Algorithmic Delegation

ALI KHODABAKHSH, University of Texas at Austin
YUANZHE LIU, Oberlin College
EMMANOUIL POUNTOURAKIS, Drexel University
SAMUEL TAGGART, Oberlin College
YICHI ZHANG, Oberlin College

This paper builds a general theory of delegation from an algorithmic perspective. In the delegation problem, an uninformed principal must consult an informed agent to make a decision. Both the agent and principal have preferences over the decided-upon action which vary based on the state of the world, and which may not be aligned. The principal may commit to a mechanism, which maps reports of the agent to actions. When this mechanism is deterministic, it can take the form of a menu of actions, from which the agent simply chooses upon observing the state. In this case, the principal is said to have delegated the choice of action to the agent. We analyze the principal's choice of menu as an algorithmic problem. Rather than study highly parametrized models, as is common in the delegation literature, we study a fully general discrete model of delegation, and show that under minimally restrictive assumptions, simple mechanisms are approximately optimal, while the exact optimal mechanism may be computationally hard to obtain (and therefore necessarily complicated). We derive measures of alignment and conflict between the principal and agent for each action in each state, and show that when alignment is insensitive to the state, it is optimal for the principal to select an action without consulting the agent. When conflict is insensitive to the state, we give tight upper and lower bounds on the approximate optimality of threshold policies. In both cases, our results offer insights into a wide range of common economic scenarios which are robust to modeling assumptions.

1 INTRODUCTION

This paper considers a ubiquitous scenario in economic decisionmaking. A decisionmaker or *principal* faces a choice in which the appropriate action is dependent on the state of the world. For example, a firm seeks to choose a new employee from a pool of applicants with varying traits, or a national health service must choose which treatment to provide to a patient who might display a range of symptoms. For practical reasons, however, the principal is unable to directly observe the state, and must rely on an *agent* to observe the state instead. Firms rely on managers who interview candidates, and the health service relies on doctors to observe patients. In such arrangements, the agent and principal tend not to have preferences which are perfectly aligned: managers may value different traits in an employee than their own supervisor, and doctors tend to overdiagnose certain conditions. How should the uninformed principal interpret information from the agent to manage the misalignment of incentives and choose the appropriate action for the state?

When the principal has commitment power, the choice of a mediating device becomes a problem of mechanism design. The agent observes the state and reports a message to the mechanism, which chooses a possibly randomized action.[1] In this setting, deterministic mechanisms hold special appeal. The *taxation principle* states that every deterministic mechanism is equivalent to a menu: the principal selects the set of allowable actions, and the agent simply chooses their preferred action upon observing the state. Such mechanisms eliminate the need for communication between the agent and principal, and are therefore so common in practice that they are often taken for granted as a managerial tool. The problem of menu design for a better-informed agent is often referred to as *delegation*, coined by [14].

Emails: ali.kh@utexas.edu, yliu8@oberlin.edu, manolis@utexas.edu, sam.taggart@oberlin.edu, yzhang4@oberlin.edu

[1] It is without loss of generality to consider a single round of communication.

Many studies of the delegation problem have focused on the case where the state is single-dimensional and preferences are single-peaked [e.g. 2]. Because of the problem's generality, many well-studied economic models can also be framed as instances of delegation. Some examples include studies of optimal taxation such as [23], tariff design [3], assortment selection [15, 25], and revenue-maximizing pricing [6]. Both types of work on delegation have produced deep theoretical results, which often involve detailed characterizations of optimal mechanisms.

In this work, we take a complementary approach. Rather than examine tailored, parametrized models such as those mentioned, we consider the delegation problem from a broad perspective. We ask: *what are minimal assumptions on the model under which simple mechanisms are optimal or approximately optimal?* In other words, by viewing the problem through the lenses of simplicity and approximation, we hope to obtain insights which are robust to modeling assumptions, and which hold explanatory power even in the case where the spaces of states, actions, and utilities are complex. Some settings in which such complexity is common include:

- **Managerial Decisionmaking.** Hierarchical structure in firms often results in a manager relying on a better-informed subordinate for hiring decisions, implementation choices, and project selection. In all three scenarios, each option may have many traits (e.g. lines on a resumé), which map onto preferences for the manager and subordinate in complicated ways.
- **Regulatory Design.** Regulators frequently wish to manage decisions by firms which are more complicated than the choice of a single price or a quantity to produce. Examples include selecting a location for a new facility or designing the suite of safety features on a new product.
- **Peer-to-Peer Platforms and Crowdsourcing.** On platforms such as Airbnb and Upwork, users are presented with a slate of options - housing and guests in the first case, workers and projects in the second. The preferences of agents on both sides of the market are typically unknown to the platform, and typically differ from those of the market designer, who might value traits such as fairness [see 16] or learning [see 20].

In the above settings, developing plausible, robust models represents a challenge for the theorist. With an algorithmic approach, we seek to circumvent these difficulties and obtain structural insights with both descriptive and prescriptive value.

1.1 Results

Our contributions are as follows.

- We introduce a fully general algorithmic model for delegation. We prove that optimal delegation in the fully general formulation is not only NP-Complete, but hard to approximate within any factor which is sublinear in the number of states or actions.
- To obtain tractable special cases, we present a method for decomposing the utility functions of the agent and principal in a way that quantifies their misalignment in incentives for each pair of state and actions, detailed in Section 3. We then study natural special cases where this alignment measure is restricted in its state-dependence. In each case, we give tight analyses of simple mechanisms. We detail these cases below.
- In Section 4, we study the case where the agent and principal agree on a value for each action which is independent of the state, but in which the agent's preferences undergo a state-dependent distortion. We prove that the optimal mechanism for the principal forgoes consulting the agent entirely, and commits to a single action ex ante. In many applications, this explains the absence of agency (or absence entirely) of informed, subordinate parties in organizations.

- In Section 5 we assume the agent and principal agree on a value for each action which is potentially state-dependent, but restrict the agent's distortion for each action to be independent of the state. This special case is APX-hard, but we show that threshold policies are approximately optimal: the principal restricts the agent to actions for which the distortion is not too large. In many applications of interest, this takes the form of budgets which limit the agent's ability to spend the principal's money. Moreover, we prove our analyses are tight with matching lower bounds.

1.2 Related Work

While many commonly-studied economic decisionmaking scenarios can be framed as delegation, the delegation problem was first articulated as such by [14]. We focus below on work which specifically elaborates on the model of [14], and make no attempt to survey the myriad well-studied areas in mechanism design with informational asymmetries that incidentally place them under the umbrella of delegation.

Much of the delegation literature has focused on the special case where the state and action space are continuous and real-valued, and where the preferences of both the agent and principal are single-peaked or more specifically quadratic loss. Notable examples include [2], which gives a detailed characterization of the optimal delegation policy, and [21], which asks when the agent's selected action is a continuous function of the state. Additionally, [18] considers randomized mechanisms, and derive conditions under which randomized mechanisms outperform deterministic mechanisms.

The cheap talk model of [8] serves as a common benchmark in the delegation literature. [8] analyzes a model with the same informational asymmetry, but a lack of commitment power on the part of the principal. Several papers, including [22] and [9] compare the payoffs of the agent and principal under delegation to their no-commitment payoffs under a single-peaked model of preferences, and derive conditions under which commitment is preferred to non-commitment and vice versa. [11] considers similar questions applied to legislative decisionmaking. We also note [19], which interpolates between zero and full commitment giving the principal the ability to commit to transfers but not actions.

Several papers consider extensions beyond the single-peaked model of preferences. [12] studies an agent and principal who jointly make many identically distributed decisions, and shows that it is beneficial for the principal to adopt a quota mechanism which resembles "grading on a curve." [4] analyzes a model of delegation in which the principal may force the agent to burn money as a condition to taking certain actions. [1] and [24] consider models in which one or more of the principal and the agent may expend effort to observe a signal about the state. These latter three papers fall outside the scope of our model, and may represent fruitful topics for future work from an algorithmic perspective. Finally, [5] considers a model for project selection with an approach that bears loose resemblance to ours. It models available projects (i.e. actions) as stochastic. The agent observes the set of realized available projects, and selects their preferred project, subject to restrictions imposed ex ante by the principal. The utilities of the agent and principal for each project are themselves random, and therefore the problem falls outside the single-peaked model.

Lastly, we note two recent papers with approaches similar in spirit to the present work. [17] studies the model of [5], and draw a connection to the prophet and Pandora's box problems, two canonical models in optimal stopping theory. They use this similarity to show that threshold policies are approximately optimal for the principal. While this paper's results are intriguingly similar to ours, model differences render them incomparable. Second, [10] studies a classical agent-principal model of contracting with unobservable actions from the perspective of approximation and simplicity, and prove that linear contracts are approximately optimal.

2 MODEL

We now outline our model of delegation. We consider a principal who seeks to choose an action from a discrete set Ω of m actions. The principal's utility is joint function of the action taken and the state of the world, which the principal does not observe. We assume the set of states S is finite, with size n. The principal may consult an agent, who observes the state, and may communicate with the principal after observation. The agent also possesses a utility function over (action, state) pairs, and this utility function may differ from that of the principal. We denote the utilities of the principal and agent by U_{P} and U_{A}, respectively, represented as m-by-n matrices. For a matrix M, we use either $m_{i,j}$ or $M(i,j)$ to denote its entry at row i and column j. For example, $U_{\mathrm{P}}(i,j)$ measures the utility of the principal if the agent picks action i in state j. To enable the study of approximation algorithms, we assume the entries of U_{P} are nonnegative.

We assume the principal has the power to commit ex ante to a mechanism for communicating with the agent and selecting an action. A mechanism is a function mapping a report of the agent (after observing the state) to a (possibly randomized) action selected by the principal. We consider deterministic mechanisms for this problem. By the taxation principle, deterministic mechanisms may be represented as menus over the set of actions. The agent observes the state and selects their preferred action from the menu. The principal then takes the selected action. Such mechanisms have the advantage that they are simple to implement, and in many settings require no communication between the agent and principal. Taking this perspective, we consider the algorithmic problem of selecting a menu A to maximize the principal's expected utility when the agent selects their preferred action according to the observed state.

Formally, when presented with action set $A \subseteq \Omega$ and after observing that the state is s, denote the agent's preferred choice by $g(A,s)$. That is, $g(A,s) = \mathrm{argmax}_{a\in A} U_{\mathrm{A}}(a,s)$. The principal is faced with a set function optimization problem. We assume the principal has a prior distribution over the state, with state s occurring with probability p_s. The principal must select a menu A for the agent which maximizes their own expected utility over the state. That is, the principal solves:

$$\underset{A\subseteq\Omega}{\text{maximize}} \;\; f(A) := \sum_{s\in S} p_s \cdot U_{\mathrm{P}}(g(A,s),s)$$

Example 2.1. While the model above is extremely general, we instantiate it with a stylized model of a firm seeking to buy a new piece of equipment. In the example, the firm acts as principal, and must decide which of three possible items, A, B, and C to consider, with fixed prices p_A, p_B, and p_C. The firm's management must rely on the specialist who is to use the equipment for recommendations as to quality. This operator serves the role of the informed agent. Hence the firm delegates the decision to the agent by offering a nonempty subset of the items, from which the operator will make a recommendation. Both parties benefit from the equipment being well-suited to its task, but only the firm internalizes the monetary costs, leading to a misalignment of incentives.

We consider three states corresponding to three different profiles of needs to which each piece of equipment is more or less suited, which we can represent by the following matrix, which we populate with values for concreteness.

	1	2	3
A	2	5	6
B	4	6	7
C	7	4	6

The first column corresponds to state of the world in which equipment C is most suitable, while in the other states, B performs best. Instantiating the items' prices as $p_A = 1$, $p_B = 3$, and $p_C = 4$, we produce the following utility matrix for the principal, whose utility is value minus price (U_P, at left):

$$U_\text{P} = \begin{bmatrix} 1 & 4 & 5 \\ 1 & 3 & 4 \\ 3 & 0 & 2 \end{bmatrix}, \quad U_\text{A} = \begin{bmatrix} 2 & 5 & 6 \\ 4 & 6 & 7 \\ 7 & 4 & 6 \end{bmatrix}.$$

That is, purchasing item A under state 1 gives the principal utility $2 - 1 = 1$, and the other entries of the matrix follow similarly. Since the operator does not pay for the equipment, his utility matrix is simply determined by the value of equipment in each state. The agent's utility matrix is at right, above (U_A). Note that because the prices change the principal's utility for the items, the principal's and agent's incentives are misaligned for example under state 2 - the operator prefers to buy item B, while the principal wants to buy cheaper equipment A.

Under the state indexed j, the operator picks their most desired item $g(X, j)$ from the set X of available items. For example, if X is the set of all items, we would have:

$$g(X, j) = \begin{cases} \text{C} & j = 1, \\ \text{B} & j = 2, \\ \text{B} & j = 3. \end{cases}$$

The principal could decide not to make all the items available, and instead offer $Y = \{A, C\}$, for example. In this case, the operator would prefer item C under job type 1, and would prefer item A under the second and third states. Assuming a uniform distribution over states, the expected utility of the principal is maximized by offering Y, and is calculated as:

$$f(Y) = \frac{1}{3} \sum_{j \in \{1,2,3\}} U_\text{P}(g(Y, j), j) = \frac{1}{3} U_\text{P}(\text{C}, 1) + \frac{1}{3} U_\text{P}(\text{A}, 2) + \frac{1}{3} U_\text{P}(\text{A}, 3) = \frac{3 + 4 + 5}{3} = 4.$$

To preview our results, we show that in settings such as the one described above, it is always approximately optimal for the firm to impose a budget on the agent. In this example, there are two reasonable choices for optimal budget: 1 and 4, both of which give the firm expected utility of 10/3.

In Appendix A.3, we give another example which models a firm which must select a project for a worker to undertake. The worker observes and internalizes the project-specific labor costs, whereas the firm does not. For settings such as this one, we will show that the optimal mechanism always selects a single action, which is imposed on the agent.

Computational Hardness of the General Case. To conclude our preliminary discussion of the model, we briefly analyze the complexity of the fully general version of the problem discussed above. Unfortunately, as stated, the problem is computationally intractable in a very strong sense, stated below.

Theorem 2.2. *The delegation problem is NP-hard and also hard to approximate within any factor less than* $\min\{m, n\}$.

A proof can be found in Appendix A.1. The reduction used to prove Theorem 2.2 rules out approximation based on most interesting parameters: the reduction has a uniform distribution over states, the principal's utility is binary (zero or one), and the agent's utility is limited to three different values. Hence, it is necessary to place restrictions on the problem to permit approximation. We discuss one natural way of doing so in the next section.

3 UTILITY DECOMPOSITION

In light of the hardness result of Theorem 2.2, we turn to special cases of the delegation problem. To identify special cases which are still sufficiently broad to apply to the motivating settings mentioned in Section 1, we decompose the principal's and agent's utility matrices to isolate the portion of their utility which is mutual, and the portion which is derived from player-specific distortion of utilities. Formally:

Definition 3.1. Given an instance (U_P, U_A) of the delegation problem, a *utility decomposition* for (U_P, U_A) is given by constants a_1, a_2, a_3, and a_4 and matrices V and D such that $U_P = a_1 V - a_2 D$ and $U_A = a_3 V + a_4 D$. We refer to a_1, a_2, a_3, and a_4 as the *decomposition coefficients* of $(a_1, a_2, a_3, a_4, V, D)$, V as the *value matrix*, and D as the *distortion matrix*.

For any decomposition coefficients (as long as $a_1 a_4 + a_2 a_3 \neq 0$) we may obtain a corresponding utility decomposition by setting

$$V = (a_1 a_4 + a_2 a_3)^{-1}(a_4 U_P + a_2 U_A)$$
$$D = (a_1 a_4 + a_2 a_3)^{-1}(-a_3 U_P + a_1 U_A).$$

Many settings, however, naturally admit particular utility decompositions in which the value or distortion matrix exhibit considerable structure. In the example in Appendix A.3, a firm (the principal) must select one of several projects for a team (the agent) to undertake. While both the team and firm benefit from undertaking a project, only the team observes and internalizes the project-specific costs. Hence, the benefits of the projects are a natural choice for the value matrix V, and the state-specific effort costs paid by the agent serve as the distortion D, yielding decomposition coefficients $a_1 = a_3 = 1$, $a_2 = 0$, and $a_4 = -1$. Note that because the value of each project is known in advance, all columns of V are identical. This will often be the case when the source of distortion are effort costs that depend on information which is unknown to the principal, but where benefits are known to all parties in advance. We refer to this condition as *restricted value*. Formally:

Definition 3.2. An instance of delegation with utility decomposition $(a_1, a_2, a_3, a_4, V, D)$ satisfies *restricted value* if all columns of V are identical.

In procurement setting of Example 2.1, the benefits to the firm from new equipment are a natural choice for the value matrix V, in which case the monetary costs paid by the firm serve as distortion D, yielding decomposition coefficients of $a_1 = a_2 = a_3 = 1$ and $a_4 = 0$. Note that in this case, the costs are state-invariant; that is, all columns of D are identical. This will often be the case in settings with transfers that are known in advance. We refer to this condition as *restricted distortion*. Formally:

Definition 3.3. An instance of delegation with utility decomposition $(a_1, a_2, a_3, a_4, V, D)$ satisfies *restricted distortion* if all columns of D are identical.

The restricted distortion and restricted value conditions still capture a variety of settings of interest (and in particular those for which existing models are ill-suited). In what follows, we study the effectiveness of simple delegation policies under each condition. In Section 4, we show that under the restricted value condition, an extremely simple policy is always optimal: pick the best single action and force the agent to take it. Therefore, under the restricted value condition, a principal may not choose to delegate authority in the first place, and may do away with the informed agent altogether if possible. Hence, what appears as a single-agent problem may only appear so because of incentive problems the principal seeks to avoid.

In Section 5, we consider the restricted distortion condition, and show that no simple optimal strategy generally exists. In particular, we show that the delegation problem under restricted distortion is APX-hard. Unlike in the general case, however, we show that there are simple strategies that are *approximately optimal* for the principal. Since the distortion of each action is state-independent, the principal may sort actions by their distortion, and restrict the agent to actions with low distortion, i.e. below some threshold. In the procurement example, this takes the form of a budget imposed on the agent. We show that the best such threshold policy is a $O(\log p_{\min}^{-1})$-approximation, where $p_{\min}$ is the probability of the least likely state. Hence, the performance of simple policies degrades smoothly with the degree of uncertainty faced by the principal. We also show in Section 6 that when distortions are nonpositive and agent utilities are nonnegative, this improves to a 2-approximation.

On Decomposition Coefficients. For the remainder of the paper, we assume $a_1 = a_2 = a_3 = a_4 = 1$, and assume restricted value or distortion hold with respect to these coefficients. We do so only for simplicity of exposition. All results extend straightforwardly to any other nontrivial choice of decomposition coefficients. Consequently, the restricted value result of Theorem 4.4 holds for the project selection setting instantiated in Appendix A.3 and the results of Section 5 hold for the procurement setting of Example 2.1.

4 RESTRICTED VALUE

In this section, we study the restricted value case. We assume the decomposition coefficients are $a_1 = a_2 = a_3 = a_4 = 1$. That is, the utilities can be decomposed as $U_P = V - D$ and $U_A = V + D$, and all columns of V are identical. As discussed at the end of the previous section, our analysis extends straightforwardly to other choices of decomposition coefficients. Since each column of V is identical, for the rest of this section we let v_i denote the value which populates the i^{th} row of V, i.e. the value associated with action i.

Our main result under the restricted value assumption is that the optimal mechanism for the principal is to choose the action that maximizes their expected utility ex ante, and impose this action on the agent as the singleton delegation set. To prove this result, we derive an intuitive consequence of the restricted value condition. Because the state only affects the utilities of the agent and principal in a way that causes their preferences to diverge, it follows that any time the agent would change their decision based on the state, the principal would prefer the opposite. Moreover, if the agent was restricted in their strategy to choose each action with a certain frequency, they would allocate their action "quotas" in a way that exactly minimizes the principal's utility over all strategies satisfying the frequency restriction. We formalize this notion as follows:

Definition 4.1. Given a delegation set A, the *frequency* of an action a under A, denoted q_a, is the probability the agent takes action a, given by $\sum_{s \in S} p_s \mathbb{1}[g(A, s) = a]$.

Definition 4.2. Given a delegation set A and action frequencies $(q_1, \ldots, q_m)$ for A, an *action reassignment* is an $m \times n$ matrix R such that:

(1) For all $a \in \Omega$ and $s \in S$, $r_{as} \geq 0$.
(2) For all $s \in S$, $\sum_{a \in \Omega} r_{as} = 1$.
(3) For all $a \in \Omega$, $\sum_{s \in S} p_s r_{as} = q_a$.

Action reassignments represent feasible solutions to the following thought experiment: given a delegation set A, allow the principal to observe the state and assign actions in A to states in a probabilistic fashion, subject to the constraint that the frequency of each action is the same as if the agent was choosing the action for each state. Because the only state-dependent part of the agent's utility is exactly misaligned with that of the principal, any reassignment of actions to states is preferable for the principal to the assignment selected by the agent. This can be interpreted as an "anti-alignment" condition, and functions as an obverse to the alignment condition discussed in e.g. [12]. Formally:

Lemma 4.3. *If utilities satisfy the restricted value condition, then for any delegation set A and any action reassignment R, the principal's utility satisfies:*

$$\sum_{s\in S} p_s U_P(g(A,s),s) \leq \sum_{s\in S}\sum_{a\in\Omega} p_s r_{sa} U_P(a,s)$$

Proof. In state s, faced with action set A, the agent chooses action $g(A,s)$. An alternate, randomized strategy the agent could employ would be to choose action a in state s with probability r_{as}, for every state s. This deviation is suboptimal for the agent, so we have:

$$\sum_{s\in S} p_s v_{g(A,s)} + \sum_{s\in S} p_s D(g(A,s),s) \geq \sum_{s\in S} p_s \sum_{a\in\Omega} r_{as} v_a + \sum_{s\in S} p_s \sum_{a\in\Omega} r_{as} D(a,s). \tag{1}$$

Note that the first term on each side of (1) is only a function of the number of times each action is taken under the agent's optimal strategy. But note that by condition (3) of Definition 4.2, action reassignments preserve these frequencies. Hence,

$$\sum_{s\in S} p_s v_{g(A,s)} = \sum_{a\in\Omega} q_a v_a = \sum_{s\in S} p_s \sum_{a\in\Omega} r_{as} v_a \tag{2}$$

It follows that the agent prefers taking $g(A,s)$ in state s to $g(A,\pi(s))$ solely because of the component of their utility which is misaligned with the principal's, i.e.:

$$\sum_{s\in S} p_s D(g(A,s),s) \geq \sum_{s\in S} p_s \sum_{a\in\Omega} r_{as} D(a,s). \tag{3}$$

Combining (2) and (3) yields the following inequality, which is equivalent to the lemma:

$$\sum_{s\in S} p_s v_{g(A,s)} - \sum_{s\in S} p_s D(g(A,s),s) \geq \sum_{s\in S} p_s \sum_{a\in\Omega} r_{as} v_a - \sum_{s\in S} p_s \sum_{a\in\Omega} r_{as} D(a,s). \tag{4}$$

□

Theorem 4.4. *If utilities satisfy the restricted value condition, then there is an optimal delegation set which is a singleton.*

Proof. Let A^* be an optimal delegation set, and let $(q_1^*, \ldots q_m^*)$ denote the action frequencies under A^*. It is a feasible action reassignment to take action i with probability q_a^*, regardless of the state. That is, define an action reassignment R^* such that $r_{as}^* = q_a^*$ for all $s \in S$. By Lemma 4.3, we have:

$$\begin{aligned}
\sum_{s\in S} p_s U_{\mathrm{P}}(g(A^*,s),s) &\leq \sum_{s\in S} p_s \sum_{a\in\Omega} r_{as}^* U_{\mathrm{P}}(a,s) \\
&= \sum_{s\in S} p_s \sum_{a\in\Omega} q_a^* U_{\mathrm{P}}(a,s) \\
&= \sum_{a\in\Omega} q_a^* \sum_{s\in S} p_s U_{\mathrm{P}}(a,s)
\end{aligned}$$

The last line above can be interpreted as the expected utility from choosing an action a with probability q_a^* and playing this action in every state. Since this is at least the principal's utility from allowing the agent to choose from the set A^*, it follows that there must be an action a^* such that playing a^* in every state yields at least as much utility as allowing the agent to choose from A^*. This proves the theorem. □

5 RESTRICTED DISTORTION

We now turn our attention to the restricted distortion assumption. That is, we assume the columns of the distortion matrix D are identical.[2] For notational convenience, for a given action a we will denote by d_a the value populating the ath row of the distortion matrix D, and we will refer to d_a as the *distortion of action a*. For a given state i, we will denote by p_i the probability of the state i got picked.

Given the previous section's main result, that there is always a simple optimal mechanism under restricted values, one may expect the same to be true for restricted distortion. Our first result is that this is not the case. This version of the problem is not just NP-hard, but APX-hard, ruling out constant approximations below a certain factor.

THEOREM 5.1. *If utilities satisfy restricted distortion condition, the principal's delegation problem is APX-hard, even under uniformly distributed states.*

PROOF. See Appendix A.4. □

While Theorem 5.1 rules out mechanisms which are both simple and optimal, it does not rule out mechanisms which are simple and *approximately optimal*. Our second result of this section is that such mechanisms generally exist. Under the restricted distortion condition, it is possible to rank actions based on their distortion. The most natural, simple class of delegation mechanisms to analyze, then, is *threshold policies*, which include actions with distortion below a certain level, and exclude those above. Since there are polynomially many such thresholds, analyzing the performance of best threshold is of special interest. We formalize this in Algorithm 1.

In the rest of this section, we prove tight approximation guarantees for the best threshold of Algorithm 1, and show that it is a $\Theta(\log p_{\min}^{-1})$-approximation, where $p_{\min}$ denotes the probability of the least likely state. Hence, threshold policies perform well under low levels of uncertainty, and their performance gradually degrades as the uncertainty grows more extreme. We state our results formally below.

THEOREM 5.2. *Algorithm 1 is a $2\log(p_{min}^{-1})$-approximation under restricted distortion, where p_{min} is the probability of the least likely state.*

THEOREM 5.3. *There exists a family of instances where no threshold policy is better than a $\Omega(\log p_{min}^{-1})$-approximation.*

To prove Theorem 5.2, we first argue in Section 5.1 for the general case where a known outside option is introduced. A known outside option is a special action which is always available to the agent. In other words, the principle must include the outside option in the action set. We subsequently show how a restricted distortion case without outside option can be turned into one with known outside option. A family of examples proving Theorem 5.3 can be found in Appendix A.5.

5.1 Upper Bound for Known Outside Option

We begin our proof of the our upper bound by considering known outside option. We derive the following:

THEOREM 5.4. *Algorithm 1 is a $2\log(p_{min}^{-1})$-approximation algorithm*[3] *under the restricted distortion assumption with known outside option.*

[2]As in the previous section, we argue only for the decomposition with $a_1 = a_2 = a_3 = a_4 = 1$, though our results hold for all other decompostions as well. See Section 3.

[3]Logarithms are with respect to base 2.

ALGORITHM 1: Restricted Distortion Delegation

Input: Set of actions Ω, set of states S, value (V) and distortion (D) matrices.
Output: Set of actions A that approximately maximizes $f(A)$.
$U_{\mathrm{P}} = V - D$, $U_{\mathrm{A}} = V + D$.
for *each* $k \in \{1, 2, \dots, m\}$ **do**
  $threshold = d_k$.
  Set $A_k \longleftarrow \{i \in \Omega : d_i \leq threshold\}$.
end
Let $\mathcal{A} := \{A_1, \dots, A_m\}$.
Return $\mathrm{argmax}_{A \in \mathcal{A}} f(A)$.

The proof of Theorem 5.4 relies on partitioning the set of states into two subsets of high and low distortion. Let OPT be the optimal solution and $d_{\max} = \max_{i \in OPT} d_i$ be the highest distortion in OPT.[4] We will ignore actions with higher distortion than $d_{\max}$, as we may achieve the desired approximation guarantee while doing so.

Let $\Delta > 0$ be a constant to be selected shortly. We say that a state is *high-distortion* (HD) if OPT chooses an action with distortion $d \in [d_{\max} - \Delta, d_{\max}]$ in that state, and otherwise the state is called *low-distortion* (LD). These definitions give a partition of the states into two sets S^{HD} and S^{LD} with total probability P_1 and P_2 such that $S = S^{\mathrm{HD}} \cup S^{\mathrm{LD}}$, and $P_1 + P_2 = 1$. As P_1 and P_2 are increasing and decreasing functions of Δ, respectively, there exists a Δ such that $P_1 = P_2 = 1/2$.[5] Choose Δ such that this holds.

Let OPT^{HD} and OPT^{LD} be the utility of the principal from each partition, i.e.,

$$OPT^{\mathrm{HD}} = \sum_{i \in S^{\mathrm{HD}}} p_i(v_{a_i,i} - d_{a_i}), \qquad OPT^{\mathrm{LD}} = \sum_{i \in S^{\mathrm{LD}}} p_i(v_{a_i,i} - d_{a_i}), \tag{5}$$

where a_i is the index of the action that OPT picks in state i, and clearly $OPT = OPT^{\mathrm{HD}} + OPT^{\mathrm{LD}}$. The following lemma states that the portion of principal's optimal utility (OPT) coming from high-distortion states can be approximated up to a constant factor, by picking the better of two threshold solutions.

Lemma 5.5. *Let $APX_1 = \{i \in \Omega : d_i \leq d_{max}\}$ and $APX_2 = \{i \in \Omega : d_i \leq d_{max} - \Delta\}$, i.e., sets of actions that has distortion no more than d_{max} and $d_{max} - \Delta$, respectively.*[6] *Then:*

$$\max\{f(APX_1), f(APX_2)\} \geq \frac{1}{2} OPT^{HD}.$$

Proof. Let b_i denote the action picked by the approximate solution (which can be either APX_1 or APX_2) in state i. Considering APX_1, in each state $i \in S$ we have

$$v_{b_i,i} + d_{b_i} \geq v_{a_i,i} + d_{a_i},$$

[4] Interchangeably, we use OPT (APX) to denote both the optimal (approximate) set of actions that the principal should allow, and also the optimal (approximate) utility that he achieves by doing so.
[5] In the case that $P_1 \neq P_2$, the arguments hold by choosing Δ such that $P_1 < 1/2$, $P_2 > 1/2$ and putting the highest distortion action in S^{LD} to S^{HD} will result in $P_1 > 1/2$ and $P_2 < 1/2$. For simplicity of exposition, we assume that $P_1 = P_2$ exists. If there is no value of Δ where in $n_1 + n_2 = n$, then there is a jump in n_1 at some $\Delta = \Delta_0$. Then there are k states in which the action picked by the agent has distortion equal to $d_{\max} - \Delta_0$. In that case, we can pick Δ_0 as our distortion margin (Δ) and divide those k states between S^{HD} and S^{LD} such that they have the desired sizes.
[6] Note that $d_{\max}$ and Δ require knowledge of the optimal solution. They are being used solely for analysis, rather than as part of our algorithm.

since the agent decided to choose b_i, while a_i was also allowed in APX_1. This implies the following helpful inequality which relates the loss in principal's utility to the distortion difference of chosen actions:

$$v_{b_i,i} - d_{b_i} \geq v_{a_i,i} - d_{a_i} + 2(d_{a_i} - d_{b_i}). \tag{6}$$

Equation (6) has the following interpretations: First, the principal gains (in a particular state i) if the agent picks an action with lower distortion in the approximate solution ($d_{b_i} \leq d_{a_i}$). Second, in case the agent picks an action with higher distortion than that of the optimal solution, the loss in principal's utility is bounded by twice the distortion difference. However, this difference is bounded by Δ in the high-distortion states, because for $i \in S^{\mathrm{HD}}$ we have $d_{a_i} \geq d_{\max} - \Delta \geq d_{b_i} - \Delta$, which implies

$$d_{b_i} - d_{a_i} \leq \Delta. \tag{7}$$

Using these two inequalities we can get the following lower bound on APX_1:

$$\begin{aligned} f(APX_1) = \sum_{i \in S} p_i(v_{b_i,i} - d_{b_i}) &\geq \sum_{i \in S^{\mathrm{HD}}} p_i(v_{b_i,i} - d_{b_i}) && \text{(8a)} \\ &\geq \sum_{i \in S^{\mathrm{HD}}} p_i(v_{a_i,i} - d_{a_i} + 2(d_{a_i} - d_{b_i})) && \text{(8b)} \\ &\geq \sum_{i \in S^{\mathrm{HD}}} p_i(v_{a_i,i} - d_{a_i} - 2\Delta) = OPT^{\mathrm{HD}} - 2P_1\Delta = OPT^{\mathrm{HD}} - \Delta, && \text{(8c)} \end{aligned}$$

where in the first inequality we just restricted the summation to the set of high-distortion actions, in (8b) we used inequality (6), and in (8c) we used inequality (7) and the fact that $P_1 = 1/2$.

Considering APX_2, we still have inequality (6) for low-distortion states, because by the definition of low-distortion state, action a_i has a distortion no more than $d_{\max} - \Delta$ and therefore it is allowed in APX_2. However, we need a new bound on the distortion differences. Note that by the positivity assumption of principal's utility, the utility of the agent from picking the highest-distortion action in OPT is at least $2d_{\max}$ (for all $i \in S$), hence any action that is picked in any state i has to satisfy at least

$$v_{a_i,i} + d_{a_i} \geq 2d_{\max},$$

which can be rearranged as

$$d_{\max} - d_{a_i} \leq \frac{v_{a_i,i} - d_{a_i}}{2}.$$

Since we allow actions up to the distortion of $d_{\max} - \Delta$ with an exception–the outside option, we get this bound on the distortion difference if we ignore the outside option:

$$d_{b_i} - d_{a_i} \leq d_{\max} - \Delta - d_{a_i} \leq \frac{v_{a_i,i} - d_{a_i}}{2} - \Delta. \tag{9}$$

Then if we include the outside option, there are two circumstances–the outside option is in S^{LD} or in S^{HD}. If outside option is in S^{LD}, the above statement holds. If the outside option is in S^{HD}, for $i \in S^{\mathrm{LD}}$, the agent won't pick the outside option. Therefore, the above statement holds in either circumstances. Now we can find a lower bound on APX_2.

$$\begin{aligned} f(APX_2) = \sum_{i \in S} p_i(v_{b_i,i} - d_{b_i}) &\geq \sum_{i \in S^{\mathrm{LD}}} p_i(v_{b_i,i} - d_{b_i}) && \text{(10a)} \\ &\geq \sum_{i \in S^{\mathrm{LD}}} p_i(v_{a_i,i} - d_{a_i} + 2(d_{a_i} - d_{b_i})) && \text{(10b)} \\ &\geq \sum_{i \in S^{\mathrm{LD}}} p_i \left(v_{a_i,i} - d_{a_i} - 2\Big(\frac{v_{a_i,i} - d_{a_i}}{2} - \Delta\Big) \right) = 2P_2\Delta = \Delta, && \text{(10c)} \end{aligned}$$

where this time we restricted the summation to the set of low-distortion actions, in (10b) we used inequality (6), and in (10c) we used inequality (9) and the fact that $P_2 = 1/2$.

We finish the proof by showing that the better solution among APX_1 and APX_2 is a 2-approximation for OPT^{HD}.

$$\max\{f(APX_1), f(APX_2)\} \geq \frac{1}{2}\left(f(APX_1) + f(APX_2)\right) \geq \frac{1}{2}(OPT^{\text{HD}} - \Delta + \Delta) = \frac{1}{2}OPT^{\text{HD}}.$$

□

Now we are ready to analyze the performance of Algorithm 1. We use an inductive argument to show that the previous 2-approximation of OPT^{HD} implies an $O(\log n)$-approximation for OPT.

Proof of Theorem 5.4. We prove by induction (on the number of states) that there exists a $2\log n$-approximation algorithm for OPT. Intuitively, if OPT^{HD} is the dominant part of OPT, then $\max\{f(APX_1), f(APX_2)\}$ from Lemma 5.5 is a good approximation for OPT too. On the other hand if OPT^{LD} is dominant, we can ignore all the high-distortion states without losing much, and this reduces the number of states by half.

For the base case, we show that there exists a 2-approximation for $n = 2$ states. Clearly, the expected principal's utility is the average of the utilities in the two states, and therefore one of these states has a utility of at least OPT. Let $i^* \in \{1, 2\}$ be that particular state, and a_{i^*} be the corresponding action in the optimal solution. The solution that consists of a_{i^*} and all actions with lower distortion will be a 2-approximation. It is guaranteed to generate at least OPT utility in state i^* which happens with probability half, because it cannot go to higher distortion levels and its utility will be at least equal to what OPT provides, based on (6). Now assume that we have a $2\log\frac{n}{2}$-approximation for any problem with $n/2$ states. For the problem with n states, one of the following happens:

Case I (Dominant OPT^{HD}). If we have $OPT^{\text{HD}} \geq \frac{1}{\log(p_{min}^{-1})}OPT$, then

$$\max\{f(APX_1), f(APX_2)\} \geq \frac{1}{2}OPT^{\text{HD}} \geq \frac{1}{2\log(p_{min}^{-1})}OPT,$$

which is the desired approximation ratio.

Case II (Dominant OPT^{LD}). In this case we have the opposite assumption $OPT^{\text{HD}} < \frac{1}{\log(p_{min}^{-1})}OPT$, which implies:

$$OPT^{\text{LD}} \geq \left(1 - \frac{1}{\log(p_{min}^{-1})}\right)OPT = \frac{\log(p_{min}^{-1})/2}{\log(p_{min}^{-1})}OPT.$$

Now using the $2\log\frac{(p_{min}^{-1})}{2}$-approximation (inductive assumption) on the set S^{LD} will give a solution APX^{LD} such that

$$f(APX^{\text{LD}}) \geq \frac{1}{2\log(p_{min}^{-1})/2}OPT^{\text{LD}}.$$

Combining the two inequalities we get:

$$f(APX^{\text{LD}}) \geq \frac{1}{2\log(p_{min}^{-1})/2} \times \frac{\log p_{min}^{-1}/2}{\log(p_{min}^{-1})}OPT = \frac{1}{2\log(p_{min}^{-1})}OPT,$$

which again gives the desired approximation ratio.

Finally, note that even though $d_{\max}$, Δ, S^{LD}, and S^{HD} in the proof are all defined based on the optimal solution and hence are not known, all the proposed approximate solutions (i.e., APX_1, APX_2, or the base case solution) have a simple structure. There exists a distortion threshold such that every action with lower distortion is allowed, and everything else is excluded. Therefore, the approximate delegation solution can be found by the linear-time search of Algorithm 1. This proves the theorem. □

Matching lower bound. We conclude by emphasizing that the above analysis is tight. We show in Appendix A.5 that our analysis in Theorem 5.4 is tight up to a constant factor. We do so by providing instances where no threshold policy can outperform the logarithmic approximation ratio.

5.2 Extension To Unknown Outside Option

In reality, principles do not have enough information about the outside options, so this leads to the idea of unknown outside option. In unknown outside option case, instead of knowing the exact details of the outside option, the principle only know a range of the outside option's distortion. Then the principle wants to maximize its minimum utility from picking the worst outside option.

THEOREM 5.6. *Restricted distortion with unknown outside option is equivalent to known outside option case and thus have* $2\log P_{min}^{-1}-$*approximation rate*

To prove theorem 5.6, we first observe that an outside option can be make 'worse' for the principle by making the principle utility smaller, i.e., with the small distortion, an outside option with all zero principle utility is worse than one with all non-zero principle utility. Then to make it even worse, we want the agent to actually pick this bad outside option, so we let the agent utility to be as large as possible. We can choose the largest distortion such that the agent utility is maximized while the principle utility is zero. Then this outside option is always the worst for principle. Then we have showed the unknown outside option case is equivalent to known outside option.

5.3 Extension To Non-Uniform States

We now extend the approximation result for restricted distortion from uniformly distributed to arbitrarily distributed states. We do so via a reduction that loses a factor of 2. For uniformly distributed states, the approximation ratio depended on the number of states. For non-uniform states, the natural analog is the inverse of the probability of the least likely state.

Let U_P, U_A be an instance of the delegation with states distributed according to $(p_1, \ldots, p_n)$. Let $p_{\min} = \min_{s\in S} p_s$ be the probability of the least likely state. We will construct a new instance of the problem with the same actions, a new set of uniformly distributed states, and where the principal's utility from any delegation set is within a constant factor in the two instances. The number of states in the new instance will be at most $p_{\min}^{-1}$. Furthermore, if the original instance satisfied restricted distortion, so too will the new instance.

The rough approach to the reduction will be rounding each state's probability down to the nearest factor of $p_{\min}$ (and normalizing). Doing so will distort the utility from any delegation set by at most a factor of 2. We may then divide each state evenly into many identical copies which each occur with probability $p_{\min}$.

Formally, for each state s, let $\bar{p}_s = p_{\min}\lfloor p_s/p_{\min}\rfloor$ be p_s rounded down to the nearest factor of $p_{\min}$. In our new instance, we will replace s with $k(s) = \lfloor p_s/p_{\min}\rfloor$ identical states, labeled $s^1, \ldots, s^{k(s)}$. The new probabilities will be given by $p'_{s^1} = p'_{s^2} = \ldots = p'_{s^{k(s)}} = p_{\min}/c$, where $c = \sum_{s\in S}\bar{p}_s$ is a normalizing constant to ensure probabilities sum to 1. Our new utilities will be given by $U_{\mathrm{P}}'(a, s^i) = cU_{\mathrm{P}}(a, s)$ and $U_{\mathrm{A}}'(a, s^i) = cU_{\mathrm{A}}(a, s)$ for all $a \in \Omega$, $s \in S$, and $i \in \{1, \ldots, k(s)\}$. Note that in state s^i for any s and $i \in \{1, \ldots, k(s)\}$, the agent's preferences over actions are identical to their preferences in state s of the original instance. We now show that because $p_{\min}$ was smaller than all other probabilities, rounding in the manner discussed does not distort utilties significantly for the principal. In particular:

Lemma 5.7. *For any delegation set A, let $f(A)$ denote the principal's utility from A in the original instance, and $f'(A)$ the principal's utility from A in the transformed instance. The following inequality holds:*

$$2f'(A) \geq f(A) \geq f'(A).$$

Proof. We first prove $f(A) \geq f'(A)$. We have:

$$\begin{aligned} f'(A) &= \sum_{s\in S}\sum_{i=1}^{k(s)} p'_{s^i} U_P'(g(A,s^i),s^i) \\ &= \sum_{s\in S} \bar{p}_s U_P(g(A,s),s) \\ &\leq \sum_{s\in S} p_s U_P(g(A,s),s) = f(A) \end{aligned}$$

In fact, if we show $2\bar{p}_s \geq p_s$ for all $s \in S$, then the above analysis also implies $2f'(A) \geq f(A)$. This latter fact follows from the observation that $p_s - \bar{p}_s \leq p_{\min} \leq \bar{p}_s$. □

The above mapping from delegation instances with arbitrarily distributed states to those with uniformly distributed states immediately implies that we may use approximation algorithms from the latter setting to obtain approximately optimal solutions to the former. This will immediately yield Theorem 5.2 as a corollary. Formally:

Theorem 5.8. *Let $\alpha(\cdot)$ be a nondecreasing function that is at least 1. Then for any $\alpha(n)$-approximation algorithm for delegation instances with uniformly distributed states, there exists a $2\alpha(p_{min}^{-1})$-approximation algorithm for delegation instances with arbitrarily distributed states.*

Proof. Let an approximation algorithm for uniform states be given. Consider running the algorithm in question on the uniform instance constructed above, and let A and A^* be the set of actions produced by the algorithm and an optimal solution, respectively. Further let A' denote an optimal solution to the uniform instance. Note that the number n' of states in the uniform instance is at most $p_{\min}^{-1}$. By Lemma 5.7, the we have the following sequence of inequalities, which implies the result.

$$f(A) \geq f'(A) \geq \tfrac{1}{\alpha(n')} f'(A') \geq \tfrac{1}{\alpha(p_{\min}^{-1})} f'(A') \geq \tfrac{1}{\alpha(p_{\min}^{-1})} f'(A^*) \geq \tfrac{1}{2\alpha(p_{\min}^{-1})} f(A^*).$$

□

We conclude by re-emphasizing that Theorem 5.2 is asymptotically tight. The uniform states example in Appendix A.5 also applies in the more general setting of non-uniform states. Moreover, note that combining identical states yields an example with many fewer states, ruling out the possibility of an approximation factor which depends only on the number of states and not their distribution.

6 SPECIAL CASES WITH RESTRICTED DISTORTION

We finally note two additional assumptions which together immediately imply threshold policies are a *constant*-approximation. First, throughout the paper we have assumed that the principal's utility for every action and every state was nonnegative; this was to make approximation sensible. For the agent, on the other hand, the sign of their utility was not important. The only thing that mattered was their preference ordering over actions in each state. In this section, we add the assumption that the agent's utilities are also nonnegative, i.e. $U_A(a,s) \geq 0$ for all $a \in \Omega$ and $s \in S$.

Second, we note that prior to this section, the utility decomposition might yield distortions which are negative. In some settings, distortions are nonpositive. This might occur, for example, if distortions take the form of transfers from the agent to the principal, or as costs felt only by the agent. In this section, we assume that $d_a \leq 0$ for all $a \in \Omega$.

Under these two assumptions (and taking the decomposition with coefficients $a_1 = a_2 = a_3 = a_4 = 1$), we may write $U_P = V + D - 2D$. The agent's utility is given by $V + D$, and is maximized by the threshold strategy that offers the full set of actions. Meanwhile, selecting the action with the most negative distortion maximizes the term $-2D$, and is also a threshold strategy. This yields a simple argument that threshold strategies are a 2-approximation. In fact, a more nuanced characterization can be obtained, and is summarized in the following two theorems, proved in Appendices A.6 and A.7.

THEOREM 6.1. *Algorithm 1 is a $(2 - p_{min})$-approximation algorithm for the delegation problem under restricted distortion assumption, when the distortions are nonpositive and the agent's utilities are nonnegative.*

THEOREM 6.2. *There exists a family of instances satisfying restricted, nonpositive distortion and nonnegative agent utilities such that Algorithm 1 is at best a $(2 - p_{min})$-approximation.*

ACKNOWLEDGMENTS

The authors would like to thank Bobby Kleinberg for observing the approximation-preserving reduction from Independent Set, and Hamsa Bastani and Rediet Abebe for helpful discussions. Part of this work was done while Ali Khodabakhsh and Emmanouil Pountourakis were visiting the Simons Institute for the Theory of Computing.

REFERENCES

[1] Philippe Aghion and Jean Tirole. 1997. Formal and real authority in organizations. *Journal of political economy* 105, 1 (1997), 1–29.
[2] Ricardo Alonso and Niko Matouschek. 2008. Optimal delegation. *The Review of Economic Studies* 75, 1 (2008), 259–293.
[3] Manuel Amador and Kyle Bagwell. 2010. On the optimality of tariff caps. *Unpublished working paper* (2010).
[4] Attila Ambrus and Georgy Egorov. 2017. Delegation and nonmonetary incentives. *Journal of Economic Theory* 171 (2017), 101–135.
[5] Mark Armstrong and John Vickers. 2010. A model of delegated project choice. *Econometrica* 78, 1 (2010), 213–244.
[6] Shuchi Chawla, Jason D Hartline, and Robert Kleinberg. 2007. Algorithmic pricing via virtual valuations. In *Proceedings of the 8th ACM conference on Electronic commerce*. ACM, 243–251.
[7] Andrea EF Clementi and Luca Trevisan. 1999. Improved non-approximability results for minimum vertex cover with density constraints. *Theoretical Computer Science* 225, 1-2 (1999), 113–128.
[8] Vincent P Crawford and Joel Sobel. 1982. Strategic information transmission. *Econometrica: Journal of the Econometric Society* (1982), 1431–1451.
[9] Wouter Dessein. 2002. Authority and communication in organizations. *The Review of Economic Studies* 69, 4 (2002), 811–838.
[10] Paul Dütting, Tim Roughgarden, and Inbal Talgam-Cohen. 2018. Simple versus Optimal Contracts. *arXiv preprint arXiv:1808.03713* (2018).
[11] David Epstein and Sharyn O'Halloran. 1999. Asymmetric information, delegation, and the structure of policy-making. *Journal of Theoretical Politics* 11, 1 (1999), 37–56.
[12] Alexander Frankel. 2014. Aligned delegation. *American Economic Review* 104, 1 (2014), 66–83.
[13] Johan Håstad. 1999. Clique is hard to approximate within $n^{1-\varepsilon}$. *Acta Mathematica* 182, 1 (1999), 105–142.
[14] Bengt Robert Holmstrom. 1978. *On Incentives and Control in Organizations.* Ph. D. Dissertation.
[15] Nicole Immorlica, Brendan Lucier, Jieming Mao, Vasilis Syrgkanis, and Christos Tzamos. 2018. Combinatorial Assortment Optimization. In *International Conference on Web and Internet Economics*. Springer, 218–231.
[16] Sampath Kannan, Michael Kearns, Jamie Morgenstern, Mallesh Pai, Aaron Roth, Rakesh Vohra, and Zhiwei Steven Wu. 2017. Fairness Incentives for Myopic Agents. In *Proceedings of the 2017 ACM Conference on Economics and Computation*. ACM, 369–386.
[17] Jon Kleinberg and Robert Kleinberg. 2018. Delegated Search Approximates Efficient Search. In *Proceedings of the 2018 ACM Conference on Economics and Computation*. ACM, 287–302.

[18] Eugen Kovác and Tymofiy Mylovanov. 2009. Stochastic mechanisms in settings without monetary transfers: The regular case. *Journal of Economic Theory* 144, 4 (2009), 1373–1395.

[19] Vijay Krishna and John Morgan. 2008. Contracting for information under imperfect commitment. *The RAND Journal of Economics* 39, 4 (2008), 905–925.

[20] Yishay Mansour, Aleksandrs Slivkins, Vasilis Syrgkanis, and Zhiwei Steven Wu. 2016. Bayesian exploration: Incentivizing exploration in bayesian games. *arXiv preprint arXiv:1602.07570* (2016).

[21] David Martimort and Aggey Semenov. 2006. Continuity in mechanism design without transfers. *Economics Letters* 93, 2 (2006), 182–189.

[22] Nahum D Melumad and Toshiyuki Shibano. 1991. Communication in settings with no transfers. *The RAND Journal of Economics* (1991), 173–198.

[23] James A Mirrlees. 1971. An exploration in the theory of optimum income taxation. *The review of economic studies* 38, 2 (1971), 175–208.

[24] Dezsö Szalay. 2005. The economics of clear advice and extreme options. *The Review of Economic Studies* 72, 4 (2005), 1173–1198.

[25] Kalyan Talluri and Garrett Van Ryzin. 2004. Revenue management under a general discrete choice model of consumer behavior. *Management Science* 50, 1 (2004), 15–33.

[26] David Zuckerman. 2006. Linear degree extractors and the inapproximability of max clique and chromatic number. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*. ACM, 681–690.

A SUPPLEMENTARY MATERIALS

A.1 General Hardness Result

Here we analyze the delegation problem in full generality. We show that the problem is hard to approximate within any sublinear factor. This holds even if the distribution over the states is uniform, the principal's utility is binary (zero or one), and the agent's utility is limited to three different values.

Proof. We reduce from the Maximum Independent Set (MIS) problem. Given an instance of the MIS problem $\mathcal{G} = (\mathcal{V}, \mathcal{E})$[7] with $|\mathcal{V}| = N$, we build an instance of the delegation problem with N actions and N states; i.e., one action and one state corresponding to each vertex of $\mathcal{G}$. We also assume that the states happen with the same probability $1/N$. We set the utility of the principal to be the identity matrix, i.e., $U_{\mathrm{P}} = \mathbb{I}_{N\times N}$. For the agent, the utility of action i at state j is defined by:

$$U_{\mathrm{A}}(i,j) = \begin{cases} 1 & \text{if } i = j, \\ 2 & \text{if } \{i,j\} \in \mathcal{E}, \\ 0 & \text{otherwise.} \end{cases} \tag{11}$$

Choosing a set of actions by the principal can be thought as choosing a subset $\mathcal{S} \subseteq \mathcal{V}$. We claim that the utility of the principal is less than or equal to $|\mathcal{S}|/N$, with equality if and only if $\mathcal{S}$ is an independent set in $\mathcal{G}$. The inequality is clear because at each state $i \in \mathcal{S}$, the principal can get at most one by the choice of U_{P} (and zero in other states). For the equality to occur, the agent should choose the action i at each state $i \in \mathcal{S}$. This requires that no neighbor of i is available, and hence $\mathcal{S}$ is an independent set.

By the previous argument, the maximum independent set maximizes the utility of the principal. On the other hand, given a solution of the delegation problem with (principal's) utility value k/N, one can find an independent set of size k (by just selecting the states where the utility of the principal is non-zero). This implies that our reduction also preserves the objective values (besides normalizing by N). Therefore, delegation inherits the hardness of approximation of the MIS problem [13, 26]. □

[7]Throughout the paper, we use calligraphic letters for graph instances to distinguish between $\mathcal{V}$ (set of vertices) and V (value matrix in delegation instance).

A.2 Hardness Result for Restricted Value with Known Outside Option

Here we will show that given a known outside option, restricted value case is hard to approximate within any sublinear factor, which is similar to the general case.

Proof. In this section, we construct a restricted value case that is equivalent to example in the previous section. We set the utility of the agent to be the utility of the example in the previous section. For the principle, the utility of action i at state j is defined by:

$$U_{\mathrm{A}}(i,j) = \begin{cases} 1 & \text{if } i = j, \\ 0 & \text{if } \{i,j\} \in \mathcal{E}, \\ 2 & \text{otherwise.} \end{cases} \tag{12}$$

Note that for any action i state j pair, $U_{\mathrm{A}}(i,j) + U_{\mathrm{P}}(i,j) = 2$, so this case is restricted value. Then we set the outside option in the following way: let agent utility be $0 < \epsilon < 1$ for all states; let principle utility be 0 for all states. Then if $i = j$ or $\{i,j\} \notin \mathcal{E}$, the agent will not pick the outside option. Otherwise only if $\{i,j\} \in \mathcal{E}$, the agent will pick the outside option. Then the principle's utility is equivalent to an identity matrix. Therefore, this case is equivalent to the example above and we can do a same reduction from MIS problem. □

A.3 Restricted Value Example

Example A.1. Consider a firm (e.g. a law firm) who wants to undertake a new project (e.g. a new case), from three possible projects A, B, and C. The firm (principal) will get paid a fixed amount upon completion of each project, denoted by r_A, r_B, and r_C. However, there are project-specific effort costs that are only observed by the employees (lawyers) who will be undertaking the project. These costs also depend on the state, which can correspond to the statues of the available resources, unknown to the principal. As an example, consider the following cost matrix which measures these project-specific effort costs under three different states.

	1	2	3
A	1	1	2
B	1	1	2
C	1	3	1

The first column corresponds to a situation where all projects have the same effort cost (e.g., when there are enough resources for all candidate projects), while in the second column, project C becomes a challenge (e.g., due to lack of required resources), and the same for projects A and B under last column. Instantiating the revenue from each project as $r_A = 2$, $r_B = 4$, and $r_C = 5$, we may produce the following utility matrix for the agent (lawyer), whose utility is revenue minus effort costs (U_{A}, at left):

$$U_{\mathrm{A}} = \begin{bmatrix} 1 & 1 & 0 \\ 3 & 3 & 2 \\ 4 & 2 & 4 \end{bmatrix}, \quad U_{\mathrm{P}} = \begin{bmatrix} 2 & 2 & 2 \\ 4 & 4 & 4 \\ 5 & 5 & 5 \end{bmatrix}.$$

That is, taking project A under state 1 gives the agent utility $2 - 1 = 1$, and the other entries of the matrix follow similarly. We assume that the principal seeks to maximize the revenue, therefore his utility will be the state-independent matrix shown at right, above (U_{P}). Note that because the efforts change the agents's utility for the projects, the principal and agent's incentives are misaligned under state 2 - the principal prefers to pick project C, while the agent prefers project B due to less effort.

Under the state indexed j, the agent picks their most desired project $g(X, j)$ from the set X of available projects. For example, if X is the set of all projects, we would have:

$$g(X, j) = \begin{cases} \text{C} & j = 1, \\ \text{B} & j = 2, \\ \text{C} & j = 3. \end{cases}$$

The principal could decide not to make all the projects available, and instead offer $Y = \{C\}$, for example. In this case, the agent would be forced to pick project C regardless of the state, and the principal would achieve the optimal expected utility of $f(Y) = 5$. We indeed show in Section 4 that this is true in general, meaning that in the case of restricted value, it is always optimal for the principal to pick their most preferred action ex ante, without consulting the agent.

A.4 Proof of Theorem 5.1

We give a reduction from the bounded degree vertex cover problem, i.e., the vertex cover problem on graphs with degree at most B (constant). This problem is known to be APX-hard [7]. Consider an instance of the bounded degree vertex cover problem $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with $\tilde{n}$ nodes and $\tilde{m}$ edges (where $\tilde{m} \leq B \cdot \tilde{n}/2 = O(\tilde{n})$).[8] We construct an instance of the delegation problem with $\tilde{m} + \tilde{n}$ states and $\tilde{n} + 1$ actions, in which the distribution over states is uniform. For each node i, there is an action a_i, with an additional "default" action a_0. Regarding states, there is a state s_e for each edge $e = \{i, j\}$, whose value is 5 only for actions a_i and a_j (two ends of that edge), 2 for the default action, and 0 for all other actions. There is also a state s_i for each node i, whose value is 2 only for a_i and a_0 (default action), and 0 for all other actions. Moreover, the only action with non-zero distortion is the default action with $d_0 = -1$.

We claim that the optimal solution of the delegation problem produces a utility of $(5\tilde{m} + 3\tilde{n} - \tilde{k})/(\tilde{m} + \tilde{n})$ for the principal, where $\tilde{k}$ is the size of the smallest vertex cover of $\mathcal{G}$. To see this, first note that any solution $\mathcal{S} \subseteq \mathcal{V}$ can be improved by including a_0, since a_0 has a negative distortion. Any time the agent would choose a_0, it is the optimal choice for the principal as well. We therefore only consider solutions containing a_0.

Now if $\mathcal{S}$ is a vertex cover of $\mathcal{G}$ with $|\mathcal{S}| = \tilde{k}$, consider the corresponding delegation set where the principal allows actions $\{a_i : i \in \mathcal{S}\} \cup \{a_0\}$. In all the states corresponding to edges, the agent will pick the action corresponding to one end of that edge (one is guaranteed to be in the cover $\mathcal{S}$) to get a utility of 5 compared to $2 - 1$ achievable from the default action. This choice will also generate utility of 5 for the principal, which makes $5\tilde{m}$ in total. For the states $s_i, i \in \mathcal{S}$ the agent will pick action a_i which generates the utility of 2 for both principal and agent. This will make $2\tilde{k}$ in total. Finally, for states $s_i, i \in \mathcal{V} \backslash \mathcal{S}$ the agent picks the default action which generates a utility of $2 - 1$ for himself but $2 + 1$ for the principal. This will give $3(\tilde{n} - \tilde{k})$ in total. Summing up the utilities and dividing by the number of states will give $(5\tilde{m} + 3\tilde{n} - \tilde{k})/(\tilde{m} + \tilde{n})$.

For the converse, consider an optimal solution A to the delegation problem. We show that the nodes corresponding to the actions in A (excluding the default action) induce a vertex cover; otherwise the solution can be improved. Assume that there exists an edge $e = \{i, j\}$ where neither a_i nor a_j is allowed in A. If we add action a_i to A, the principal gets a utility of 5 in state s_e, compared to current utility of 3 from the default action. On the other hand, the utility of the principal decreases in state s_i from 3 to 2. So the total utility of $A \cup \{a_i\}$ is more than A which contradicts the optimality of A. Therefore A should be a vertex cover (plus default action). This in turn implies that the utility is at most $(5\tilde{m} + 3\tilde{n} - \tilde{k})/(\tilde{m} + \tilde{n})$ where $\tilde{k}$ is the size of the minimum vertex cover.

[8] To distinguish between the parameters of the vertex cover instance and the delegation instance, we use tilde ($\sim$) for the graph instance.

Since $\tilde{m} = \Theta(\tilde{n})$ and the minimum vertex cover has size at least $\tilde{m}/B = \Omega(\tilde{n})$, a constant factor gap in the bounded degree vertex cover problem translates into a constant factor gap in the optimal solution of the delegation problem, which yields the desired APX-hardness result.

A.5 Proof of Theorem 5.3

In this appendix, we show that the logarithmic approximation factor of Algorithm 1 is tight, up to a constant factor. That is, no threshold algorithm can perform better than $(\log n)/2$. To prove the tightness of our analysis in Section 5.1, we construct an infinite family of instances. Instead of presenting value and distortion matrices, we define the utility of the principal in (13) and also restricted distortion of each action in (14). Consequently, the agent's utility and the value matrix can be calculated as $U_{\mathrm{A}} = U_{\mathrm{P}} + 2D$ and $V = U_{\mathrm{P}} + D$. For any $k \geq 2$, consider an instance with $m = 2k - 1$ actions (rows) and $n = 2^k - 1$ states (columns), with the following principal's utility matrix:

$$U_{\mathrm{P}} = \left[\begin{array}{c|c|c|c|c|c|c|c|c}
2^k & & & & & & & & \\ \hline
2^{k-1}+3\epsilon & & & & & & & & \\ \hline
 & 2^{k-1} & 2^{k-1} & & & & & & \\ \hline
2^{k-2}+4\epsilon & 2^{k-2}+4\epsilon & 2^{k-2}+4\epsilon & & & & & & \\ \hline
 & & & 2^{k-2} & 2^{k-2} & 2^{k-2} & 2^{k-2} & & \\ \hline
2^{k-3}+5\epsilon & 2^{k-3}+5\epsilon & 2^{k-3}+5\epsilon & 2^{k-3}+5\epsilon & 2^{k-3}+5\epsilon & 2^{k-3}+5\epsilon & 2^{k-3}+5\epsilon & & \\ \hline
\vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \\ \hline
2+O(k\epsilon) & 2+O(k\epsilon) & 2+O(k\epsilon) & 2+O(k\epsilon) & 2+O(k\epsilon) & 2+O(k\epsilon) & 2+O(k\epsilon) & \cdots & \\ \hline
 & & & & & & & \cdots & \underbrace{2 \ \cdots \ 2}_{2^{k-1}}
\end{array}\right] \tag{13}$$

Note that all the empty entries in the above matrix are zero, and are removed to make the structure of the matrix more apparent. The optimal solution is to select set of odd actions (with size k). The colored entries indicate (state,action) pairs that contribute to the optimal principal's utility (OPT). The optimal utility is equally divided between the odd rows, leaving 2^k for each one: the first row has 2^k in the first column, the third row has 2^{k-1} in columns 2 and 3, the fifth row has 2^{k-2} over the next 4 columns and so on. In the example, the even actions are constructed to lower the principal's utility whenever they are included in a threshold solution. In every state (column), we divide the colored utility by 2 to find the utility of the next row, and keep dividing by 2 to complete the subsequent even rows. The ϵ terms are added to break the ties and are of little importance.

Next, we define the distortion matrix. Since we are in the restricted distortion setting, it is sufficient to determine the (state-independent) distortion of each action. We set $d_1 = 0$, and the rest of actions have the following distortions:

$$d_{2i+1} = \sum_{j=1}^{i} 2^{k-1-j}, \quad d_{2i} = d_{2i+1} - \epsilon, \qquad i \in \{1, \dots, k-1\}. \tag{14}$$

Now that all the parameters are set, it is easy to verify that given the set of odd actions (rows), the agent will indeed pick the colored entries. This generates the optimal utility, since it is optimal in every single state. Assuming uniform distribution over states, the optimal expected utility is equal to:

$$OPT = \frac{k \times 2^k}{n}.$$

However, the best threshold solution in the constructed instance is to allow the entire set of actions (Ω). To see this, assume that the principal allows actions with distortion less than or equal to $d_{2\ell-1}$ for some $\ell \leq k$. (Thresholds set at even-indexed actions can be easily shown to be suboptimal.) In this case, the principal will get utility of $2^{k-\ell+1} + O(\ell\epsilon)$ from the first $2^\ell - 1$ states, and zero from the remaining states. Observe that the overall utility $(2^\ell - 1) \times 2^{k-\ell+1}$ is an increasing function in ℓ, meaning that the best strategy for the principal is to not limit the agent. In this case, the agent will pick the penultimate action in the first half of columns, and the last action for the second half, generating utility of (almost) 2 for principal in every state. More precisely, we have:

$$APX = 2 + O(k\epsilon).$$

We get the desired lower bound by dividing the above objectives:

$$\frac{OPT}{APX} = \frac{k \times 2^k}{2n + O(nk\epsilon)} \cong \frac{k}{2} \cong \frac{\log n}{2}.$$

Example A.2. In order to make sure that the above construction is clear, here we present the full matrices for the case of $k = 3$, which translates into $m = 5$ actions and $n = 7$ states. The principal's utility in this case is:

$$U_{\mathrm{P}} = \begin{bmatrix} 8 & 0 & 0 & 0 & 0 & 0 & 0 \\ 4+3\epsilon & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 4 & 4 & 0 & 0 & 0 & 0 \\ 2+4\epsilon & 2+4\epsilon & 2+4\epsilon & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 2 & 2 & 2 & 2 \end{bmatrix}$$

Calculating the distortions in (14) results in:

$$D = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 2-\epsilon & 2-\epsilon & 2-\epsilon & 2-\epsilon & 2-\epsilon & 2-\epsilon & 2-\epsilon \\ 2 & 2 & 2 & 2 & 2 & 2 & 2 \\ 3-\epsilon & 3-\epsilon & 3-\epsilon & 3-\epsilon & 3-\epsilon & 3-\epsilon & 3-\epsilon \\ 3 & 3 & 3 & 3 & 3 & 3 & 3 \end{bmatrix}$$

It is clear that the value matrix $V = U_{\mathrm{P}} + D$ is non-negative, and the agent's utility $U_{\mathrm{A}} = U_{\mathrm{P}} + 2D$ will be:

$$U_{\mathrm{A}} = \begin{bmatrix} 8 & 0 & 0 & 0 & 0 & 0 & 0 \\ 8+\epsilon & 4-2\epsilon & 4-2\epsilon & 4-2\epsilon & 4-2\epsilon & 4-2\epsilon & 4-2\epsilon \\ 4 & 8 & 8 & 4 & 4 & 4 & 4 \\ 8+2\epsilon & 8+2\epsilon & 8+2\epsilon & 6-2\epsilon & 6-2\epsilon & 6-2\epsilon & 6-2\epsilon \\ 6 & 6 & 6 & 8 & 8 & 8 & 8 \end{bmatrix}$$

Observe that $OPT = 24/7$ by the set of odd actions $\{1, 3, 5\}$, while $APX = (14 + 12\epsilon)/7$ from the entire set of actions $\Omega = \{1, 2, 3, 4, 5\}$.

A.6 Proof of Theorem 6.1

The algorithm is similar to Algorithm 1: try every threshold policy, and adopt the best. We show that this results in an approximation ratio of $2 - p_{\min}$.

Proof. Let OPT be an optimal solution, and define a_i to be the index of action chosen by the agent, in the state i, given that he is allowed the actions in OPT. Let $\overline{d}$ be the maximum distortion in OPT. Without loss of generality, we can assume that the corresponding action is taken in at least one state, otherwise we can remove it without changing the optimal utility. Therefore we can write,

$$\overline{d} = \max_{i \in OPT} d_i = \max_{i \in S} d_{a_i}.$$

Let APX be the subset of actions within distortion range of $[\underline{d}, \overline{d}]$, where $\underline{d}$ is the lowest distortion in Ω. Although we do not know $\overline{d}$, the algorithm is guaranteed to find APX at some iteration. Let b_i be the index of action taken in state i, given that the agent is allowed to pick actions in APX.

We show that if the agent picks a lower distortion compared to the optimal solution (i.e., $d_{b_i} \leq d_{a_i}$) while the action a_i is allowed, the principal is guaranteed to get a utility as big as the optimal solution in that particular state i. On the other hand, if the agent picks a higher distortion, we show that the principal loses at most twice the distortion difference (i.e., $2(d_{b_i} - d_{a_i})$). To show this, note that we have:

$$v_{b_i,i} + d_{b_i} \geq v_{a_i,i} + d_{a_i}, \tag{15}$$

because $a_i \in APX$ (by choice of $\overline{d}$), but the agent decided to pick b_i. This inequality can be re-written as

$$v_{b_i,i} - d_{b_i} \geq v_{a_i,i} - d_{a_i} - 2(d_{b_i} - d_{a_i}),$$

which justifies both cases mentioned above. Multiplying by corresponding probabilities and summing up over all states, we can form the utility of the principal in both optimal and approximate solutions:

$$\sum_{i \in S} p_i(v_{b_i,i} - d_{b_i}) \geq \sum_{i \in S} p_i(v_{a_i,i} - d_{a_i}) - 2 \sum_{i \in S} p_i(d_{b_i} - d_{a_i}).$$

Therefore,

$$APX \geq OPT - 2 \sum_{i \in S} p_i(d_{b_i} - d_{a_i}), \tag{16}$$

where we abused notation to write OPT instead of $f(OPT)$ (same for APX).

Next, we have to bound the distortion differences. By non-positivity of distortion, any term $d_{b_i} - d_{a_i}$ is at most $|\underline{d}|$. Moreover, $d_{b_i} - d_{a_i}$ is non-positive for at least one of the states. This is because $d_{b_i} \leq \overline{d}$ (for all $i \in S$), and $\overline{d} = d_{a_{i'}}$ for one particular state i'. In conclusion,

$$\sum_{i \in S} p_i(d_{b_i} - d_{a_i}) \leq |\underline{d}| \sum_{i \neq i'} p_i \leq |\underline{d}|(1 - p_{\min}).$$

Using this bound in (16), we get

$$APX \geq OPT - 2(1 - p_{\min})|\underline{d}|. \tag{17}$$

In parallel, consider the singleton solution of lowest distortion. In that case, the agent is forced to pick that low-distortion action in every single state. By non-negativity of (agent) utilities, the value of that action is at least $|\underline{d}|$ in every state, generating the minimum utility of $2|\underline{d}|$ for the principal. If we denote this singleton solution by APX', we have

$$APX' \geq 2|\underline{d}|. \tag{18}$$

The final step is to take a weighted average of (17) and (18), as the better of the two will be better than the average. We multiply the inequalities by $\frac{1}{2-p_{\min}}$ and $\frac{1-p_{\min}}{2-p_{\min}}$, respectively, and add them together.

$$\max\{APX, APX'\} \geq \left(\frac{1}{2 - p_{\min}}\right)APX + \left(\frac{1 - p_{\min}}{2 - p_{\min}}\right)APX' \geq \frac{OPT}{2 - p_{\min}}.$$

□

A.7 A Matching Lower Bound for Theorem 6.1

Here we show that the $(2 - p_{\min})$-approximation factor of Theorem 6.1 is tight. In fact, no thresholding algorithm can perform better than $(2 - p_{\min})$, where by thresholding algorithm we mean any algorithm that outputs actions within a continuous range of distortion. In other words, the thresholding algorithm is not allowed to have gaps between actions: if two actions are chosen, any other action with a middle distortion should be included.

Example A.3. Consider an example with $m = 3$ actions and $n \geq 2$ (uniform) states. The value matrix is given by

$$V = \left[\begin{array}{c|c|c|c|c} (2-\epsilon)n & 0 & 0 & \dots & 0 \\ \hline \epsilon & 3\epsilon & 3\epsilon & \dots & 3\epsilon \\ \hline 1 & \underbrace{1+\epsilon \;\; 1+\epsilon \;\; \dots \;\; 1+\epsilon}_{n-1} \end{array}\right],$$

and the restricted distortions are $(d_1, d_2, d_3) = (0, -\epsilon, -1)$. It is straightforward to see that the optimal set of actions is $OPT = \{1, 3\}$ with objective value

$$f(OPT) = \frac{1}{n}\left[(2-\epsilon)n + (2+\epsilon)(n-1)\right] = \frac{4n-2-\epsilon}{n}.$$

On the other hand, Algorithms 1 will return $\Omega = \{1, 2, 3\}$ with objective value

$$f(\Omega) = \frac{1}{n}\left[(2-\epsilon)n + 4\epsilon(n-1)\right] = \frac{(2+3\epsilon)n - 4\epsilon}{n}.$$

As ϵ goes to zero, we get our desired lower bound for the approximation ratio, which matches the result of Theorem 6.1.

$$\lim_{\epsilon \to 0} \frac{f(OPT)}{f(\Omega)} = \frac{4n-2}{2n} = 2 - \frac{1}{n}.$$