

Optimal Delegation in a Multidimensional World

Andreas Kleiner*

Preliminary Draft
September 26, 2022

We study a model of delegation in which a principal takes a multidimensional action and an agent has private information about a multidimensional state of the world. The principal can design any direct mechanism, including stochastic ones. We provide necessary and sufficient conditions for an arbitrary mechanism to maximize the principal's expected payoff. We also discuss simple conditions which ensure that some *convex delegation set* is optimal. A key step of our analysis shows that a mechanism is incentive compatible if and only if its induced indirect utility is convex and lies below the agent's first-best payoff.

*Department of Economics, Arizona State University. Email: andreas.kleiner@asu.edu. I thank Navin Kartik, Alejandro Manelli, and Kun Zhang for useful comments and discussions.

1. Introduction

In many economic and political environments, a principal faces a better-informed but biased agent. The principal can choose a permissible set of actions and ‘delegate’ the decision to the agent: a firm appoints a manager to select investment levels in different projects; US Congress delegates power to federal agencies; a legislative forms a committee to draft bills; a regulator lets a monopolist choose prices. The principal may yield some discretion to the agent to utilize his informational advantage but may impose restrictions on the agent’s actions to counter his bias. Following Holmström (1977), an extensive literature models such delegation problems by assuming that both the action and the state of the world lie in a one-dimensional space. A main result of this literature characterizes when it is optimal for the principal to constrain the agent’s choice to lie in an interval, and this conclusion has been used to explain why managers face spending caps, regulators impose price ceilings, and trade agreements specify maximum tariff levels.

The assumption that the action and state space are one-dimensional is made for tractability. In many applications, the underlying states and actions are more complex and more realistically modeled as multidimensional: managers invest in several projects, Congress delegates many decisions to the EPA, and committees draft multiple bills. What mechanisms are optimal in such multidimensional settings? How robust are conclusions obtained for one-dimensional models? And can we still expect that relatively simple mechanisms are often optimal?

To study these questions, we consider a principal that takes a multidimensional action and faces an agent with private information about a multidimensional state of the world (the agent’s *type*). Payoffs depend on the action and the state of the world, and transfers are infeasible. The principal can design arbitrary mechanisms, including stochastic ones, to maximize her expected payoff. Our main result characterizes, for an arbitrary mechanism, when this mechanism is optimal. Often, it is optimal to delegate the decision to the agent but to constrain the agent by requiring that her action lies in some set. For convex delegation sets, we provide a simple characterization, which is a direct analog of conditions characterizing when interval delegation is optimal in one-dimensional models. Our results illustrate how a principal can benefit from optimally bundling independent decisions. Even for one-dimensional models, this approach provides new insights: our main result characterizes for arbitrary mechanisms—not just interval delegation sets—when this mechanism is optimal. And as corollaries, we obtain novel conditions under which some interval delegation set will

be optimal.

A key step to deriving our results lies in obtaining a simple characterization of the set of feasible mechanisms. Given a mechanism, the corresponding *indirect utility* assigns to any type the payoff this type would get by choosing his report optimally. This payoff must be less than the *first-best payoff*, i.e., the payoff this type would get if he could choose the action without any restrictions. Moreover, our assumption that the agent’s utility is an affine function of the state implies that the indirect utility must be a convex function because it is the maximum of a family of affine functions. Lemma 1 shows that any function satisfying these two properties is the indirect utility of an incentive-compatible mechanism.

This characterization is easy to use and already helpful for one-dimensional delegation models. Our formulation differs from the previous literature, which often considered only deterministic mechanisms. Since the convex combination of two incentive-compatible deterministic mechanisms is not necessarily incentive compatible, the set of deterministic mechanisms is not even convex.¹ Moreover, a common approach is to first treat the model as one with transfers and then impose that these transfers are zero (or negative). Compared to this approach, formulating the problem via indirect utilities is more direct and provides valuable geometric insights into which mechanisms can be optimal. For the multidimensional problem, the approach via indirect utilities provides additional benefits because it circumvents intricate characterizations of incentive compatibility (see Rochet, 1987).

To find the optimal mechanism, we formulate the principal’s problem in terms of indirect utilities. In this formulation, the problem becomes a linear program, and we use linear programming duality to derive necessary and sufficient conditions for a given mechanism to be optimal. Typically, optimal mechanisms pool certain types, and our main result shows that a mechanism is optimal if conditional on any pooling region, a stochastic dominance condition (using the convex order) is satisfied. Intuitively, this condition requires that, restricted to the pooling region (where the indirect utility function is affine), any convex indirect utility yields a lower payoff. If the pooling regions are at most one-dimensional, the stochastic dominance condition has a simple formulation in terms of majorization. Using this observation, we provide necessary and sufficient conditions for a convex delegation set with a smooth boundary to be optimal. These conditions are easy to check and are straightforward extensions of

¹ Some earlier papers also consider stochastic mechanisms (or allow for money burning/restricted transfer) and obtain a convex set of mechanisms; see, for example, Amador and Bagwell (2013), Kovác and Mylovánov (2009), Ambrus and Egorov (2017), Amador and Bagwell (2020), Kartik, Kleiner, and Van Weelden (2021), and Kleiner, Moldovanu, and Strack (2021).

conditions that ensure the optimality of interval delegation sets in one-dimensional models (see Amador and Bagwell, 2013).

Related Literature The literature on delegation has focused mainly on problems in which the principal delegates a single one-dimensional decision and therefore assumed that both the action and state spaces are one-dimensional; see, for example, Holmström (1977), Holmström (1984), Melumad and Shibano (1991), Alonso and Matouschek (2008), Amador and Bagwell (2013), and Kolotilin and Zapechelnuyk (2019).

A few delegation papers do consider richer action and/or type spaces. Armstrong (1995) considers an agent with two-dimensional private information and discusses several applications. Since the principal’s action is assumed to be one-dimensional (and only interval delegation sets are considered), there is only limited scope to screen two-dimensional types in his analysis. Koessler and Martimort (2012) characterize the optimal mechanism in a setting where two decisions depend on a single-dimensional underlying state. Galperti (2019) considers multidimensional information and actions but restricts the principal’s choice to a particular class of “budgeting mechanisms”. The closest paper to ours is Frankel (2016), which studies the delegation of several independent decisions, which yield multidimensional action and state spaces. For quadratic preferences with a constant bias, he shows that if the states are independently and identically distributed according to normal distributions then it is optimal to delegate a ‘half space’. Without the normality assumption, he shows that the principal’s payoff from such a mechanism converges to the first-best as the number of independent decision problems grows. Frankel (2014) also considers multidimensional delegation problems and characterizes the max-min optimal mechanism, which maximizes the principal’s payoff against the worst-case preference type of the agent.

The elicitation of information about multiple independent decisions from a biased agent has been studied in general mechanism design (e.g., Jackson and Sonnenschein, 2007) and cheap talk environments (Chakraborty and Harbaugh, 2007; Lipnowski and Ravid, 2020). Jackson and Sonnenschein (2007) show that by linking independent decisions, the principal’s payoff converges to the first-best as the number of decisions grows. Our results can be used to show how the principal should optimally link decisions, which can be important if there is a limited number of decisions.

On a methodological level, our work is related to the literature on multidimensional mechanism design, and in particular on multiproduct monopolists (see, e.g., Rochet, 1987;

Manelli and Vincent, 2006; Manelli and Vincent, 2007; Daskalakis, Deckelbaum, and Tzamos, 2017; Haghpanah and Hartline, 2021).

2. Model

A principal chooses an action $a \in \mathbb{R}^n$. An agent is privately informed about the state of the world $s \in S$, where $S \subseteq \mathbb{R}^n$ is compact and convex. The agent's and principal's payoffs depend on both the action and the state of the world, and are given by

$$\begin{aligned} u_A(a, s) &:= a \cdot s + b(a) \\ u_P(a, s) &:= a \cdot g(s) + \kappa b(a), \end{aligned}$$

respectively, where $b : \mathbb{R}^n \rightarrow \mathbb{R}$ is strictly concave, differentiable with a Lipschitz-continuous gradient mapping, and satisfies $\lim_{\|a\| \rightarrow \infty} \frac{b(a)}{\|a\|} = -\infty$,² $g : S \rightarrow \mathbb{R}^n$ is Lipschitz-continuous, and $\kappa > 0$.

We assume that the state s is distributed according to a probability distribution F with differentiable density f and support S . The principal aims to maximize her expected payoff and can design arbitrary mechanisms.

The revelation principle applies and we define a *mechanism* to be a function $m : S \rightarrow \Delta(\mathbb{R}^n)$ such that expected payoffs are integrable.³ To simplify notation, we extend the domain of $b(\cdot)$ and $u_i(\cdot, s)$ linearly to include probability distributions over \mathbb{R}^n , so that $b(m(s)) = \mathbb{E}_{m(s)}[b(a)]$ and analogously for $u_i(\cdot, s)$. A mechanism is *incentive compatible* if for all s and s' in S ,

$$u_A(m(s), s) \geq u_A(m(s'), s).$$

²Since the set of actions is unbounded, some assumptions are needed to ensure that for every type $s \in S$ there is an optimal action. This is ensured by the assumptions on b ; weaker conditions could be used but would complicate some arguments.

³We denote by $\Delta(\mathbb{R}^n)$ the Borel σ -algebra on \mathbb{R}^n .




Figure 1: The function U satisfies $U(s) \leq h(s)$ for all $s \in S$ but does not correspond to a feasible mechanism. To see this, note that there is no convex extension of U to \mathbb{R} such that the extension lies below h . Lemma 1 then implies that U is not the indirect utility of any feasible mechanism.

3. Characterizing incentive-compatible mechanisms

We characterize the set of incentive-compatible mechanisms in terms of their indirect utilities. To any incentive-compatible mechanism m corresponds an *indirect utility* $U : \mathbb{R}^n \rightarrow \mathbb{R}$ defined by

$$U(s) := \sup_{s' \in S} \mathbb{E}[m(s')] \cdot s + b(m(s')).$$

Which indirect utilities correspond to some incentive-compatible mechanism? First, any indirect utility is convex as the supremum of a family of functions that are affine in the state s . Second, in the absence of transfers the agent's utility cannot be higher than if he was free to choose his action. Defining the *first-best payoff* $h : \mathbb{R}^n \rightarrow \mathbb{R}$ by

$$h(s) := \sup_{a \in \mathbb{R}^n} a \cdot s + b(a),$$

$U \leq h$ is clearly necessary.⁴ The following result shows that these two conditions characterize the set of feasible indirect utilities.

Lemma 1. *An indirect utility U corresponds to an incentive-compatible mechanism if and only if U is convex and lies below the first-best payoff: $U \leq h$.*

⁴We denote the pointwise order by \leq , so $U \leq h$ means $U(s) \leq h(s)$ for all s in the common domain of U and h .

Intuitively, if U is convex then it would correspond to an incentive-compatible mechanism if transfers were available and the agent had quasi-linear preferences. If the required transfers are all negative then we can use the agent's risk aversion (coming from the strict concavity of b) to simulate these transfers via stochastic actions. One can show that $U \leq h$ implies that the required transfers are negative. This last step relies on the domain of U and h being large enough and it would not suffice to require only that $U(s) \leq h(s)$ for all $s \in S$. Figure 1 illustrates a convex function U which lies below h on all of S , but which does not correspond to a mechanism because the lotteries assigned to low types would yield a strictly higher payoff than the first-best payoffs for some hypothetical types, an impossibility.

Proof. Let us first recall basic observations from convex analysis. The convex conjugate of a function U is denoted by U^* and defined by $U^*(a) := \sup_{s \in \mathbb{R}^n} a \cdot s - U(s)$. We will use the following facts, which follow immediately from this definition: (i) $h = (-b)^*$, (ii) $U \leq h$ implies $h^* \leq U^*$, and (iii) $a \in \partial U(s)$ implies $U^*(a) = a \cdot s - U(s)$.⁵

Suppose U is convex and satisfies $U \leq h$. Let the mechanism m assign to any type $s \in S$ a lottery with expected value $a \in \partial U(s)$ that yields the payoff $a \cdot s + b(a) - U^*(a) + h^*(a)$. Such a lottery exists because $a \cdot s + b(a)$ would be the payoff for type s from always getting action a , because fact (ii) implies that the agent's payoff is lower, and because b is strictly concave.⁶ Then facts (i) and (iii) imply that the payoff of a truthful type s is $U(s)$:

$$u_A(m(s), s) = s \cdot a + b(a) - U^*(a) + h^*(a) = U(s).$$

It remains to show that m is incentive compatible. For all s and s' ,

$$\begin{aligned} u_A(m(s), s) &= U(s) \geq U(s') + \mathbb{E}[m(s')] \cdot (s - s') \\ &= \mathbb{E}[m(s')] \cdot s' + b(m(s')) + \mathbb{E}[m(s')] \cdot (s - s') = u_A(m(s'), s), \end{aligned}$$

where the first inequality follows since $\mathbb{E}[m(s)] \in \partial U(s')$.

Q.E.D.

⁵ Here, $\partial U(s)$ denotes the subdifferential of U at s . To see (iii), note that the definition of U^* implies $U^*(a) \geq a \cdot s - U(s)$. Conversely, convexity of U and $a \in \partial U(s)$ imply that for all s' , $a \cdot s - U(s) \geq a \cdot s' - U(s')$. Taking the supremum of the right-hand side with respect to s' yields $a \cdot s - U(s) \geq U^*(a)$.

⁶ More formally, strict concavity of b implies that for any $a \in \mathbb{R}^n$ and nonzero $d \in \mathbb{R}^n$ there is $\varepsilon > 0$ such that $1/2[b(a+d) + b(a-d)] < b(a) - \varepsilon$. It follows that for any $\alpha > 1$, $1/2[b(a+\alpha d) + b(a-\alpha d)] \leq b(a) - \alpha\varepsilon$. Therefore, by choosing α arbitrarily large, one can design lotteries with expected value a that yield arbitrarily low payoff to the agent.

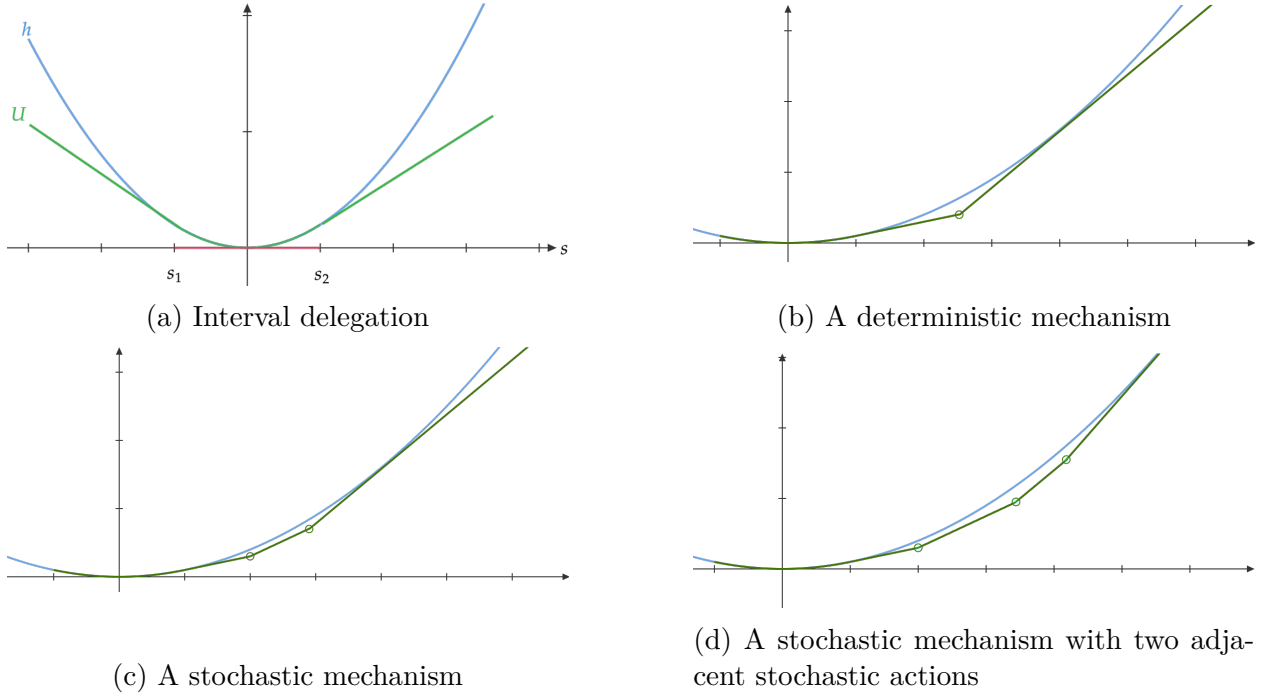


Figure 2: Examples of indirect utilities. The blue curves show the function h for one-dimensional types and quadratic payoffs (i.e., assuming $b(a) = -\frac{a^2}{2}$). The green curves show indirect utilities corresponding to incentive-compatible mechanisms.

Figure 2 illustrates the result for one-dimensional types and quadratic payoffs. It shows four indirect utilities that, according to Lemma 1, correspond to incentive-compatible mechanisms. In Figure 2a, all types between s_1 and s_2 obtain their first-best utility and U is affine below s_1 and above s_2 . This indirect utility can be obtained by letting types choose their preferred action from the interval of deterministic actions $[s_1, s_2]$. In Figure 2b, the menu of actions from which the agent can choose contains an additional deterministic action above s_2 . The indirect utility in Figure 2c contains an affine piece that lies strictly below the graph of h . This part of the indirect utility corresponds to types that obtain a (nondegenerate) stochastic action, which yields no type its first-best payoff. Finally, Figure 2d illustrates an indirect utility corresponding to a mechanism in which types in two adjacent regions obtain a stochastic action.

4. Characterizing optimal mechanisms

We characterize the optimal mechanisms in this section. To do so, we first formulate the principal’s problem in terms of indirect utilities (Section 4.1). We then state the main characterization of optimal mechanisms in Section 4.2 and illustrate the result for particular mechanisms. Finally, we outline the proof of the main result in Section 4.3.

4.1. Formulating the principal’s problem

Consider an indirect utility U that corresponds to some incentive-compatible mechanism. In general, there are many incentive-compatible mechanisms that induce the same indirect utility; however, all such mechanism induce the same payoff for the principal. To see this, let m be an incentive-compatible mechanism with corresponding indirect utility U . Using $\nabla U(s) = \mathbb{E}[m(s)]$ (by an Envelope theorem) and $U(s) = \nabla U(s) \cdot s + b(m(s))$, the principal’s payoff from mechanism m in state s is completely determined by U :

$$\mathbb{E}[m(s)] \cdot g(s) + \kappa b(m(s)) = \nabla U(s) \cdot [g(s) - \kappa s] + \kappa U(s).$$

This observation implies that the principal’s payoff is a linear function of U . Therefore, a solution to the principal’s problem can be found at an extreme point of the feasible set. Returning to Figure 2, it is easy to see that the indirect utilities in Figures 2a–2c are extremal in that they cannot be written as a nontrivial convex combination of two feasible indirect utilities. In contrast, the indirect utility in Figure 2d can be written as such a convex combination. This implies that whenever this mechanism is optimal, there is another (and simpler) mechanism which is also optimal. Intuively, one can write this indirect utility as a convex combination because two adjacent regions obtain distinct stochastic actions. This insight shows how, without loss of optimality, one can restrict attention to a smaller class of mechanism.⁷ For multidimensional settings, analogous arguments show that many complicated mechanisms are not extremal and therefore the principal need not consider these mechanisms.

As is standard in multidimensional mechanism design (see, for example, Rochet and

⁷Kleiner, Moldovanu, and Strack, 2021 develop this point more formally in the context of one-dimensional types/actions and quadratic utilities and characterize the set of extremal mechanisms. Formulating the problem in terms of indirect utilities and using our Lemma 1, one can obtain this characterization more directly. It would be interesting to extend the characterization of extremal mechanisms to the multidimensional setting.

Choné, 1998), we can use the divergence theorem to reformulate the objective function as follows:

$$\begin{aligned} & \int [\kappa U(s) + \nabla U(s) \cdot [g(s) - \kappa s]] dF(s) \\ &= \int U(s) [\kappa f(s) - \operatorname{div}[(g(s) - \kappa s)f(s)]] ds + \int_{\operatorname{bd} S} U(s)[g(s) - \kappa s]f(s) \cdot \hat{\mathbf{n}}_S(s) d\mathcal{H}(s), \end{aligned}$$

where div denotes the divergence of a function, for any set A its boundary is denoted by $\operatorname{bd} A$, \mathcal{H} denotes the $n - 1$ -dimensional Hausdorff measure on the boundary of S , and $\hat{\mathbf{n}}_S(s)$ denotes the outward normal vector to the convex set S at $s \in \operatorname{bd} S$.⁸

This allows us to write the principal's problem as⁹

$$\begin{aligned} & \max_{U \text{ convex}} \int U(s) d\mu(s) \\ & \text{s. t. } U \leq h, \end{aligned}$$

where the measure μ is defined by

$$\mu(E) = \int_E \nu(s) d\lambda(s),$$

λ is the Lebesgue measure on S plus the Hausdorff measure on the boundary of S ,¹⁰ and

$$\nu(s) := \begin{cases} \kappa f(s) - \operatorname{div}[(g(s) - \kappa s)f(s)] & \text{if } s \in \operatorname{int} S \\ [g(s) - \kappa s]f(s) \cdot \hat{\mathbf{n}}_S(s) & \text{if } s \in \operatorname{bd} S. \end{cases}$$

Clearly, ν plays an important role in determining which mechanisms are optimal. Heuristically, $\nu(s)$ measures how much the principal's payoff increases if the indirect utility of type s is increased, but where types on the boundary get extra weight.

To illustrate ν and for later use, let us compute ν for a one-dimensional type space

⁸ Since U , g , and f are bounded and Lipschitz continuous functions on the compact and convex set S , all requirements of the divergence theorem in Pfeffer (1991, Theorem 5.19) are satisfied and his result implies $\int_{\operatorname{bd} S} f(s)U(s)[g(s) - \kappa s] \cdot \hat{\mathbf{n}}_S(s) d\mathcal{H}(s) = \int_S \operatorname{div} [f(s)U(s)[g(s) - \kappa s]] ds$. Using the definition of the divergence and rearranging terms yields our expression.

⁹ The existence of a maximizer follows from standard arguments.

¹⁰ That is, $\lambda(E) = \mu(E) + \mathcal{H}(E \cap \operatorname{bd} S)$ for any measurable set E .

$S = [\underline{s}, \bar{s}]$:

$$\nu(s) := \begin{cases} \kappa f(s) - [g'(s) - \kappa]f(s) - [g(s) - \kappa s]f'(s) & \text{if } s \in (\underline{s}, \bar{s}) \\ [g(\bar{s}) - \kappa \bar{s}]f(\bar{s}) & \text{if } s = \bar{s} \\ [\kappa \underline{s} - g(\underline{s})]f(\underline{s}) & \text{if } s = \underline{s}. \end{cases} \quad (1)$$

Example 1. Suppose $S = [-\frac{1}{2}, \frac{1}{2}]^n$ and F is the uniform distribution on S . Let us assume payoffs are quadratic and that $g(s) = \alpha s$ for some $\alpha \in (0, \kappa]$; this implies that the principal is biased towards the ex-ante optimal action 0. In that case, ν simplifies to

$$\nu(s) := \begin{cases} \kappa + (\kappa - \alpha)n & \text{if } s \in \text{int } S \\ (\alpha - \kappa)s \cdot \hat{\mathbf{n}}_S(s) & \text{if } s \in \text{bd } S. \end{cases} \quad (2)$$

4.2. Optimal mechanisms

Given an indirect utility U , we let \mathcal{Q} denote a coarsest partition of \mathbb{R}^n such that U is affine on each partition element. We denote by $\{\mu|_Q\}_{Q \in \mathcal{Q}}$ a conditional measure of μ given Q .

Theorem 1. *Let U be a feasible indirect utility. Then U is optimal if for a.e. $Q \in \mathcal{Q}$, $\mu|_Q(Q) \geq 0$ and $\mu|_Q \leq_{cx} \delta_Q$, where δ_Q is a point mass of mass $\mu|_Q(Q)$ at s if there is $s \in Q$ satisfying $U(s) = h(s)$ and δ_Q is the zero-measure otherwise.*

Moreover, this condition is necessary for U to be optimal if U is differentiable $|\mu|$ -almost everywhere.

Two comments on the conditions in Theorem 1 are in order. First, if there is $s \in Q$ such that $U(s) = h(s)$ then it is unique because U is affine on Q and h is strictly convex. Second, for the necessity result, observe that U is differentiable Lebesgue-almost everywhere since it is a convex function. Since $|\mu|$ is absolutely continuous with respect to the Lebesgue measure on the interior of S , U is differentiable $|\mu|$ -almost everywhere if, for example, the density f is zero on the boundary of S or if U is differentiable \mathcal{H} -almost everywhere on the boundary of S . In the one-dimensional case, this last condition can always be satisfied.

Why are the conditions in Theorem 1 sufficient for U to be optimal? Consider a partition element $Q \in \mathcal{Q}$ and suppose δ_Q is a point mass at s^* . Then any feasible indirect utility V will be convex and lie below U at s^* . Also, $\mu|_Q \leq_{cx} \delta_Q$ implies $\int V(s) d\mu|_Q(s) \leq \int V(s) d\delta_Q(s)$. Moreover, if a is an affine function that coincides with V at the barycenter of $\mu|_Q$ then we

get $\int V(s) d\mu|_Q(s) \leq \int a(s) d\mu|_Q(s)$. Since U restricted to Q is affine, lies above V at s^* , and $\mu|_Q(Q) \geq 0$, this implies that conditional on the type belonging to Q , the principal's expected payoff under U is higher than under V . And if δ_Q is the zero measure then V might lie above U but the same conclusion follows since $\mu|_Q(Q) = 0$ and conditional on Q , adding a constant to the indirect utility does not change the principal's payoff. The conclusion that these conditions are also essentially necessary shows that the problem can in some sense be decomposed: whenever the principal can improve U conditional on Q , she can extend this improved version to a feasible indirect utility that yields unconditionally a higher payoff.

A particularly simple mechanism is if the principal delegates the decision to the agent, potentially restricting the agents action to belong to some set A . Note that any deterministic mechanism can be implemented as an indirect mechanism in this way. For a closed set $A \subseteq S$, we say that *delegating to A* is optimal if an optimal mechanism takes the form that any type in A gets her first-best action, and any other type gets her most preferred action among the first-best actions of types in A . For example, if $n = 1$ and $A = [s_1, s_2]$ then delegating to A is optimal if there is an optimal mechanism in which any type below s_1 gets the first-best action of type s_1 , any type in $[s_1, s_2]$ gets her first best action, and any type above s_2 gets the first-best action of type s_2 . In the following we will specialize Theorem 1 and discuss under what conditions such a mechanism is optimal.

We can simplify the conditions in Theorem 1 by recalling that the convex order has a simple structure for one-dimensional spaces. A cdf H_1 on a one-dimensional interval $[x, y]$ dominates a cdf H_2 in the convex order if and only if H_2 *majorizes* H_1 :

$$\int_s^y H_1(z) dz \leq \int_s^y H_2(z) dz$$

for all $s \in [x, y]$ with equality for $s = x$ (Shaked and Shanthikumar, 2007, Theorem 3.A.1). This observation simplifies the characterization in Theorem 1 whenever U is affine on at most one-dimensional sets. As we will see, this is useful even if the type space is multidimensional. To illustrate the simpler conditions, we first consider when interval delegation is optimal with one-dimensional types (for earlier characterizations, see Alonso and Matouschek, 2008; Amador and Bagwell, 2013).

Corollary 1. *Suppose $n = 1$ and $s_1, s_2 \in S$ with $s_1 < s_2$. Delegating to the interval $[s_1, s_2]$ is optimal if and only if*

$$(i) \nu(s) \geq 0 \text{ for all } s \in [s_1, s_2],$$




Figure 3: Optimality of interval delegation

- (ii) $\int_{\bar{s}}^{\infty} (x - s)\nu(x) d\lambda(x|x \geq s_2) \leq 0$ for all $s \geq s_2$ with equality for $s = s_2$, and
- (iii) $\int_{\underline{s}}^s (s - x)\nu(x) d\lambda(x|x \leq s_1) \leq 0$ for all $s \leq s_1$ with equality for $s = s_1$.

The corollary follows from Theorem 1 by rewriting the convex order constraint using majorization (for the following heuristic discussion, we drop any ‘almost everywhere’-qualifiers). The partition \mathcal{Q} that is induced by delegating to the interval $[s_1, s_2]$ contains the singletons $\{s\}$ for all $s \in (s_1, s_2)$ and the additional partition elements $(-\infty, s_1]$ and $[s_2, \infty)$. For any $Q = \{s\}$ with $s \in (s_1, s_2)$, Condition (i) is equivalent to $\mu|_Q(Q) \geq 0$. And since $\mu|_Q$ is a point mass, $\mu|_Q \leq_{cx} \delta_Q$ is trivially satisfied.

For $Q = [s_2, \infty)$, $\mu|_Q \leq_{cx} \delta_Q$ can be rewritten as Condition (ii) using the above-mentioned formulation of the convex order in terms of majorization. Moreover, $\mu|_Q(Q) \geq 0$ follows from Condition (ii) after observing that the derivative of the left-hand side with respect to s and evaluated at s_2 is negative (since the left-hand side equals zero for $s = s_2$ and is negative for $s \geq s_2$) and equals $\int_{s_2}^{\bar{s}} -\nu(x) d\lambda(x|x \geq s_2) = -\mu_Q(Q)$. An analogous argument for $Q = (-\infty, s_1]$ establishes the result.

Figure 3 illustrates Condition (ii) of Corollary 1. Suppose that starting with interval delegation (represented by the solid indirect utility), the principal changes the mechanism and assigns a lottery with expected value strictly above $\nabla U(s_2)$ to all types above s . This tilts the indirect utility starting at s upwards (see the dashed indirect utility) and therefore increases the indirect utility for every type $x \geq s$ in proportion to $x - s$. The change in the principal’s expected payoff is therefore proportional to $\int_{\bar{s}}^{\infty} (x - s)\nu(x) d\lambda(x|x \geq s_2)$. Consequently, condition (ii) ensures that such changes are not profitable. Equality for $s = s_2$ implies, in addition, that it would not be profitable to marginally reduce the action for all

types above s_2 either.

Interestingly, the conditions identified in Corollary 1 are in our setting equivalent to the ones obtained in Amador and Bagwell (2013, Proposition 2a). This might initially be surprising since we characterize optimality of interval delegation in the class of stochastic mechanisms and Amador and Bagwell characterize optimality in the class of deterministic mechanisms (and stochastic mechanisms can do strictly better in general). Figure 3 illustrates why the conditions are the same: Suppose the principal strictly benefits from deviating to the dashed indirect utility, which represents a stochastic mechanism. Since her payoff is linear in U , the arguments in the previous paragraph imply that she also benefits from deviating to the dotted indirect utility. Since the dotted linear utility corresponds to a deterministic mechanism, we conclude that conditions (ii) in (iii) in Corollary 1 are necessary for interval delegation to be optimal in the class of deterministic mechanisms (and necessity of condition (i) can be shown easily). Later, it will become clear that this equivalence is specific to the one-dimensional setting.

Corollary 2. *If $n = 1$ and $\{s \in S : \nu(s) \geq 0\}$ is an interval, then delegating to an interval is optimal.*

The key insight for this result is that any pooling region (i.e., any Q such that $Q \cap S$ is not a singleton) must contain types s with $\nu(s) \geq 0$ (since $\mu|_Q(Q) \geq 0$) and types s with $\nu(s) \leq 0$ (since no point measure δ_Q can dominate a distinct positive measure in the convex order). If ν is positive on an interval, it follows that there can be at most two pooling regions. A simple argument then shows that delegating to an interval is an optimal mechanism.

Corollary 2 extends Proposition 2(a) in Amador, Bagwell, and Frankel (2018), which in our notation requires ν to be positive on (\underline{s}, \bar{s}) . A simple implication of our result is the following, which can be useful for applications.

Corollary 3. *Suppose the type space is one-dimensional (i.e., $n = 1$), $\kappa = 1$, and the agent has a constant bias (i.e., $g(s) = s + \beta$ for some $\beta \in \mathbb{R}$). If f is logconcave then delegating to an interval is optimal.*

As another illustration, let us return to Example 1 specializing to a one-dimensional type space.

Example 1 (continued). *For $n = 1$ and $\kappa = 1$, the objective function simplifies to $\nu(s) = 2\kappa - \alpha$ for $s \in (\underline{s}, \bar{s})$, $\nu(\underline{s}) = (\kappa - \alpha)\underline{s}$, and $\nu(\bar{s}) = (\alpha - \kappa)\bar{s}$. Since ν is positive on an interval,*




Figure 4: Indirect utility for delegation to a convex set.

Corollary 2 implies that delegating to an interval is optimal, and it only remains to find the best interval.

The optimal interval must satisfy Condition (ii) as an equality for $s = s_2$, which requires

$$(2 - \alpha) \left[\frac{1}{8} - \frac{1}{2}s_2^2 - s_2 \left(\frac{1}{2} - s_2 \right) \right] = 0,$$

and simple algebra yields $s_2 = \frac{\alpha}{2-\alpha}$. Using symmetry, it follows that it is optimal to delegate to the interval $\left[-\frac{\alpha}{2-\alpha}, \frac{\alpha}{2-\alpha} \right]$.

For a one-dimensional type space, the approach used in Corollary 1 can be used to simplify the conditions in Theorem 1 for any mechanism, not just interval delegation. More generally, this approach is useful even with multidimensional types. To see this, let A be a closed and convex set and, for $s \in \text{bd } A$, let $N_A(s)$ denote the normal cone to A at s . With quadratic payoffs, if the principal delegates A and $s \in \text{bd } A$, then all types in $s + N_A(s)$ will choose action s . Moreover, if the boundary of A is differentiable then $N_A(s)$ is a (one-dimensional) ray and we can again use majorization to simplify the convex dominance conditions in Theorem 1.

Corollary 4. *Suppose payoffs are quadratic and $A \subseteq S$ is closed, convex, has nonempty interior and a differentiable boundary. Delegating to A is optimal if and only if*

- (i) $\nu(s) \geq 0$ for all $s \in A$ and

(ii) for all $s \in \text{bd } A$ and $z > 0$,

$$\int_z^\infty (x - z)\nu(s + x\hat{\mathbf{n}}_A(s)) \, d\lambda(s + x\hat{\mathbf{n}}_A(s)|s + N_A(s)) \leq 0$$

with equality for $z = 0$.

The conditions in Corollary 4 closely resemble those in Corollary 1. Indeed, Condition (i) in either case requires that ν is positive on the set of types that obtain their first-best payoffs, and Condition (ii) (and Conditions (ii) and (iii), respectively) impose that for each point on the boundary the analogous stochastic dominance condition holds.

The economic interpretation of Condition (ii) is analogous to how we interpreted Condition (ii) in Corollary 1. This condition ensures that the principal does not benefit from marginally tilting the indirect utility along line segments that are orthogonal to the boundary of A , e.g., the solid line segment in Figure 4. Observe that there is a stochastic mechanism in which the indirect utility is increased only in a small neighborhood of the solid line segment (by Lemma 1). On the other hand, there is no deterministic mechanism achieving this because for any deterministic action the indirect utility would have to increase significantly along the solid line segment (in order to reach the first-best payoff for some type) and convexity then requires that all types in a neighborhood of the line segment obtain higher indirect utilities. This indicates that our characterization relies in the multidimensional setting on stochastic mechanisms being feasible.

Example 1 (continued). *Consider a two-dimensional example and recall that F is the uniform distribution and $g(s) = \alpha s$ for some $\alpha \in [0, \kappa)$. We assume quadratic payoffs; in that case, the problem is separable across dimensions: the principal's optimal action in dimension 1 depends only on the first component of the state and is independent of the action in dimension 2.*

Suppose first that there are two agents: For $i = 1, 2$, agent i has private information about s_i (but not s_j for $j \neq i$) and cares only about the action and state in dimension i . It follows that the principal faces two independent delegation problems, and our earlier analysis implies that it is optimal to let each agent choose any action in $\left[-\frac{\alpha}{2-\alpha}, \frac{\alpha}{2-\alpha}\right]$. In effect, the agents' choice will be the action in the red square in Figure 5 that is closed to the realized state.

Now compare this to the situation where there is only one agent. This agent has private information about both dimensions of the state and cares about both dimensions of the




Figure 5: Optimal bundling

action. How can the principal improve her expected payoff? Intuitively, she could offer the agent to take more extreme actions in one dimension if he moderates his action in the other dimension. How can the principal optimally bundle the two decision problems?

Corollary 4 provides insights into how to solve the problem: if one can find an A satisfying the conditions stated there, delegating to this set will be an optimal mechanism. Since ν is positive on the interior of S and strictly negative on the boundary of S , Condition (i) will be satisfied if $A \subseteq \text{int } S$ and Condition (ii) will be satisfied if, for every $s \in \text{bd } A$, equality holds in Condition (ii) for $z = 0$. This yields a second-order differential equation, whose solution describes the boundary of the optimal delegation set, see the blue curve in Figure 5 for an illustration.

4.3. Proof Sketch

To prove Theorem 1, we use duality in linear programming. To formulate the dual program, it is more convenient to work with indirect utilities that are defined on a compact domain. But recall that it is not enough in Lemma 1 to only require that $U(s) \leq h(s)$ for all $s \in S$. The following technical result ensures that we can restrict the indirect utilities to have a compact domain as long as this domain is chosen large enough.

Lemma 2. *There is a compact $X \subseteq \mathbb{R}^n$ such that the principal's problem can be written as $\max\{\int U d\mu | U : X \rightarrow \mathbb{R}, U \text{ convex}, U \leq h\}$.*

Formally, we show that if X is chosen large enough then for any solution to the above problem there is a corresponding solution to the original problem. For a convex function U defined on S , we consider the smallest convex function defined on \mathbb{R}^n that extends U . If this extension lies below h on a large set X then $h(y) < U(y)$ for some y is possible only if $\|\nabla U(s)\|$ is large for some $s \in S$, i.e., the expected action for some type is large. We show that this implies that the principal's expected payoff is low, contradicting that U is a solution.

Now let X be as in the above lemma and denote by \mathcal{U} the set of convex continuous functions that map X to \mathbb{R} and by \mathcal{M}_+ the set of positive measures on X . We can formulate the principal's problem as follows (and call this formulation the primal problem):

$$\begin{aligned} & \max_{U \in \mathcal{U}} \int U(s) d\mu(s) & (\text{P}) \\ & \text{s. t. } U \leq h \end{aligned}$$

The dual problem We will show that the following problem is the dual problem:

$$\begin{aligned} & \inf_{\gamma \in \mathcal{M}_+} \int h(s) d\gamma(s) & (\text{D}) \\ & \text{s. t. } \gamma \geq_{cx} \mu, \end{aligned}$$

where \geq_{cx} denotes the convex order on the space of measures.

Note that h is a convex function; therefore, if μ was a positive measure, this would be a trivial problem with solution $\gamma = \mu$. However, since μ is a signed measure and γ has to be a positive measure, μ is not feasible in general.

It is easy to see that weak duality holds, i.e., that the value of the primal problem (P) is always below the value of the dual problem (D). Indeed, for any feasible U and γ ,

$$\int U(s) d\mu(s) \underbrace{\leq}_{(i)} \int U(s) d\gamma(s) \underbrace{\leq}_{(ii)} \int h(s) d\gamma(s) \quad (3)$$

since (i) U is convex and $\mu \leq_{cx} \gamma$ and (ii) γ is a positive measure and $U \leq h$. The following result shows that strong duality holds, that is the optimal values of both problems are equal and the dual problem has a solution.

Lemma 3 (Strong duality). *A feasible mechanism U is optimal if and only if there exists a positive measure $\gamma \geq_{cx} \mu$ such that*

$$U(s) = h(s) \text{ for } \gamma\text{-almost every } v \tag{4}$$

$$\int U(s) d\mu(s) = \int U(s) d\gamma(s). \tag{5}$$

This result is an analogue of a result in the revenue-maximization problem of a multi-product monopolist (see Theorem 2 in Daskalakis, Deckelbaum, and Tzamos, 2017). Our formulation of the delegation problem allows us to easily deduce strong duality. Note that there is a convex function U such that $h(x) - U(x) > 0$ for all $x \in X$. Therefore, Slater's condition is satisfied and standard results from linear programming imply that the dual problem has a solution and that the optimal solutions of the primal and dual problems achieve the same value. Since both inequalities in (3) have to hold as equalities, Lemma 3 follows.

Proof idea for Theorem 1. It is easy to show that the conditions in Theorem 1 imply that U is optimal: by aggregating the measures δ_Q , one obtains a positive measure γ satisfying the complementary slackness conditions (4) and (5) and $\gamma \geq_{cx} \mu$. Lemma 3 then implies that U is optimal.

For the converse direction, suppose U is optimal. By Lemma 3, there is a positive measure γ such that the complementary slackness conditions (4) and (5) hold and $\gamma \geq_{cx} \mu$. Letting μ^+ (μ^-) denote the positive (negative) part of μ , this last condition is equivalent to $\gamma + \mu^- \geq_{cx} \mu^+$. Strassen's theorem then implies that $\gamma + \mu^-$ is a mean-preserving spread of μ^+ : one can obtain the measure $\gamma + \mu^-$ by taking, for every s , the mass μ^+ puts on s and spreading it according to a probability measure D_s with expected value s . Since U is convex, Jensen's inequality implies that $U(s) \leq \int U(x) dD_s(x)$ and equality holds only if U is affine on the convex hull of the support of D_s . Since equality must hold by (5), we obtain that for all $Q \in \mathcal{Q}$ and $s \in Q$, the support of D_s is contained in the closure of Q . To simplify this informal discussion, suppose that for all $Q \in \mathcal{Q}$ and $s \in Q$, the support of D_s is actually contained in Q (and not just the closure of Q) and consider a partition element Q of positive measure. Then the conditional measure $\gamma|_Q$ is positive (since γ is positive)

and satisfies $\gamma|_Q + \mu^-|_Q \geq_{cx} \mu^+|_Q$ (since the left-hand side is a mean-preserving spread of the right-hand side). Moreover, by (5) we get $U(s) = h(s)$ for every s in the support of $\gamma|_Q$. Since h is strictly convex and U is affine on Q , there is at most one $s \in Q$ with $U(s) = h(s)$ and therefore $\gamma|_Q$ is a point mass at this s or the zero measure. It follows that $\mu|_Q \leq_{cx} \delta_Q$, where δ_Q is a point mass at s or is the zero measure. The proof in the Appendix follows this sketch but uses the additional assumption in Theorem 1 and additional arguments to deal with the case where the support of D_s is a subset of the closure of Q but not a subset of Q .¹¹

References

- Alonso, Ricardo and Niko Matouschek (2008). Optimal Delegation, *Review of Economic Studies* 75.1, pp. 259–293. (Cit. on pp. 3, 11).
- Amador, Manuel and Kyle Bagwell (2013). The Theory of Optimal Delegation with an Application to Tariff Caps, *Econometrica* 81.4, pp. 1541–1599 (cit. on pp. 2, 3, 11, 13).
- (2020). Money burning in the theory of delegation, *Games and Economic Behavior* 121, pp. 382–412 (cit. on p. 2).
- Amador, Manuel, Kyle Bagwell, and Alex Frankel (2018). A Note on Interval Delegation, *Economic Theory Bulletin* 6.2, pp. 239–249 (cit. on p. 13).
- Ambrus, Attila and Georgy Egorov (2017). Delegation and Nonmonetary Incentives, *Journal of Economic Theory* 171, pp. 101–135 (cit. on p. 2).
- Armstrong, Mark (1995). Delegating Decision-Making to an Agent with Unknown Preferences. unpublished (cit. on p. 3).
- Bogachev, Vladimir Igorevich (2007a). Measure theory. Vol. 1. Springer (cit. on p. 25).
- (2007b). Measure theory. Vol. 2. Springer (cit. on p. 23).
- Chakraborty, Archishman and Rick Harbaugh (2007). Comparative cheap talk, *Journal of Economic Theory* 132.1, pp. 70–94 (cit. on p. 3).
- Daskalakis, Constantinos, Alan Deckelbaum, and Christos Tzamos (2017). Strong Duality for a Multiple-Good Monopolist, *Econometrica* 85.3, pp. 735–767 (cit. on pp. 4, 18).
- Dragomirescu, F and C Ivan (1992). The smallest convex extensions of a convex function, *Optimization* 24.3-4, pp. 193–206 (cit. on p. 22).
- Dunford, Nelson and Jacob T. Schwartz (1988). Linear Operators, Part 1: General Theory. John Wiley and Sons (cit. on p. 22).
- Frankel, Alex (2014). Aligned delegation, *American Economic Review* 104.1, pp. 66–83 (cit. on p. 3).
- (2016). Delegating multiple decisions, *American Economic Journal: Microeconomics* 8.4, pp. 16–53 (cit. on p. 3).

¹¹ If s lies in the closure of Q and $Q' \neq Q$, then U is not differentiable at s and, therefore, $U(s) < h(s)$. It follows from (4) that such points have measure zero under γ . The additional assumption ensures that such points also have measure zero under μ^+ and μ^- , and hence play no role.

- Galperti, Simone (2019). A theory of personal budgeting, *Theoretical Economics* 14.1, pp. 173–210 (cit. on p. 3).
- Haghpanah, Nima and Jason Hartline (2021). When is pure bundling optimal?, *The Review of Economic Studies* 88.3, pp. 1127–1156 (cit. on p. 4).
- Hiriart-Urruty, Jean-Baptiste and Claude Lemaréchal (2004). *Fundamentals of convex analysis*. Springer Science & Business Media (cit. on p. 22).
- Holmström, Bengt (1977). On Incentives and Control in Organizations. Ph.D. Thesis, Stanford University (cit. on pp. 1, 3).
- (1984). On the Theory of Delegation. In: *Bayesian Models in Economic Theory*. Ed. by M Boyer and R Kihlstrom. North Holland, pp. 115–141 (cit. on p. 3).
- Jackson, Matthew O and Hugo F Sonnenschein (2007). Overcoming incentive constraints by linking decisions, *Econometrica* 75.1, pp. 241–257 (cit. on p. 3).
- Kartik, Navin, Andreas Kleiner, and Richard Van Weelden (2021). Delegation in veto bargaining, *American Economic Review* 111.12, pp. 4046–87 (cit. on p. 2).
- Kleiner, Andreas, Benny Moldovanu, and Philipp Strack (2021). Extreme points and majorization: Economic applications, *Econometrica* 89.4, pp. 1557–1593 (cit. on pp. 2, 8).
- Koessler, Frédéric and David Martimort (2012). Optimal delegation with multi-dimensional decisions, *Journal of Economic Theory* 147.5, pp. 1850–1881 (cit. on p. 3).
- Kolotilin, Anton and Andriy Zapechelnjuk (2019). Persuasion meets Delegation. unpublished (cit. on p. 3).
- Kováč, Eugen and Tymofiy Mylovanov (2009). Stochastic Mechanisms in Settings Without Monetary Transfers: The Regular Case, *Journal of Economic Theory* 144.4, pp. 1373–1395 (cit. on p. 2).
- Lipnowski, Elliot and Doron Ravid (2020). Cheap talk with transparent motives, *Econometrica* 88.4, pp. 1631–1660 (cit. on p. 3).
- Manelli, Alejandro M and Daniel R Vincent (2006). Bundling as an optimal selling mechanism for a multiple-good monopolist, *Journal of Economic Theory* 127.1, pp. 1–35 (cit. on p. 4).
- (2007). Multidimensional mechanism design: Revenue maximization and the multiple-good monopoly, *Journal of Economic theory* 137.1, pp. 153–185 (cit. on p. 4).
- Marshall, Albert W, Ingram Olkin, and Barry C Arnold (2010). *Inequalities: theory of majorization and its applications*. Vol. 143. Springer (cit. on p. 24).
- Melumad, Nahum and Toshiyuki Shibano (1991). Communication in Settings with No Transfers, *RAND Journal of Economics* 22.2, pp. 173–198 (cit. on p. 3).
- Pfeffer, Washek F (1991). The Gauss-Green theorem, *Advances in Mathematics* 87.1, pp. 93–147 (cit. on p. 9).
- Phelps, Robert R (2001). *Lectures on Choquet’s theorem*. Springer (cit. on p. 23).
- Preston, Chris (2008). *Some Notes on Standard Borel and Related Spaces* (cit. on p. 24).
- Rochet, Jean-Charles (1987). A necessary and sufficient condition for rationalizability in a quasi-linear context, *Journal of Mathematical Economics* 16.2, pp. 191–200 (cit. on pp. 2, 4).

Rochet, Jean-Charles and Philippe Choné (1998). Ironing, sweeping, and multidimensional screening, *Econometrica*, pp. 783–826 (cit. on p. 8).

Shaked, Moshe and J George Shanthikumar (2007). Stochastic orders. Springer (cit. on p. 11).

Shapiro, Alexander (2010). On duality theory of conic linear problems. In: *Semi-Infinite Programming: Recent Advances*. Ed. by M. A. Goberna and M. A. Lopez. Springer, U.S.A. (cit. on p. 22).

A. Omitted Proofs

Proof of Lemma 2. Let B_r denote a ball of radius r around 0 and let U be a solution to $\max\{\int U d\mu | U : B_r \rightarrow \mathbb{R}, U \text{ convex}, U \leq h\}$. We will show that U can be extended to a solution to the principal's original problem. Let \tilde{U} denote the smallest convex extension to \mathbb{R}^n of the restriction of U to S (see Dragomirescu and Ivan, 1992). If \tilde{U} is not feasible for the original problem then there is $y \notin B_r$ such that $\tilde{U}(y) > h(y)$ and there is $s \in S$ such that $\tilde{U}(y) = U(s) + \nabla U(s) \cdot (y - s)$ (since \tilde{U} is the smallest convex extension). Using strong convexity of h (which follows since b has Lipschitz-continuous gradients, see Theorem E.4.2.2 in Hiriart-Urruty and Lemaréchal (2004)), one can show that $U(s) < h(s) - z(r)$, where $z(r) \rightarrow \infty$ as $r \rightarrow \infty$.¹² Then either $U(s') \leq h(s) - z(r)/2$ for all $s' \in S$ or, on a set of positive Lebesgue-measure, $\nabla U(s') \notin B_{r/c}$ for some constant $c > 0$ independent of r, s and s' . Since $\lim_{\|a\| \rightarrow \infty} b(a) = -\infty$ by assumption, this implies that in either case for r large enough, the principals payoff from U will be less than her payoff from taking the ex-ante optimal action. This contradicts our assumption that U was optimal. Hence, any solution can be extended to a solution of the original problem. Q.E.D.

Proof of Lemma 3. Let $\mathcal{C}(X)$ denote the vector space of continuous functions on X with the supremum norm and recall that its dual space is the space of (Radon) measures on X , which we denote by $\mathcal{M}(X)$. Let $\mathcal{V} := \{g \in \mathcal{C}(X) : \forall x \in V, g(x) \geq 0\}$; the polar cones of \mathcal{U} and \mathcal{V} are defined by

$$\begin{aligned} \mathcal{U}^* &:= \{\gamma \in \mathcal{M}(X) : \forall U \in \mathcal{U}, \int U d\gamma \geq 0\} \\ \mathcal{V}^* &:= \{\gamma \in \mathcal{M}(X) : \forall g \in \mathcal{V}, \int g d\gamma \geq 0\}. \end{aligned}$$

The principal's problem can be written as $\max_{U \in \mathcal{U}} \int U d\mu$ subject to $h - U \in \mathcal{V}$. This is a conical linear program and its dual is $\inf_{\gamma \in \mathcal{V}^*} \int h d\gamma$ subject to $\mu - \gamma \in \mathcal{U}^*$ (e.g., Shapiro, 2010). Since $\mathcal{V}^* = \mathcal{M}_+(X)$ by the Riesz representation theorem (Dunford and Schwartz, 1988, p. 265) and $\mu - \gamma \in \mathcal{U}^*$ is equivalent to $\mu \geq_{cx} \gamma$, (D) is the dual problem.

Since there is $U \in \mathcal{U}$ such that $h - U$ is in the interior of \mathcal{V} , Slater's condition is satisfied and standard results imply that strong duality holds (e.g., Shapiro, 2010, Proposition 2.8).

¹² Let c' denote modulus of convexity of h . Then, for all $y \in B_r$ that lie on the line segment from s to x , and all $t \in \partial(h - \tilde{U})(y)$, $h(s) - \tilde{U}(s) \geq [h(y) - \tilde{U}(y)] + t \cdot (s - y) + \frac{c'}{2} \|y - s\|^2$. Since $\tilde{U}(s) = U(s)$ and the first two terms of the RHS are positive, the claim follows.

Given (3), it follows that U is optimal if and only if there is a positive measure $\gamma \geq_{cx} \mu$ such that $\int U d\mu = \int U d\gamma$ and $\int U d\gamma = \int h d\gamma$, which implies the result. Q.E.D.

Proof of Theorem 1. Given $s \in X$, we denote by $Q(s)$ the partition element of \mathcal{Q} that contains s .

Sufficiency: Let $\gamma := \int \delta_{Q(s)} d|\mu|(s)$. Given the properties of $\mu|_Q$, we conclude that $\gamma \in \mathcal{M}_+$ and $\text{supp } \gamma \subseteq \{s : U(s) = h(s)\}$. Moreover, for all $c \in \mathcal{U}$,

$$\int c(x) d\gamma(x) = \int \int c(x) d\delta_{Q(s)}(x) d|\mu|(s) \geq \int \int c(x) d\mu|_{Q(s)} d|\mu|(s) = \int c(x) d\mu(x).$$

and equality holds for $c \equiv U$ because (i) U is affine on each $Q \in \mathcal{Q}$ and (ii) $\delta_Q \geq_{cx} \mu|_Q$ implies $\int a(x) d\delta_Q = \int a(x) d\mu|_Q$ for any affine function $a \in \mathcal{C}(X)$. Therefore, γ is feasible for the dual problem and satisfies the complementary slackness conditions (4) and (5). We conclude that U is optimal.

Necessity: By Lemma 3, U is optimal if and only if there is $\gamma \in \mathcal{M}_+$ satisfying (4), (5), and $\gamma \geq_{cx} \mu$. Letting μ^+ and μ^- denote the positive and negative parts of μ , respectively, the last condition is equivalent to $\gamma + \mu^- \geq_{cx} \mu^+$. Since both sides of the inequality are positive measures, Strassen's theorem (see, for example, Phelps, 2001, p. 93-94) implies that there is a dilation D_s (that is, for each s , D_s is a probability measure with barycenter s) satisfying $\gamma + \mu^- = \int D_s d\mu^+(s)$.

Let $\mu|_Q$ be a (regular, proper) system of conditional measures (such conditional measures exist by Example 10.4.11 in Bogachev, 2007b), which by definition satisfies

$$\int_X c(s) d\mu(s) = \int_X \int_X c(y) d\mu|_{Q(s)}(y) d|\mu|(s)$$

for all $c \in \mathcal{C}(X)$. Letting $\alpha_Q := \int D_s d\mu|_Q^+(s) - \mu|_Q^-$, we claim that there is $\mathcal{Q}' \subseteq \mathcal{Q}$ such that \mathcal{Q}' has $|\mu|$ -measure 0 and, for all $Q \in \mathcal{Q} \setminus \mathcal{Q}'$, α_Q is a positive measure and its support is a subset of $Q \cap \{s : U(s) = h(s)\}$.

Before proving this claim, we show that it implies the necessity result: From the definition of α_Q it follows that if α_Q is a positive measure then $\mu|_Q(Q) \geq \alpha_Q(Q) \geq 0$. Also, $\alpha_Q \geq_{cx} \mu|_Q$ since D_s is a dilation. Moreover, if α_Q has support in $Q \cap \{s : U(s) = h(s)\}$ then α_Q is either a point mass at $Q \cap \{s : U(s) = h(s)\}$ or the zero measure, which implies the result.

To prove the claim, we show first that the support of α_Q is a subset of the closure of $Q \cap \{s : U(s) = h(s)\}$ for $|\mu|$ -a.e. Q . To obtain a contradiction, suppose there is $\mathcal{Q}' \subseteq \mathcal{Q}$

with strictly positive $|\mu|$ -measure such that, for all $Q' \in \mathcal{Q}'$, the support of $\alpha_{Q'}$ is not a subset of the closure of Q . Fix arbitrary $Q' \in \mathcal{Q}'$. Since the support of $\alpha_{Q'}$ is not contained in the closure of Q' , there is a set $A \subseteq Q'$ of strictly positive $\mu|_{Q'}^+$ -measure such that, for all $x \in A$, the support of D_x is not contained in the closure of Q' . Since Jensen's inequality is strict whenever the convex function is not affine on the convex hull of the support (Marshall, Olkin, and Arnold, 2010, Proposition 16.C.1), we obtain

$$\int U(s) d\mu|_{Q'}^+(s) < \int \left[\int U(x) dD_s(x) \right] d\mu|_{Q'}^+(s).$$

This yields

$$\begin{aligned} \int U(s) d\mu(s) &= \int \left[\int U(x) d\mu|_{Q(s)}^+(x) - \int U(x) d\mu|_{Q(s)}^-(x) \right] d|\mu|(s) \\ &< \int \left[\int \left(\int U(y) dD_x(y) \right) d\mu|_{Q(s)}^+(x) - \int U(x) d\mu|_{Q(s)}^-(x) \right] d|\mu|(s) \\ &= \int \int U(y) dD_s(y) d\mu^+(s) - \int U(x) d\mu^-(s) \\ &= \int U(s) d\gamma(s), \end{aligned}$$

which contradicts (5). We conclude that, except possibly on a $|\mu|$ -Null set, the support of α_Q is a subset of the closure of Q .

Second, we show that $\alpha_{Q(s)}$ is a positive measure for $|\mu|$ -almost every s : Let

$$B := \{s \in X : s \in \text{cl } Q \cap \text{cl } Q' \text{ for } Q \neq Q'\},$$

and note that for any $s \in B$, U is not differentiable at s and therefore $U(s) < h(s)$. Since $\text{supp } \gamma \subseteq \{s : U(s) = h(s)\}$ by (4), $\gamma(B) = 0$. Moreover, $\mu^-(B) = 0$ because U is continuously differentiable $|\mu|$ -almost everywhere by assumption. Let \mathcal{G} denote the σ -algebra generated by \mathcal{Q} and note that the Borel σ -algebra on X is generated by some countable algebra $\{A_1, A_2, \dots\}$ (Preston, 2008, Propositions 3.1 and 3.3). For each n and $G \in \mathcal{G}$,

$$\begin{aligned} \int_G \alpha_{Q(s)}(A_n) d|\mu|(s) &= \int_G \int_X D_{s'}(A_n) d\mu|_{Q(s)}^+(s') - \mu|_{Q(s)}^-(A_n) d|\mu|(s) \\ &= \int_G D_s(A_n) d\mu^+(s) - \mu^-(A_n \cap G) \\ &\geq \left[\int_X D_s(A_n \cap G) d\mu^+(s) - \mu^-(A_n \cap G) \right] - \int_{X \setminus G} D_s(A_n \cap G) d\mu^+(s). \end{aligned}$$

The bracketed term equals $\gamma(A_n \cap G)$ and is therefore positive. The last term is zero since

$$\int_{X \setminus G} D_s(A_n \cap G) d\mu^+(s) \leq \int D_s(A_n \cap G \cap B) d\mu^+(s) - \mu^-(A_n \cap G \cap B) = \gamma(A_n \cap G \cap B) = 0$$

(recall that $\gamma(B) = \mu^-(B) = 0$). Since $\alpha_{Q(s)}(A_n)$ is \mathcal{G} -measurable in s , it follows that there is a $|\mu|$ -Null set Z_n such that $\alpha_{Q(s)}(A_n) \geq 0$ for all $s \in X \setminus Z_n$. Letting $Z := \bigcup_{n=1}^{\infty} Z_n$, for all $s \in X \setminus Z$ and Borel sets A , $\alpha_{Q(s)}(A) \geq 0$ by Caratheodory's extension theorem (see Bogachev, 2007a, Theorem 1.5.6 and the comment afterward).

Finally, if it is not true that for $|\mu|$ -almost every s , the support of $\alpha_{Q(s)}$ is a subset of $\{s : U(s) = h(s)\}$, then $\int U d\gamma < \int h d\gamma$, contradicting (4). Moreover, for any $s \in \text{cl } Q \setminus Q$, U is not differentiable at s and therefore $U(s) < h(s)$ (because $U \leq h$ and h is differentiable). We conclude that there is a collection $\mathcal{Q}' \subset \mathcal{Q}$ with $|\mu|$ -measure 0 such that, for all $Q \in \mathcal{Q} \setminus \mathcal{Q}'$, α_Q is a positive measure that has support on $Q \cap \{s : U(s) = h(s)\}$. Q.E.D.

Proof of Corollary 1. Note that the partition \mathcal{Q} induced by U has elements $(-\infty, s_1]$, $[s_2, \infty)$, and $\{s\}$ for all $s \in (s_1, s_2)$. For all $s \in (s_1, s_2)$, $\nu(s) \geq 0$ is equivalent to $\mu|_Q(Q) \geq 0$ and $\mu|_Q \leq_{cx} \delta_Q$ for $Q = \{s\}$.¹³ Now consider $Q = [s_2, \infty)$ and let $\lambda(x|x \geq s_2)$ denote the conditional distribution of λ conditional on $x \geq s_2$. Since δ_Q is a point mass of mass $\mu|_Q(Q)$ at s_2 , we can use majorization to rewrite $\mu|_Q \leq_{cx} \delta_Q$ as

$$\int_s^{\bar{s}} \int_x^{\bar{s}} \nu(z) d\lambda(z|z \geq s_2) dx \leq 0$$

for all $s \geq s_2$ with equality for $s = s_2$. Integrating by parts, this becomes condition (ii). Moreover, since the derivative with respect to s of the left-hand side of the above inequality evaluated at s_2 is negative, we obtain $\mu|_Q(Q) \geq 0$. The argument for $Q = (-\infty, s_1]$ is analogous. Q.E.D.

Proof of Corollary 2. Let U be an optimal indirect utility and \mathcal{Q} a corresponding partition. Since $\mu|_Q(Q) \geq 0$ and $\mu|_Q \leq_{cx} \delta_Q$, any pooling region¹⁴ $Q \in \mathcal{Q}$ must contain types with $\nu(s) \geq 0$ and types with $\nu(s) \leq s$.

If $\nu(\underline{s}) \geq 0$ and $\nu(\bar{s}) \geq 0$, ν is positive everywhere and the claim follows. So suppose $\nu(\underline{s}) < 0$; then there is a pooling region $Q := [x, y] \in \mathcal{Q}$ which contains \underline{s} and some s

¹³ For $s \in \{s_1, s_2\}$, if $s \in \text{int } S$ then $\nu(s) \geq 0$ follows because ν is continuous on the interior of S . And if $s \in \text{bd } S$, there is $Q \in \mathcal{Q}$ with $Q \cap S = \{s\}$ and hence $\mu|_Q(Q) \geq 0$ implies $\nu(s) \geq 0$.

¹⁴ That is, any Q such that $Q \cap S$ contains strictly more than one element.

with $\nu(s) > 0$. If $\nu(y) < 0$, then $[x, y] \subseteq Q$ must hold and the claim follows. Therefore, assume $\nu(y) \geq 0$. The measure δ_Q from Theorem 1 must be a point mass at some $z \in Q$ with $\nu(z) \geq 0$ (if δ_Q were the zero measure or a point mass at z' with $\nu(z') < 0$, then $\int x - x^* d\mu|_Q > \int x - x^* d\delta_Q$ whenever $x^* = \inf\{x : \nu(x) \geq 0\}$, which contradicts $\mu|_Q \leq_{cx} \delta_Q$). It follows that $U(z) = h(z)$.

If $\nu(\bar{s}) \geq 0$ then $\nu(s) \geq 0$ for all $s \in [z, \bar{s}]$ and delegating to $[z, \bar{s}]$ is optimal. If $\nu(\bar{s}) < 0$, repeating our previous argument implies that there is an interval $[x', y'] \in \mathcal{Q}$ which contains \bar{s} and some z' with $\nu(z') \geq 0$ and $U(z') = h(z')$. Since $\nu(s) \geq 0$ for all $s \in [z, z']$, delegating to $[z, z']$ is optimal. If $\nu(\underline{s}) \geq 0$ and $\nu(\bar{s}) < 0$, a symmetric argument applies. Q.E.D.

Proof of Corollary 3. It follows from (1) that $\nu(s) = f(s) \left[1 - \beta \frac{f'(s)}{f(s)}\right]$ for $s \in (\underline{s}, \bar{s})$. If $\beta \geq 0$ then ν is singlecrossing from below on (\underline{s}, \bar{s}) (since f is logconcave) and $\nu(\underline{s}) \leq 0$. The claim then follows from Corollary 2. Q.E.D.

Proof of Corollary 4. The corresponding indirect utility induces the partition with the following elements: for any a in the interior of A , $\{a\}$, and for any $a \in \text{bd } A$, the normal cone $N_A(a)$, which is a ray through a and orthogonal to $\text{bd } A$. For any such normal ray Q , condition (ii) is equivalent to $\mu|_Q \geq_{cx} \delta_Q$ by the same argument as in Corollary 1.

“ \Leftarrow ”: Condition (i) ensures that $\mu|_{\{a\}}$ is positive for all a in the interior of A . Since it has singleton support, $\mu|_{\{a\}} \geq_{cx} \delta_{\{a\}}$. For any normal ray Q , $\mu|_Q \geq_{cx} \delta_Q$ by condition (ii) and $\mu|_Q(Q) \geq 0$ since $\int_0^\infty \nu(s + x\hat{n}_A(s)) d\lambda(s + x\hat{n}_A(s)|ray) \geq 0$ follows from condition (ii). It follows from Theorem 1 that U is optimal.

“ \Rightarrow ”: If $\nu(a) < 0$ for some a in the interior of A then there is a subset of A with positive $|\mu|$ -measure on which ν is strictly negative, which implies $\mu|_Q(Q) < 0$ on a set of positive measure, which contradicts optimality of U . Similarly, if $\nu(a) < 0$ for some $a \in \text{bd } A$ then it can be shown that $\mu|_Q(Q) < 0$ on a set of positive measure, which contradicts optimality of U by Theorem 1.

If condition (ii) is violated, $\mu|_Q \not\geq_{cx} \delta_Q$ on a set of positive measure, which again contradicts optimality of U by Theorem 1. Q.E.D.