree on the bottom right shows the principal's subjective view of the game after becoming aware of action y. In the branches of the different trees we also indicate the principal's awareness. After updating the principal can contemplate how the game would have unfolded if she had not learned about y but she cannot contemplate branches of the objective game tree following disclosures of actions different than y. Solid squares (resp. circles) denotes the principal's (resp. agent's) information sets. The principal has exactly one information set in each of the bottom trees.

also generate the same moves at the corresponding nodes in trees with higher levels of awareness. Each function $D(\cdot)$ embeds all information for a complete description of the principal's strategy in the generalised game. In particular, given a function $D(\cdot)$, we can describe the principal's strategy restricted to a T_Y-partial game by considering the behaviour of D restricted to the domain $Y' \subseteq Y$.

B.2 Rationalizability and Prudent Rationalizability

In this section we show that *the equilibrium described in the main text is the unique* Δ*-rationalizable outcome of the generalised game we described above whenever we restrict to pure strategies and assume the two tie-breaking rules reported in the main text to be commonly believed.*

System of Beliefs A system of beliefs associates to each information set h of a player, his/her beliefs over the opponent's strategies restricted to the $T-$partial game that includes the information set, assigning probability one to the opponent's strategies that lead to the information set h. The system is 'prudent' whenever it gives otherwise full support to all 'conceivable' strategies of the opponent (see Heifetz et al. (2020) for details).

First, we describe the *principal's set of beliefs.* When confronted with awareness Y, the principal must attach probability one to all strategies of the agent that include the first move $\sigma_{Y'} = Y$ (so only to the node $\sigma_Y = Y$ in the tree T_Y) and then must form beliefs about the continuation of the agent's strategy $y_{Y'}$ in the last stage for all $Y' \subseteq Y$. However note, that the entry Y in the functions describing the component of the agent's strategy at the last stage is only in trees T_Y or more expressive trees. But such more expressive trees are not part of the T_Y-partial game. This implies that the principal must assign probability zero to all such functions. Moreover, since the agent's move at the root is taken before knowing the realisation of θ, all of the equilibrium considerations are useless for the principal in terms of inference based on forward induction considerations. As a consequence, what matters for the principal's beliefs at the (sole) information set in tree T_Y is the agent's strategy y_Y in that tree alone. This is what we will keep track of. So, our 'prudent rationalizability' sets omit some (redundant and irrelevant) elements compared to the defintion by Meier and Schipper (2014).

In our problem, the only *agent's beliefs* that need to be specified are those at the root about the principal's rationalizable strategies $D(\cdot)$. Those refer to actions taken over different trees. However, as argued above, they can be described by a single function.

To define prudent rationalizability, consider hence the following sequence of sets:

$$\begin{aligned} R_0^A(Y) &= \left\{(\sigma_Y, y_Y):\ \sigma_Y \in \mathcal{Y}, \sigma_Y \subseteq Y,\ \text{for } Y' \subseteq Y, D \subseteq Y',\ y_Y : (Y', \theta, D) \mapsto D\right\}, \\ R_0^P &= \left\{D:\ \forall Y \in \mathcal{Y},\ D(Y) \in \bar{\mathcal{D}}(Y)\right\}, \end{aligned}$$

and for $k > 0$:

$$R_k^A(Y) = \left\{ (\sigma_Y, y_Y) \in R_{k-1}^A(Y) : y_Y(Y', \theta, D) \in BR^A(\theta, D)\ \exists \mu^A \in \bar{\Delta}(R_{k-1}^P); \sigma_Y \in \arg\max_{\hat{Y} \supseteq Y^p} \int_{R_{k-1}^P} V^A(\hat{Y}, D(\hat{Y}), y_Y) d\mu^A \right\},$$

$$R_k^P = \left\{ D \in R_{k-1}^P : \forall Y \in \mathcal{Y}\ \exists \mu^P \in \bar{\Delta}(R_{k-1}^A(Y));\ D(Y) \in \arg\max_{\hat{D} \subseteq Y} \int_{R_{k-1}^A(Y)} V^P(Y, \hat{D}, y_Y) d\mu^P \right\}.$$

In the above:

1. The entry Y in $R^A(\cdot)$ indicates strategies defined on information sets located in the tree T_Y. Recall that those are the only ones we will keep track of in the T_Y-partial game.
2. $\bar{\Delta}(R_{k-1}^A(Y))$ is the set of distributions over $R_{k-1}^A(Y)$ that gives probability one to $\sigma_Y = Y$ and has 'full support' over all functions y_Y in $R_{k-1}^A(Y)$ that satisfy the tie-breaking rule;
3. $\bar{\Delta}(R_{k-1}^P)$ is the set of distributions over R_{k-1}^P with 'full support' over the delegation sets in it that satisfy the tie-breaking rule;
4. The set $BR^A(\theta, D)$ indicates the best response of the agent as defined in (1).
5. For $i = A, P$, $V^i(Y, D, y_Y) := \int_0^1 U^i(\theta, y_Y(Y, \theta, D)) dF(\theta)$.

The sets of prudent-rationalizable strategies for player $i = A, P$ are defined as:

$$R_\infty^A(Y) := \cap_{k=0}^\infty R_k^A(Y) \qquad \text{and} \qquad R_\infty^P := \cap_{k=0}^\infty R_k^P.$$

The technical complications involved in the definition of 'full support' are not analysed in detail because, as we will see next, the sets quickly shrink to singletons.

Consider indeed the sets $R_1^A(\cdot)$ and note that R_0^P allows for all $D \subseteq Y$ in each information set $\sigma = Y$, so all agent's information sets are reached. Furthermore, it is immediate to see that the objective function and the feasibility sets defining y_Y are identical across all sets. With an appropriate restriction of the domain (to be expressible within the tree T_Y) each of these rules can hence be described concisely by a function describing the agent's best strategy in all trees. The function y_Y would then coincide with this extended function when its domain is restricted to $Y' \subseteq Y$ and $D \subseteq Y'$. Moreover, the commonly believed tie-breaking rule implies that such function is unique. Recalling the notation used in Appendix A.5, let's denote this function by y^*.

Next, while R_1^P is still a rich set, we now argue that R_2^P is a singleton. The argument in the previous paragraph implies that for each Y, the belief μ^P must be concentrated on the singleton $(\sigma_Y, y_Y) = (Y, y_Y^*)$ (where we denote by y_Y^* the reduced domain version of y^*). This implies that, for each Y, when deciding D, the principal only considers the objective function $V^P(Y, D, y_Y^*)$ (with full probability weight on this combination alone). This, together with the tie-breaking rule implies that the principal has only one solution for each Y. This is what we denote in the main text as $D^*(Y)$. The solutions across the principal's information sets defines the function $D^*(\cdot)$

Finally, consider the sets $R_3^A(\cdot)$. We already argued that the last stage component can be described by the function y^*. Moreover, we argued in the previous paragraph that R_2^P is a singleton. The objective of the agent is hence $V^A(Y, D^*(Y), y_Y^*)$ with full weight on this combination alone.

Recalling that for all Y, D, θ we have $y_Y^*(Y, D, \theta) = y^*(Y, D, \theta)$, our findings can be summarised as follows:

Proposition 14 (Rationalizable Outcome)**.** *Let Y^P given. The 'relevant' set of (prudently) rationalizable strategies can be compactly defined as the sets (R^A, R^P) satisfying:*

$$
\begin{aligned}
R^P &= \left\{ D^* : \ \forall Y \in \mathcal{Y}, \ D^*(Y) \in \arg\max_{\hat{D} \subseteq Y} V^P(Y, \hat{D}) \text{ together with the tie-breaking rule} \right\}, \\
R^A &= \left\{ (Y, y^*) : \ y^*(Y, \theta, D) \in BR(\theta, D) \text{ with tie-breaking rule}, \ \ Y \in \arg\max_{\hat{Y} \supseteq Y^p} \ V^A(\hat{Y}, D^*(\hat{Y})) \right\},
\end{aligned}
$$

where for all $Y \in \mathcal{Y}$ and $D \in \bar{\mathcal{D}}(Y)$:

$$V^A(Y, D) \ := \ \int_0^1 U^A(\theta, y^*(Y, \theta, D)) dF(\theta); \tag{31}$$

$$V^P(Y, D) \ := \ \int_0^1 U^P(\theta, y^*(Y, \theta, D)) dF(\theta). \tag{32}$$

The proof is now complete, it is indeed immediate to see that the singletons described in R^P and R^A coincide with the Perfect Bayesian Nash equilibrium strategies described in the main text.

B.3 The Generalised Game with Renegotiation

In this subsection, we describe in detail the generalised extensive-form game representing our model with the possibility of renegotiation. An example of the forest of trees is reported in Figure 4.

There are two instances in the generalised game where the principal may change her view. The first one is after the initial move by the agent, as in the 'pure' delegation game we described in Section B.1. A second instance where the principal may change her awareness is at the renegotiation stage, whenever the agent proposes an action outside the awareness of the principal. The generalised game is hence constituted by more trees compared to the 'pure' delegation case.

Recall that we assume that the principal is initially unaware of the possibility of renegotiation. This assumption, first of all, implies that the generalised game includes trees very similar to the 'pure' delegation case. They represent situations where the principal only becomes aware of new actions right after the agent's move at the root and there is no proposal at the renegotiation stage. To be consistent with the previous notation, we denote all such threes as T_Y. However, to be consistent with a richer continuation of the tree in the renegotiation stage, we append to them a singleton move by the agent and a singleton move by the principal, the agent's move can be seen as an 'artificial' no-proposal which is followed by the principal's move assigning again the original D. These moves lead to a final stage that is identical to that in the 'pure' delegation case,

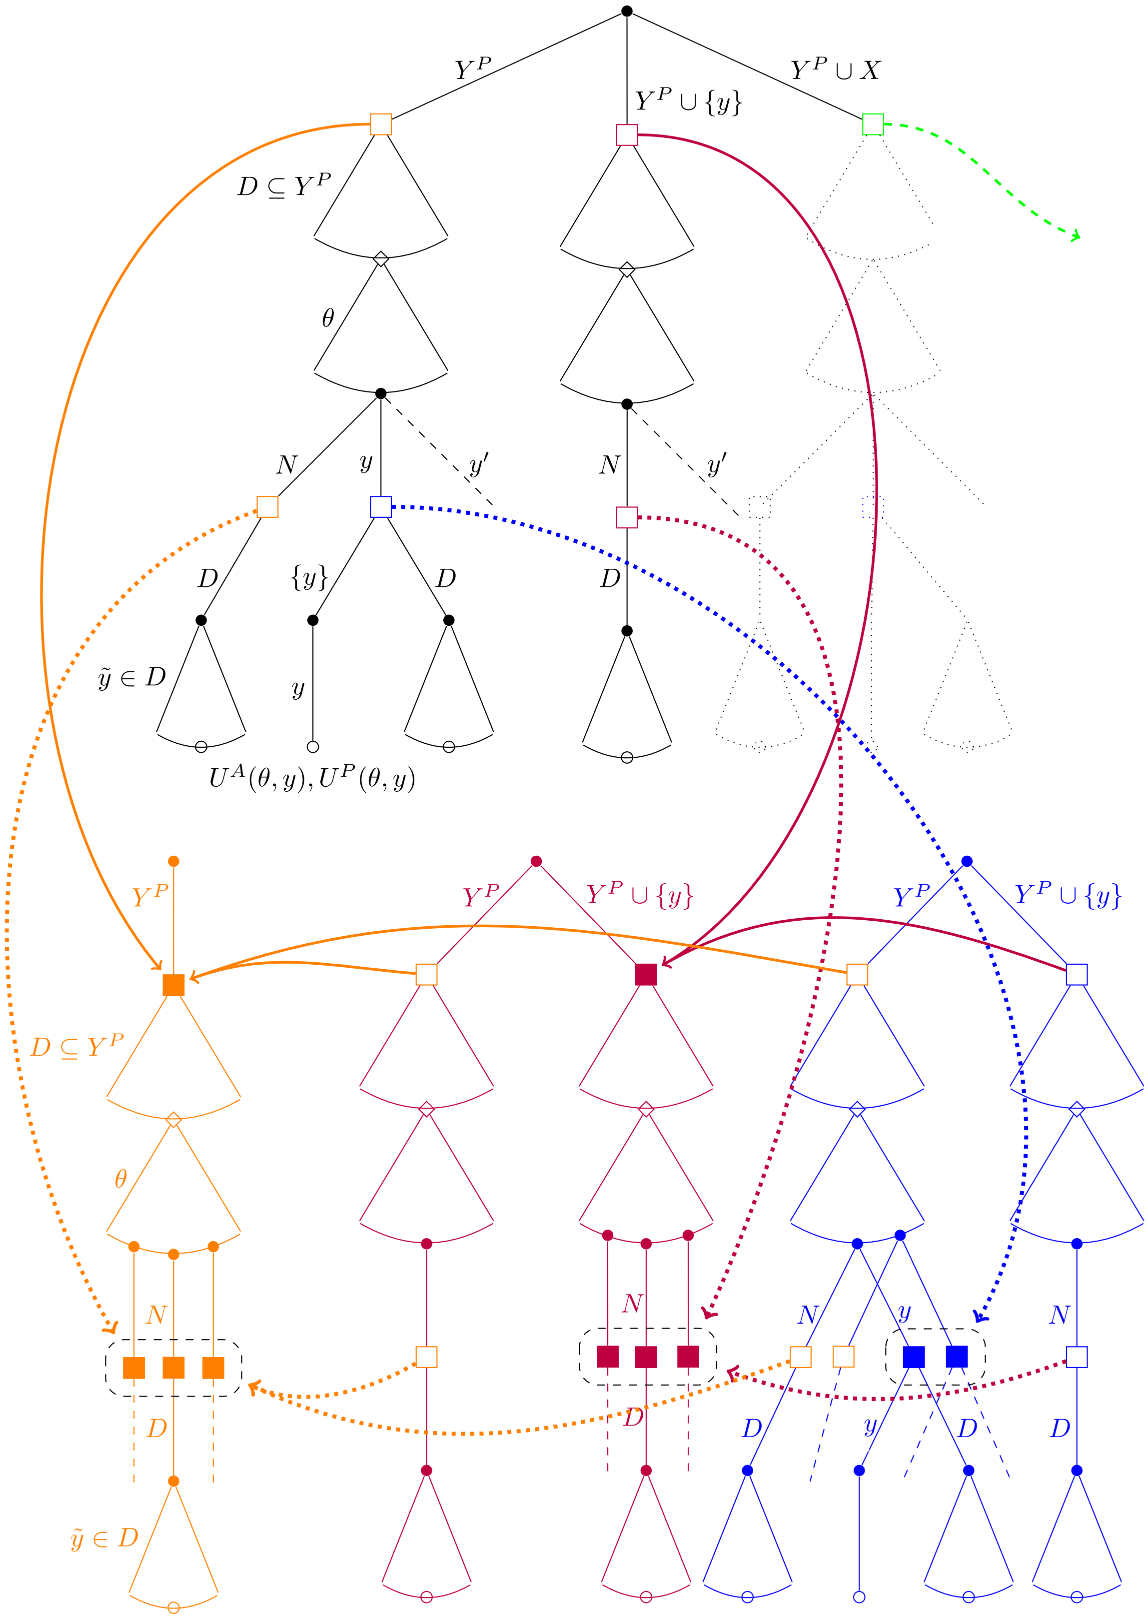

Figure 4: Objective and subjective trees for the game with renegotiation.

where the agent can pick any $y \in D$. The tree T_Y differs from the tree that would represent the view of the principal who has been made aware of Y at the initial stage and does not receive any proposal outside her awareness set at the renegotiation stage, but she is aware of the possibility of renegotiation. The trees representing the latter situation/view are denoted by $T_{Y,\emptyset}$, where the second entry in the subscript indicates that there is no proposal *outside the awareness set at the renegotiation stage*, so the principal's awareness over the set of actions is still Y, although, in this case, she is aware of the possibility or renegotiation. The notation also permits to compute the final awareness of the principal in the tree: in this case, $Y \cup \emptyset = Y$.

Recall further that the agent may only propose singletons at the renegotiation stage. Consider awareness level Y after the agent's move at the root, and a proposal $y \notin Y$ at the renegotiation stage. This leads to a principal's information set that belongs to the game tree denoted as $T_{Y,y}$. In this tree, the principal contemplates the fact that at the beginning of the game the agent could have revealed $\sigma = Y \cup \{y\}$ (followed by N at the renegotiation stage) or any subset of $Y \cup \{y\}$ which is a superset of Y^P. Similarly, the principal realises that the agent could have revealed any other $y' \in Y$ at the renegotiation stage had he started at the root with $\sigma' = Y' \subset Y$, $\{y\} \cap Y' = \emptyset$, and so on. Our notation implies that the tree representing both the modeller's and the agent's view is denoted by $T_{Y^A,\emptyset}$. This tree will, of course, be identical to all other trees where the principal has full awareness and she is also aware of the renegotiation stage. As anticipated, with the exclusion of the trees T_Y, where the principal remains unaware of the possibility of a renegotiation stage, all other trees are identical as long as the principal's final awareness is the same.[28] Despite these similarities, in order to describe the principal's information sets, all such trees must be treated separately. This is because within a generalised game, information sets model both information and awareness.

Information sets and Strategies The forest of trees $\mathbf{T}$ is constituted by trees of the form (in increasing order of awareness): T_Y, $T_{Y,\emptyset}$, and $T_{Y,y'}$, for $Y \in \mathcal{Y}(Y^P)$ and $y' \in Y^A \backslash Y$.

In each tree of the form T_Y, the principal has a single information set after a specific move at the root by the agent: this is the move $\sigma_T = Y$. At these information sets the principal chooses the delegation set $D \subseteq Y$. In T_Y the principal has also one information set after the agent's 'artificial' proposal N at the renegotiation stage but the principal has a singleton move at this information set, so it is less relevant.

In each tree $T = T_{Y,y}$, the principal has information sets at the renegotiation stage alone. To describe the principal's information sets at the renegotiation stage, say, within the tree $T_{Y,y}$, recall that the agent's proposal y at the renegotiation stage occurs after he observes θ, while the principal cannot observe it. Moreover, the principal might have taken a different delegation set after the agent's first move. Given $D \subseteq Y$, a principal's information set in the tree $T_{Y,y}$ is constituted by all nodes of the form (Y, D, θ, y) for $\theta \in [0, 1]$. In other terms, in each one of the trees $T_{Y,y}$, the principal has several information sets, one for each $D \subseteq Y$, and each one of such information sets includes all nodes with the same agent's moves in the renegotiation stage (y in this case) and all

[28]For example, the tree $T_{Y\cup\{y_1\},y_2}$ is identical to the tree $T_{Y\cup\{y_2\},y_2}$, because in both games the principal is aware of the possibility of renegotiation and has the same overall awareness set.

$\theta \in [0,1]$. The *agent has information sets* at each decision node across the different trees, and they are all singletons.

A principal's strategy in the generalised game is described by a move in each information set. Given the specific structure of the principal's information sets across the trees constituting the generalised game, a principal's strategy can be described by a delegation move D_T and a reply function ρ_T for each tree T of the form T_Y. The delegation move is as in the 'pure delegation' case: it assigns a delegation set D given the awareness of the principal at this stage of the tree. As indicated, the function ρ_T for the three T_Y is reported for completeness but it is a singleton choice.

In trees of the type $T_{Y,y}$, there is no information set for D_T while ρ_T takes the particular history (Y, D, y) in the tree and maps into $\{a, r\}$, where a stays for 'accept' and r for 'reject' the proposal.[29] Note that within the tree, only the entry D can vary. This is because each alternative move of the agent $Y' \neq Y$ or $y' \neq y$ would lead to a principal's information set in a different tree.[30] In particular, although the agent is aware of the possibility of renegotiation, the information set corresponding to history (Y, D, N) in the tree $T_{Y,y}$ lies in the tree T_Y.

A complete description of the agent's set of strategies requires the indexation by the tree and the history in each tree. Fix a tree T. Then σ_T represents the agent's move at the root of the tree, x_T the proposal function and y_T the agent's last stage choice. The function x_T takes tuples (σ_T, D, θ) into a proposal $x \in Y^A$, while the function y_T takes a tuple $(\sigma_T, D, \theta, x_T, D')$ to D', where $D' \in \{D, x_T\}$. All such functions must be consistent with being in tree T, which indicates the view of the principal. So, for example, if $T = T_{Y,y}$ then $Y^P \subseteq \sigma_T \subseteq Y$, $x_T \in Y \cup \{y\}$, and of course $D \subseteq \sigma_T$ and $y_T \in D'$.

We now describe what a *T-partial game associated to a given information set* is. There are two main types of T-partial games. First of all, if the principal remains unaware of the possibility or renegotiation, her information sets lie into $T-$partial games constituted by T_Y and all less expressible trees (i.e., $T_{Y'}$ for all $Y' \subset Y$). Again, this is in full analogy to the 'pure' delegation case. We then have T-partial games that are associated (i.e., include) information sets in which the agent is aware of the possibility of renegotiation. If the information set is at node (Y, D, y) in the tree $T_{Y,y}$, the T-partial game associated to this information set is constituted by the following trees: (i) $T_{Y'}$ for all $Y' \subseteq Y \cup \{y\}$; and (ii) $T_{Y',y'}$ where $Y' \cup \{y'\} \subseteq Y \cup \{y\}$, including the tree $T_{Y \cup \{y\}, \emptyset}$.

In order to define a Bayesian equilibrium, we need to define beliefs over nodes at information sets that include more then one node. As described above, the only instance where this happens is at the principal's information sets at the renegotiation stage after the agent's proposal. Given the commonly known cumulate F, and the focus on pure strategies, the principal's beliefs at such information sets are uniquely defined by the support. The support's belief function will be denoted as $\Theta_T(\cdot)$ and, for each tree T it maps triplets (Y, D, x) into subsets of $[0,1]$. As mentioned above, such beliefs are irrelevant for the principal's choice in all trees T_Y. Given that both the agent and the principal have a singleton choice, we set w.l.o.g. $\Theta_T(\cdot) = [0,1]$ for such trees (both on and off

[29] Recall, when $\rho_T = a$ then the delegation set the agent faces is $D' = y$, and if $\rho_T = r$ then $D' = D$.

[30] If $Y' \subset Y$ and $y' \in Y$ these nodes are present in the tree as hypothetical decision nodes for the principal, but they do not constitute principal's information sets.

the equilibrium).

The literature provides several definitions of (Perfect) Bayesian Equilibrium and Sequential Equilibrium for generalised extensive-from games with unawareness of actions. Some of them are defined over slightly different frameworks. This is for example the case of Halpern and Rêgo (2014) and Feinberg (2020).[31] Our signalling game of the renegotiation stage fits into the class of dynamic games studied by Ozbay (2007) in that there is a fully aware agent who moves first.

Definition 2. *A PBE equilibrium for the generalised game is constituted by an 'assessment', that is, a pair of strategies for the agent and the principal -* $\{(\sigma_T^*, x_T^*, y_T^*)\}_{T\in\mathbf{T}}$ *and* $(\{D_{T_Y}^*\}_{Y^P\subseteq Y\subseteq Y^A}, \{\rho_T^*\}_{T\in\mathbf{T}})$ *- together with a beliefs' support system* $\{\Theta_T^*\}_{T\in\mathbf{T}}$ *for the principal at the renegotiation stage, that satisfy the following conditions. (i) At each information set, each player's choice maximises his/her expected utility, given the equilibrium strategies, conditional on his/her awareness, and the associated beliefs' support system. (ii) The strategies and beliefs' support system constitute a PBE in each subtree within any* $T-$*partial game. (iii) The beliefs' support system is consistent with the principal's awareness and equilibrium strategies whenever possible, and the conditional probabilities over each element in* Θ_T^* *are computed using the cumulate* F.

The meaning of the qualification 'conditional on his/her awareness' in condition (i) is that the choices and beliefs at each information set h only consider payoffs and strategies in the T-partial game that include the information set h. Condition (ii) is imposed for consistency with the definition given in Appendix A.5 which requires PBE at each (Y, D).

The next proposition relates the equilibrium set according to the previous definition to the equilibrium set according to the definition provided in Appendix A.5.

Proposition 15. *(i) Let* (σ^*, x^*, y^*), (D^*, ρ^*) *and* Θ^* *be an equilibrium, where* $(\sigma^*, D^*(\cdot))$ *is the solution of the contracting phase given* y^* *and* (x^*, ρ^*, y^*) *together with* Θ^* *constitute PBE equilibria according to Definition 1 in Appendix A.5. Then we can find a PBE equilibrium according to Definition 2 that is payoff equivalent and, in fact, coincides in all 'payoff-relevant' elements.*
(ii) The payoff-relevant components of the agent's best equilibrium outcome are the same across the two definitions.

Proof. (i) Let us start with the function y_T, $T \in \mathbf{T}$ describing the agent's strategy in the very last stage. It is clear that, given (θ, D'), the last stage problem of the agent is identical across all trees. We can hence describe the agent's strategy component at the last stage with functions y that map tuples (Y, D, θ, D') into D'. Given y, each functions y_T can be recovered by appropriately reducing the domain of the function y. The structure of the generalised game allows for different functions in each tree. As we have argued above, once we impose the tie-breaking rule, the agent's final mapping compatible with equilibrium will be a single function. As a consequence, the restriction we impose here is w.l.o.g. as long as we require subgame perfection. In fact, abusing notation, we

[31] For related equilibrium definitions for games with unawareness in normal-form see Meier and Schipper (2014) and Copic and Galeotti (2006).

let y^* be such function and its components that lead to $D' = D$ as it has been defined in the pure delegation model, while for $D' = x$ we obviously have $y^*(\cdot, x) = x$.

Recall that the delegation choice of the principal is a move for each Y, so it can be described—exactly as in the 'pure delegation' case—by a function mapping principal's awareness Y into a delegation set $D \subseteq Y$. Moreover, given the unawareness about the renegotiation stage at all such information sets, the principal's considerations regarding the future at these information sets are summarised by the map y^*, as in the 'pure' delegation case. In our framework, to take into account the unawareness of the renegotiation phase, the functions which the principal sees are those where the last entry is $D' = D$, which in turn is the consequence of the (artificial) 'no proposal' we assumed at the renegotiations stages in all trees T_Y. The tie-breaking rule implies that the equilibrium delegation function is unique and it coincides with the optimal one in the 'pure delegation' model, which we denoted by D^*.

So far we have shown that both equilibrium definitions use the (unique) functions D^* and y^* of the pure delegation model (appropriately adjusted for the extensive form).

We now move to the ρ function component of the principal's strategy. We want to show that using $\rho^*_T = \rho^*$, as defined in Appendix A.5, for all trees where the principal is aware of the renegotiation stage (with appropriate adjustments for the domain) and trivial choices for trees where the principal is unaware of renegotiation, constitutes an equilibrium for the generalised game.

For trees corresponding to the principal's views where she is aware of the renegotiation stage, the ρ_T component of the principal's strategy will be determined by the principal's beliefs at the information set (Y, D, x) of the tree $T_{Y,x}$. Note that all changes in Y and x that lead to changes in principal's awareness lead to information sets in different trees, so the only redundancy comes from the principal's choice D. Since the agent's initial move σ is done before knowing θ, the principal's inference at node (σ_T, D, x_T) only depends on the pair (D, x), her awareness Y, and the agent's strategy at the node. Note that the principal's beliefs and choices off the equilibrium path may matter for the agent's choice at the root. It will be useful to distinguish two type of off-the-equilibrium beliefs. First, whenever the node has the D-component off-the-equilibrium, beliefs and the principal's choices do not affect the agent's choice. On the other hand, for nodes of the form $(Y, D^*(Y), x)$, for each $Y \supseteq Y^P$, we have shown in Proposition 5 that in each PBE at the renegotiation stage where x is accepted by the principal, the beliefs' support is:[32]

$$\Theta(Y, D^*(Y), x) := \{\theta \in [0,1] : U^A(\theta, x) \geq U^A(\theta, y'), \forall y' \in D^*(Y)\}. \tag{33}$$

Note in particular, that this set does not depend on the agent's strategy and that the set of the agent's proposals that can be accepted by the principal in a PBE is again defined by Proposition 5.

We have hence shown that for trees where the principal is aware of the renegotiation stage we can set $\rho^*_T = \rho^*$, where ρ^* is the equilibrium function according to Definition 1 (mapping tuples (Y, D, x) to $D' \in \{D, x\}$); and the principal's belief at node (Y, D, x) can be the set $\Theta^*(Y, D, x)$ of feasible θ, satisfying (33) together with the marginal induced by the cumulate F obtained from Definition

[32]It is easy to see that the inference leading to the set Θ^* holds in this context as well.

1. If faced with the entry (Y, D, N), or in trees where the principal is unaware of the renegotiation stage, the function ρ_T has a singleton choice $\rho_T(\cdot) = D$ and we can set $\Theta^*_T(Y, D, N) = [0, 1]$ at these information sets. Similarly, we can set $x^*_T = x^*$ for all trees where the principal is aware of renegotiation and $x^*_T = N$ for all other trees.

It is immediate to see that the mentioned strategies together with the initial choice $\sigma^*_T = Y^*$ in all trees reached in equilibrium for some realisation of θ constitute a PBE according to Definition 2. The agent's choice at the root of trees which are not reached in equilibrium will also be specified in Defintion 2, while they are not defined in the more compact definition of Appendix A.5.

Point (ii) is a direct consequence of the fact that in each equilibrium according to Definition 2, the set of the agent's proposals that can be accepted by the principal is characterised by Proposition 5. □