Delegating Multiple Decisions

Alex Frankel*
Chicago Booth

November 2011

Abstract

In a class of delegation problems over multiple decisions, it is optimal for the principal to cap the agent's choices against the direction of her bias. Geometrically, this corresponds to a half-space delegation set.

1 Introduction

Consider a principal who delegates decisionmaking authority to an informed agent. For instance, a firm assigns a manager to be in charge of certain investment decisions. The manager has information about the appropriate level of investment for each project. But due to her intrinsic preferences, her compensation package, or her personal career concerns, she would choose different levels than what the firm would want. The principal knows the agent's biases, and exerts control by specifying a set of actions from which the agent may choose. The principal is committed to abide by the agent's choice from the set.

This so-called delegation problem was introduced by Holmström (1977, 1984). It has been used to analyze the optimal set of tariff levels allowed by a trade agreement; the prices that a regulated monopolist may charge; the policies a delegated committee or political advisor may choose; or the set of penalties a judge may impose on a convicted offender. When there is a single decision to be made, Melumad and Shibano (1991), Martimort and Semenov (2006), Goltsman et al. (2007), Alonso and Matouschek (2008), Ambrus and Egorov (2009), Kovac and Mylovanov (2009), and Amador and Bagwell (2011) give conditions under which it is optimal for the principal to *cap the agent's choices against the direction of her bias.*[1]

*Many thanks to Kyle Bagwell, Eugen Kovac, Tymofiy Mylovanov, Andy Skrzypacz, and Bob Wilson for useful comments.

[1] In other recent work, Athey et al. (2005) and Amador et al. (2006) argue for similar caps in the context of dynamic models outside of the traditional delegation framework.

If the principal knows that the manager is biased towards investing too much money in a project, he is given a budget and allowed to invest any amount up to that level. Likewise a trade agreement sets a maximum tariff cap; a monopolist is given a price ceiling; a liberal committee is prohibited from choosing left-wing policies; an overly forgiving judge is required to follow mandatory minimum sentencing guidelines.

How does this problem change if there are multiple decisions – can we do better than to treat the decisions independently?[2] And if the agent has different biases for each decision, how should we incorporate these differences into the optimal delegation set?

Suppose that there are N decisions. For decision i, the agent's ideal action on the real line is shifted from the principal's by a constant λ_i. So across all of the decisions, the principal and agent ideal actions can be represented as points in $\mathbb{R}^N$, separated from each other by a bias vector $\boldsymbol{\lambda} = (\lambda_1, ...\lambda_N)$. The principal and agent face a quadratic payoff loss from taking actions away from their ideals. Finally, suppose that the ex ante distribution of the players' ideal points is iid normal; only the agent observes the realized ideal points. The main result of the paper is that the optimal delegation set allows the agent to choose actions from a half-space in $\mathbb{R}^N$, with boundary normal to the bias vector.

In other words, just as in the single-decision problem, the optimal multiple-decision delegation set *caps the agent's choices against the direction of her bias.* This cap can be implemented by giving the agent a budget from which actions can be "purchased," where the price of increasing an action is proportional to the bias on that decision. The agent is allowed to spend any amount up to the budget level.

Frankel (2010) and Chakraborty and Harbaugh (2007) consider related problems in which the agent's bias is identical across decisions. In this case, the optimal delegation set above simply caps the unweighted sum, or average, over actions. The manager has a budget of money which she can spend as she wishes across projects.

In Frankel (2010), the principal does not know the agent's bias, and only knows that the same for each. When preferences are of the form considered here, that paper argues for a budget in which the agent is forced to spend exactly to the limit. This is a "max-min" optimal contract which protects the principal against extreme agent biases. In one of the leading examples of that paper, a school worried about grade-inflating or -deflating teachers would require the teachers to conform to a predetermined class grade point average. The current paper argues that the school could do better if it knew the magnitude and direction

[2]Without solving for optimal contracts, papers such as Chakraborty and Harbaugh (2007) and Jackson and Sonnenschein (2007) show how putting together independent, identical decisions can often improve payoffs.

of a teacher's bias. Instead of having to achieve a specified class GPA, a grade-inflating teacher would be given a GPA ceiling. GPA ceilings by class are in fact commonly used, e.g. at the business schools of Stanford and the University of Chicago.

The corporate finance literature has argued that forms of credit lines – budgets implemented over time – can be components of optimal contracts to constrain agents who want to invest too much of the principal's money; see for instance DeMarzo and Sannikov (2006), DeMarzo and Fishman (2007), and Malenko (2011). In Malenko (2011) the agent has state-independent preferences and her payoff is linear in the amount of money spent. The author points out that if the agent is known to prefer certain types of projects over others, then the budget should correct for this (as in the current paper) by putting a higher "price" on the preferred investments.

Section 2 states the model and the main results, and Section 3 discusses their interpretation. The main results are proven in Section 4. Section 5 shows how the rule-of-thumb to impose caps against the direction of the bias may extend to more general environments. For instance, similar caps are approximately optimal under arbitrary distributions of states when the agent is very strongly biased. And if we fix a distribution from which states are drawn iid and we increase the number of decisions N, then payoffs from such a cap approach first-best at a rate of $\frac{1}{N}$. This is a "worst-case asymptotic optimal" rate – there are some state distributions for which no delegation sets can approach first-best at a faster rate. Such a cap also gives approximately first-best payoffs when there are many decisions taken sequentially, instead of all at once.

2 The Model and Results

A principal and agent jointly make $N < \infty$ decisions, indexed by $i = 1, ..., N$. Each decision has an exogenous underlying state $\theta_i \in \mathbb{R}$, and an action $a_i \in \mathbb{R}$ is to be taken. Writing $\boldsymbol{a} = (a_1, ..., a_N)$ and $\boldsymbol{\theta} = (\theta_1, ..., \theta_N)$, the principal and agent payoffs are, respectively,

$$V_P(\boldsymbol{a}|\boldsymbol{\theta}) = -\sum_{i=1}^{N}(a_i - \theta_i)^2$$

$$V_A(\boldsymbol{a}|\boldsymbol{\theta}) = -\sum_{i=1}^{N}(a_i - \theta_i - \lambda_i)^2$$

So the principal wants to choose action a_i to match the state θ_i, while the agent wants

a_i equal to $\theta_i + \lambda_i$. I call $\boldsymbol{\lambda} = (\lambda_1, ..., \lambda_N)$ the *bias* of the agent. Each player has quadratic losses from taking actions away from their respective ideal points of $\boldsymbol{\theta}$ and $\boldsymbol{\theta} + \boldsymbol{\lambda}$.

The principal knows the agent's bias $\boldsymbol{\lambda}$ but is uninformed about the states of the world $\boldsymbol{\theta}$. I assume that the principal's prior belief is that states are iid normally distributed, with mean normalized to 0 and variance normalized to 1. The game is as follows:

1. The principal chooses a closed *delegation set* $D \subseteq \mathbb{R}^N$.
2. The agent observes the underlying states, $\boldsymbol{\theta}$.
3. The agent chooses a vector of actions $\boldsymbol{a}$ from D to maximize her payoff $V_A(\boldsymbol{a}|\boldsymbol{\theta})$.

The principal's goal is to choose an optimal delegation set, one which maximizes his expected payoff with respect to his beliefs.

Theorem 1. If $\boldsymbol{\lambda} = \mathbf{0}$ then the optimal delegation set is $\mathbb{R}^N$. If $\boldsymbol{\lambda} \neq \mathbf{0}$, there exists $K > 0$ such that $D^* = \{\boldsymbol{a} \mid \sum \lambda_i a_i \leq K\}$ is an optimal delegation set.

The proof of the theorem, which explains how to solve for K, is given in Section 4. The value of K is a function of the magnitude of the bias $|\boldsymbol{\lambda}|$, and does not depend on the number of decisions N.

For nonzero biases, the theorem states that the optimal delegation set is a half-space. The half-space always contains the origin, the uninformed principal's ex ante preferred action. The bounding hyperplane of this half-space, $\sum_i \lambda_i a_i = K$, is normal to the bias vector $\boldsymbol{\lambda}$ and is a distance of $\frac{K}{|\boldsymbol{\lambda}|}$ from the origin at its closest point. The further is the boundary from the origin, the larger is the set of actions that the agent is allowed to take.

Proposition 1. Writing K as a function of $|\boldsymbol{\lambda}| \in (0, \infty)$, the distance $\frac{K(|\boldsymbol{\lambda}|)}{|\boldsymbol{\lambda}|}$ of the boundary of D^* from the origin:

1. strictly decreases in $|\boldsymbol{\lambda}|$,
2. approaches 0 as $|\boldsymbol{\lambda}| \to \infty$, and
3. approaches ∞ as $|\boldsymbol{\lambda}| \to 0^+$.

So the stronger is the agent's bias, the less freedom she is given. An unbiased agent is given complete freedom, and an agent with a very strong bias is only allowed to take actions on a half-space which just barely contains the principal's ex ante preferred action.

Following Kovac and Mylovanov (2009), I show in the proof that the delegation set D^* is optimal even in an extended problem in which the agent may choose lotteries over actions rather than just actions.

3 Discussion

Figure 1: The optimal delegation sets are half-spaces with boundaries normal to the agent's bias.

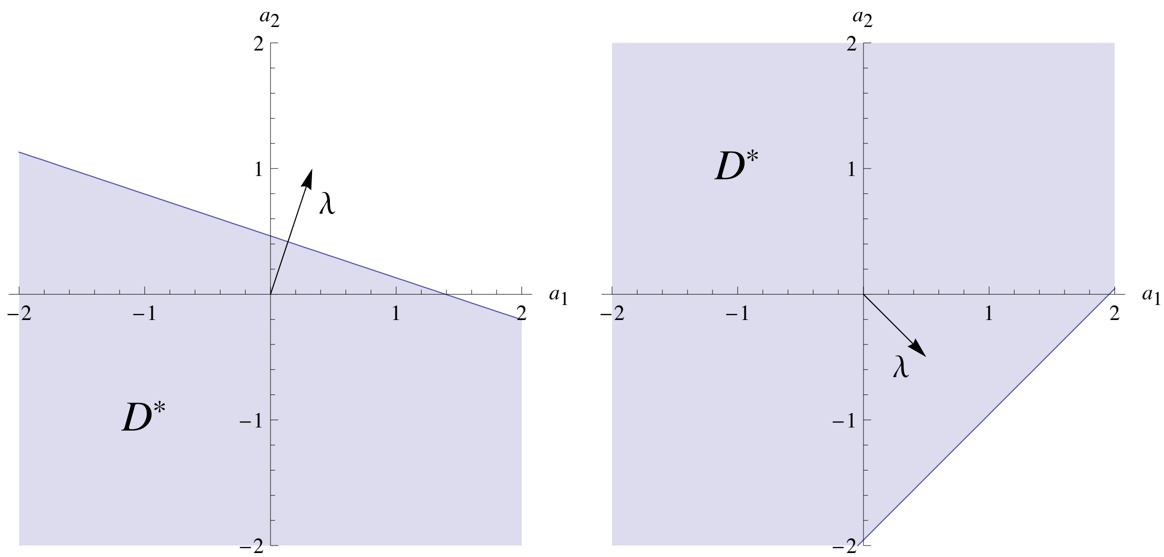

$N = 2,\ \boldsymbol{\lambda} = \left(\frac{1}{3}, 1\right)$
$D^* :\ \left\{\frac{1}{3}a_1 + a_2 \leq .464\right\}$

$N = 2,\ \boldsymbol{\lambda} = \left(\frac{1}{2}, -\frac{1}{2}\right)$
$D^* :\ \left\{\frac{1}{2}a_1 - \frac{1}{2}a_2 \leq .978\right\}$

Given a delegation set D and a state $\boldsymbol{\theta}$ in $\mathbb{R}^N$, the agent chooses the point in D closest to $\boldsymbol{\theta} + \boldsymbol{\lambda}$. The principal wants to minimize the expected distance-squared from this chosen point to $\boldsymbol{\theta}$, given that each coordinate of $\boldsymbol{\theta}$ has a standard normal distribution.

Theorem 1 states that the optimal delegation set is a half-space with boundary normal to the agent's bias $\boldsymbol{\lambda}$. We can think of $\mathbb{R}^N$ as generated by an orthogonal basis consisting of $\boldsymbol{\lambda}$ along with $N-1$ perpendicular vectors. Under the new basis, the agent has a bias of 0 in each of the $N-1$ perpendicular dimensions; she agrees completely with the principal. So the agent is given complete freedom along these dimensions of common interest. But on the dimension of disagreement parallel to $\boldsymbol{\lambda}$, the principal limits the agent's choice set. The principal caps her to prevent her from taking actions that are too extreme in the direction of the bias.

This cap helps to align incentives without the use of transfer payments. Say that the bias λ_j is positive; the agent prefers action a_j to be too high. The principal's initial response might be to cap the maximum allowed value of a_j. This corresponds to N constraints of

Figure 2: Comparing the optimal delegation set, D^*, to the optimal one with independent constraints across decisions, D^{ind}.

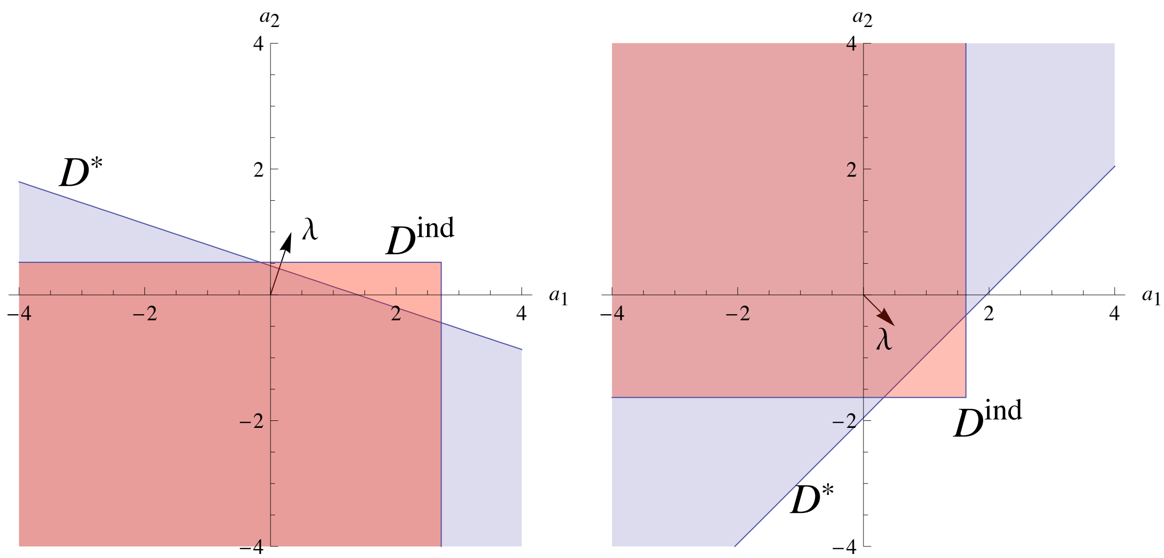

$N = 2,\ \boldsymbol{\lambda} = \left(\frac{1}{3}, 1\right)$
$\mathbb{E}[V_P|D^*] = -.678$
$\mathbb{E}[V_P|D^{ind}] = -.756$

$N = 2,\ \boldsymbol{\lambda} = \left(\frac{1}{2}, -\frac{1}{2}\right)$
$\mathbb{E}[V_P|D^*] = -.425$
$\mathbb{E}[V_P|D^{ind}] = -.489$

To give context to these payoff values, the principal's first-best outcome (from taking $\boldsymbol{a} = \boldsymbol{\theta}$) gives $V_P = 0$ while taking action $\mathbf{0}$ (an uninformed principal's ex ante first choice) gives $\mathbb{E}[V_P] = -2$, minus the sum of the variances.

In order to solve for the $\mathbb{E}[V_P]$ values in the figure, let $v^1(x)$ represent the expected payoff to the principal if the optimal delegation set is used in a single decision problem where the agent has bias $x \in \mathbb{R}$; this can be evaluated numerically. It holds that $\mathbb{E}[V_P|D^*] = v^1(|\boldsymbol{\lambda}|)$ and $\mathbb{E}[V_P|D^{\text{ind}}] = \sum_{i=1}^{N} v^1(\lambda_i)$.

the form $\lambda_i a_i \leq K_i$, one for each decision.[3] But with multiple decisions, the principal has a new tool at his disposal. He can reward the agent for reducing a_j below the cap by giving her additional freedom on her other decisions – that is, by increasing $K_{i \neq j}$. In particular, the principal links all of the decisions together through a joint constraint $\sum_i \lambda_i a_i \leq K$, corresponding to a restriction on decision j of $\lambda_j a_j \leq K_j$ for $K_j = K - \sum_{i \neq j} \lambda_i a_i$.

The principal can implement this delegation set through an artificial pricing scheme where the agent is given an amount K of resources ("delegation dollars"), and is told to spend up

[3] If we treated the decisions independently rather than having the constraints on a_j depend on the actions $\boldsymbol{a}_{-j}$, Theorem 1 applied to the $N = 1$ case tells us that the optimal policy would be to use N separate budgets.

to this level on buying a bundle of actions. The price of each unit of action a_i is set to be equal to the agent's bias along that dimension, λ_i.

If the agent has a strong positive bias on decision 1 then the price on a_1 is large and positive. Reducing a_1 by one unit saves a lot of resources, giving her much more freedom on other actions. Increasing a_1 by one unit costs a lot, restricting the other actions by a large amount. On the other hand, shifting her other actions by one unit doesn't affect the set of affordable a_1 values by very much. If the agent has a weak negative bias on decision 2 then a_2 has a small and negative price. Reducing a_2 takes away a small amount of freedom on other actions, and increasing it gives a small amount of additional freedom. Changing her other actions may affect the feasible set of a_2 values dramatically. If the agent has no bias on decision 3 then the price on a_3 is 0. Her decision along this dimension is unconstrained, no matter what other actions she chooses. And her choice of a_3 has no effect on the constraints facing other decisions.

In order to derive half-spaces as optimal out of all possible sets, I assumed that states were iid normal. This distribution of independent states is special in that it is rotationally symmetric. In Section 5, I argue that half-space delegation is an appealing benchmark for other joint distributions of states, or for more general payoff assumptions, even though it is not necessarily optimal.

4 Construction of the Optimal Delegation Set

There is no existing technique for characterizing the general properties (connectedness, convexity, etc) of an optimal multidimensional delegation set, let alone solving for the set itself. This paper contributes to the theoretical literature by showing how to derive an optimal set in a benchmark example where players have quadratic loss utilities, the agent has a constant bias, and states are iid normally distributed.

If $\boldsymbol{\lambda} = \mathbf{0}$, then the agent shares the principal's utility function and it is optimal for the principal to give the agent no constraints: the principal achieves a first-best payoff by choosing $D = \mathbb{R}^N$. For the rest of this section I suppose that $\boldsymbol{\lambda} \neq \mathbf{0}$.

I prove the theorem in two steps. First I consider an augmented game in which the principal can "cheat" and learn some information about the states before choosing D. Conditional on this extra information, the principal does at least as well as in the original problem. So the optimal payoff in the augmented game gives an upper bound on what the principal can achieve in the original game.

Then I show that there is a delegation set D which the principal can choose in the original game which exactly implements the optimal outcomes of the augmented game. Because this set achieves a theoretical upper bound on payoffs, it must be optimal.

In the augmented game, before the principal chooses a delegation set D, he learns some information about the state. In particular, he observes the projection $P_{\boldsymbol{\theta}}$ of $\boldsymbol{\theta}$ onto the hyperplane defined by $\left\{\boldsymbol{x} \in \mathbb{R}^N \middle| \sum_i \lambda_i x_i = 0\right\}$:

$$P_{\boldsymbol{\theta}} = \boldsymbol{\theta} - \frac{\sum_i \lambda_i \theta_i}{\sum_i \lambda_i^2} \boldsymbol{\lambda}$$

So in the augmented game, we add a period 0 to the original description:

0. The principal observes $P_{\boldsymbol{\theta}}$.
1. The principal chooses a closed delegation set $D \subseteq \mathbb{R}^N$.
2. The agent observes the underlying states, $\boldsymbol{\theta}$.
3. The agent chooses a vector of actions $\boldsymbol{a}$ from D to maximize her payoff $V_A(\boldsymbol{a}|\boldsymbol{\theta})$.

Learning the projection $P_{\boldsymbol{\theta}}$ is equivalent to learning on which line parallel to the agent's bias $\boldsymbol{\lambda}$ is the state $\boldsymbol{\theta}$. Denote the line parallel to $\boldsymbol{\lambda}$ intersecting $P_{\boldsymbol{\theta}}$ by $M_{\boldsymbol{\theta}}$, and parametrize points on the line by the distance $k \in \mathbb{R}$ from $P_{\boldsymbol{\theta}}$ (where $\boldsymbol{\lambda}$ is taken to be pointing in the positive direction):

$$M_{\boldsymbol{\theta}}(k) = P_{\boldsymbol{\theta}} + k \frac{\boldsymbol{\lambda}}{|\boldsymbol{\lambda}|}$$

Lemma 1. Conditional on observing $P_{\boldsymbol{\theta}}$, the principal's posterior belief on $\boldsymbol{\theta}$ is that $\boldsymbol{\theta} = M_{\boldsymbol{\theta}}(k)$ for k distributed normally with mean 0 and variance 1.

Proof. This is seen most easily by observing that, conditional on any known $\boldsymbol{\theta}_{-i}$, the coordinate θ_i has a standard normal distribution. But by the spherical symmetry of the distribution of $\boldsymbol{\theta}$, the principal's posterior belief on $\boldsymbol{\theta}$ conditional on knowing that it is on any line $M_{\boldsymbol{\theta}}$ must have a standard normal distribution on that line. □

Lemma 2. Given $P_{\boldsymbol{\theta}}$, the set $D^{\text{aug}}(P_{\boldsymbol{\theta}}) \equiv \left\{M_{\boldsymbol{\theta}}(k) \middle| k \in (-\infty, \overline{k}]\right\}$ defines an optimal delegation set in the augmented game, where $\overline{k} \geq 0$ is the unique solution to $\mathbb{E}_{\theta \sim \mathcal{N}(0,1)}\left[\theta \middle| \theta \geq \overline{k} - |\lambda|\right] = \overline{k}$.

That is, if the principal knows that $\boldsymbol{\theta}$ lies on the line $M_{\boldsymbol{\theta}}$, he only allows the agent to pick actions from a ray contained in this line. The ray is capped against the direction of the agent's bias.

Figure 3: The Augmented Game

Lemma 1

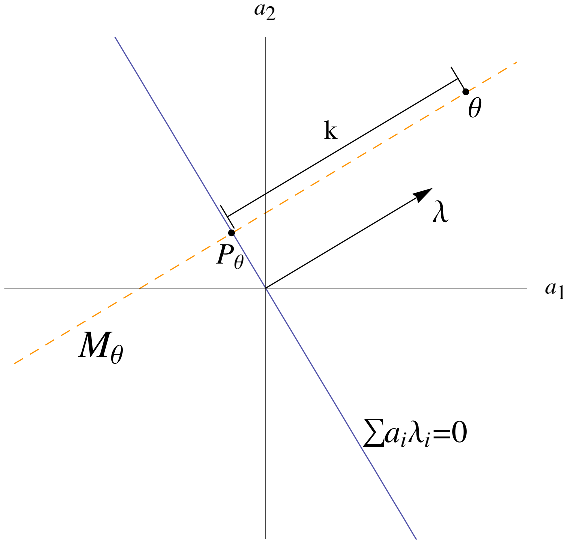

In the augmented game, before the principal chooses a delegation set he observes $P_{\boldsymbol{\theta}}$, the projection of $\boldsymbol{\theta}$ onto the hyperplane defined by $\sum_i a_i\lambda_i = 0$. That is, he learns the line $M_{\boldsymbol{\theta}}$ parallel to $\boldsymbol{\lambda}$ on which $\boldsymbol{\theta}$ lies. The principal's posterior is that the distance k from $P_{\boldsymbol{\theta}}$ to $\boldsymbol{\theta}$ has a standard normal distribution.

Lemma 2

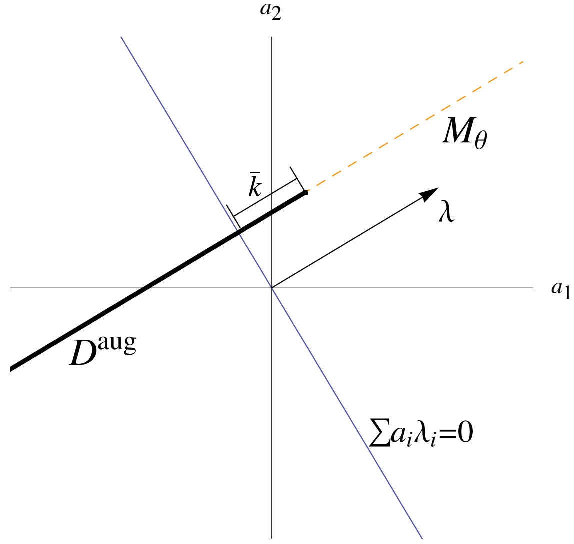

Conditional on the information that $\boldsymbol{\theta}$ lies on $M_{\boldsymbol{\theta}}$, the thick black line D^{aug} is an optimal delegation set. D^{aug} is the subset of $M_{\boldsymbol{\theta}}$ for which the distance k from the hyperplane $\sum_i a_i\lambda_i = 0$ is in the range $(-\infty, \bar{k}]$.

Lemma 3

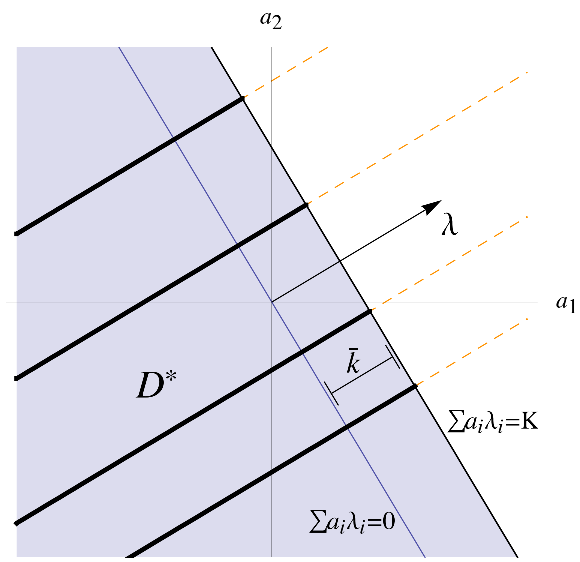

Each D^{aug} set is a translation of any other in a direction normal to $\boldsymbol{\lambda}$. Taking the union of D^{aug} sets over all $M_{\boldsymbol{\theta}}$ lines gives a halfspace $D^* = \{\sum_i a_i\lambda_i \leq K\}$. For any $\boldsymbol{\theta}$, the agent chooses identically from D^* in the original game as from D^{aug} in the augmented game. So D^* is an optimal delegation set.

Lemma 2 is the key technical result of the paper. In fact, it follows as an almost-direct corollary of Proposition 1 of Kovac and Mylovanov (2009), which can be extended to give a method for finding the optimal one-dimensional delegation set when the state is normally distributed on a line.

Proof. **Step 0:** There is a unique $\overline{k}$ satisfying $\mathbb{E}_{\theta\sim\mathcal{N}(0,1)}\left[\theta\middle|\theta\geq\overline{k}-|\boldsymbol{\lambda}|\right]=\overline{k}$ given any $|\boldsymbol{\lambda}|>0$. This follows from Appendix A.1 Claim 2, with $\overline{k}=\beta_0$.

Step 1: Suppose we had $N=1$, with agent bias $\lambda_1=\lambda>0$ and state $\theta_1=\theta$ distributed according to a standard normal. I claim that the optimal delegation set on $\mathbb{R}$ would be the interval $D=(-\infty,\overline{k}]$. Moreover, I this delegation set (which results in deterministic actions for any state θ) is actually optimal in the class of stochastic mechanisms, those which allow the agent to choose lotteries over actions.

This all follows from Kovac and Mylovanov (2009) Proposition 1, subject to checking the regularity condition which they call Assumption 1.[4] See Appendix A.1 for the application of this result to the current setting. In particular, Claim 4 completes this step of the proof.

Step 2: Show that if the state is known to lie on the line $M_{\boldsymbol{\theta}}$ and have the distribution as in Lemma 1, then $D^{\text{aug}}(P_{\boldsymbol{\theta}})$ is an optimal delegation set in $\mathbb{R}^N$.

The derivation of the one-dimensional delegation set from Step 1 immediately implies that D^{aug} is optimal over delegation sets which are contained in $M_{\boldsymbol{\theta}}$, since the agent's bias is parallel to the line with magnitude $\boldsymbol{\lambda}$. And there can be no improvements by allowing D to contain points off of $M_{\boldsymbol{\theta}}$. This is because choosing points $\boldsymbol{a}\notin M_{\boldsymbol{\theta}}$ in the augmented game is mathematically identical to choosing randomized actions in the one-dimensional problem of Step 1.

That is, suppose the agent were given some delegation set and chose $\boldsymbol{a}\notin M_{\boldsymbol{\theta}}$. Say that $\boldsymbol{a}\in\mathbb{R}^N$ projects onto $\boldsymbol{m}\in M_{\boldsymbol{\theta}}$ and is a distance of d away from the line. Both player's ideal points are on the line $M_{\boldsymbol{\theta}}$, so the distance-squared from the ideal point to the action is the distance-squared from the ideal point to $\boldsymbol{m}$, plus d^2. Therefore for each player, the payoff from taking action $\boldsymbol{a}$ is the payoff from taking action $\boldsymbol{m}$, minus d^2. This is exactly the payoff from taking a randomized action on the line $M_{\boldsymbol{\theta}}$ with mean $\boldsymbol{m}$ and variance d^2. And we know from Step 1 that there is no benefit from allowing the agent to take randomized actions. □

From the proof, we see that the delegation set D^{aug} implements the optimal outcomes

[4] Their paper also assumes that states are drawn from a compact support, but this is not necessary for the result. Normal distributions have thin tails, so the contribution to players' payoffs from states outside of $[-l,l]$ goes to 0 as $l\to\infty$.

not just of any deterministic mechanism, but of any stochastic one as well.

The cutoff value of $\overline{k}$ which determines the delegation ray depends on $|\boldsymbol{\lambda}|$, but – by the distributional assumptions on $\boldsymbol{\theta}$ – does not depend on the observation of $P_{\boldsymbol{\theta}}$. In particular, the conditional distributions of an iid normal are translation-invariant. Given $\boldsymbol{\theta} \in M_{\boldsymbol{\theta}}$, the conditional density of $\boldsymbol{\theta} = M_{\boldsymbol{\theta}}(k)$ is constant with respect to $\boldsymbol{\theta}$. This holds for any vector $\boldsymbol{\lambda}$ by the rotational symmetry of the iid normal.

Lemma 3. 1. Let $D^* \subset \mathbb{R}^N$ be the union of all possible D^{aug} rays:

$$D^* = \bigcup_{\boldsymbol{\theta} \in \mathbb{R}^N} D^{\text{aug}}\left(P_{\boldsymbol{\theta}}\right).$$

Then D^* is a half-space in $\mathbb{R}^N$ with

$$D^* = \left\{ \boldsymbol{x} \in \mathbb{R}^N \;\middle|\; \sum_{i=1}^{N} x_i \lambda_i \leq \overline{k} \cdot |\boldsymbol{\lambda}| \right\}$$

for $\overline{k}$ as defined in Lemma 2.

2. For any $\boldsymbol{\theta}$, the agent's optimal choice of $\boldsymbol{a}$ from D^* in the original game is equal to her optimal choice of $\boldsymbol{a}$ from $D^{\text{aug}}\left(P_{\boldsymbol{\theta}}\right)$ in the augmented game.

Proof. Part 1 is straightforward. For Part 2, it suffices to show that the closest point in D^* to the agent's ideal point $\boldsymbol{\theta} + \boldsymbol{\lambda}$ is in $D^{\text{aug}}(P_\theta)$, i.e., is collinear with $\boldsymbol{\theta}$ and $\boldsymbol{\theta} + \boldsymbol{\lambda}$. If $\boldsymbol{\theta} + \boldsymbol{\lambda} \in D^*$, this is clear. If $\boldsymbol{\theta} + \boldsymbol{\lambda} \notin D^*$, then the closest point to $\boldsymbol{\theta} + \boldsymbol{\lambda}$ is the projection of $\boldsymbol{\theta} + \boldsymbol{\lambda}$ onto the hyperplane bounding D^*. And so the result follows because this boundary is normal to $\boldsymbol{\lambda}$ (see Figure 3). □

This lemma proves Theorem 1. Giving the agent a delegation set D^* with K set to $\overline{k} \cdot |\boldsymbol{\lambda}|$ implements the same outcomes as giving the optimal delegation sets in the augmented game. And in expectation the principal must do weakly better in the augmented game, in which she has strictly more information than in the original game. So no delegation set in the original game can give a higher expected payoff than does D^*.

Proposition 1 then concerns the comparative statics on $\overline{k} = \frac{K}{|\boldsymbol{\lambda}|}$, the distance from the bounding hyperplane to the origin. It follows directly from Claim 2 part 3 of Appendix A.1, with $\overline{k} = \beta_0$ and $|\boldsymbol{\lambda}| = \lambda$.

5 Extensions

Throughout this section I relax the assumption that states are iid normally distributed. I introduce alternate assumptions on the joint distribution of the principal's priors as I proceed.

Halfspace delegation is not generally optimal for other iid distributions of states – eg, iid uniform – or when states are distributed according to independent normals with different variances. But the use of half-spaces to cap an agent's choices against the direction of her bias can still be thought of as a good rule of thumb for the principal even when it is not optimal. I will show in this section that these half-space delegation sets are nearly optimal when biases are extreme, or when there are many decisions. They also do well when decisions are taken sequentially instead of simultaneously, and can be generalized to work for cases where decisions are asymmetrically important.

Part of this argument is anticipated by Frankel (2010), which shows that hyperplane delegation (via "budget mechanisms") is max-min optimal against an agent with an unknown bias that is known to be identical across decisions. By the lemma below, these hyperplanes are in turn dominated by half-spaces when the direction of the agent's bias is known.[5]

For any agent bias $\boldsymbol{\lambda}$ and budget level K, let $D^H \equiv \{\boldsymbol{a} | \sum_i \lambda_i a_i = K\}$ be a hyperplane delegation set normal to $\boldsymbol{\lambda}$ and let $D^{HS} \equiv \{\boldsymbol{a} | \sum_i \lambda_i a_i \leq K\}$ be the corresponding half-space with outward normal vector of $\boldsymbol{\lambda}$.

Lemma 4. Fix an agent bias $\boldsymbol{\lambda}$ and a budget level K. Conditional on any $\boldsymbol{\theta}$:

1. The agent plays in delegation set D^H as if she had bias $\mathbf{0}$, i.e., as if she shared the principal's utility function.
2. The principal's payoff from the agent's optimal choice in D^{HS} is at least a high as his payoff from the agent's optimal choice in D^H.

All proofs of Extensions are in Appendix A.2.

In the language of Frankel (2010), part 1 of the lemma states that the delegation set D^H induces *aligned delegation*. If we can find some agent strategy in D^H which would give the principal a payoff of V_P, then the principal's payoff from the agent's optimal strategy in D^{HS} is at least as high as V_P.

In the analysis below, I will often focus on hyperplane and half-space delegation sets defined by one particular budget level. Setting $K = \sum_i \lambda_i \mathbb{E}[\theta_i]$, the hyperplane – or, in the

[5] The argument is also related to Chakraborty and Harbaugh (2010), which shows that $N-1$ dimensions of private information (e.g., a hyperplane) can be revealed by agents with state-independent preferences, corresponding to a particular form of very large bias.

case of a half-space, the bounding hyperplane – cuts through the mean state. I call this hyperplane and half-space D^{EH} and D^{EHS}, where "E" stands for "expectation."

$$D^{EH} \equiv \{\boldsymbol{a} | \sum_i \lambda_i a_i = \sum_i \lambda_i \mathbb{E}[\theta_i]\}$$

$$D^{EHS} \equiv \{\boldsymbol{a} | \sum_i \lambda_i a_i \leq \sum_i \lambda_i \mathbb{E}[\theta_i]\}$$

5.1 Large Biases

Fix a number of decisions, any joint distribution over the states $(\theta_1, ..., \theta_N)$, and fix a direction of the agent's bias. As we increase the magnitude of the bias in this direction, half-space delegation becomes approximately optimal:

Proposition 2. Consider a sequence of delegation problems in which states are drawn from some fixed joint distribution of $\boldsymbol{\theta}$ over $\mathbb{R}^N$ (with finite first and second moments). Fix an N-dimensional unit vector $\boldsymbol{\kappa}$; let the agent's biases in problem j be $\boldsymbol{\lambda}^j = c^j \boldsymbol{\kappa}$, for $c^j > 0$ going to infinity. The half-space delegation set D^{EHS} is approximately optimal as c gets large.

By approximately optimal I mean that for any $\epsilon > 0$, there exists $\bar{c}$ large enough so that if $c^j > \bar{c}$ then this delegation set gives the principal an expected payoff within ϵ of that of an optimal delegation set.

5.2 Large numbers of decisions

When there are many iid decisions, half-space delegation sets can give the principal approximately first-best payoffs. This holds for any iid distribution of states, not necessarily iid normal. In fact, the payoff loss relative to first-best goes at a rate of $\frac{1}{N}$, which is the *worst-case asymptotic optimal rate* as defined in Satterthwaite and Williams (2002).

To formalize this, I first need to define the *payoff loss relative to first-best* for a delegation set under an arbitrary joint distribution of states. If the principal received no information from the agent, he would choose the ex ante expected best point $\boldsymbol{a} = \mathbb{E}(\boldsymbol{\theta})$ and get a lifetime "no delegation" payoff of $-\sum_i \text{Var}(\theta_i)$. If he had full information, he would take $a_i = \theta_i$ in each period and receive a "first-best" payoff of 0. Let the payoff loss relative to first-best, or simply the relative payoff loss, be the proportion of the way from the first-best payoff to the no delegation payoff. First-best has relative payoff loss of 0, no delegation has relative payoff loss of 1, and delegation sets may have relative payoff losses of greater than 1. A delegation set with a lifetime principal payoff of V_P has relative payoff loss defined as

$$\frac{\text{First-best payoff} - V_P}{\text{First-best payoff} - \text{No delegation payoff}} = \frac{-V_P}{\sum_i \text{Var}(\theta_i)}$$

When states are iid with a standard normal distribution, the variance of each state is 1, and so the denominator of relative payoff loss is N. The relative payoff loss is $-V_P/N$, i.e., minus the payoff per decision.

It is easy to calculate the principal's payoff from D^{EH}, for any joint distribution of states. Any other hyperplane D^H would give the principal a lower payoff.

Lemma 5. The principal's payoff from choosing delegation set D^{EH} is $-\text{Var}\left(\sum_i \frac{\lambda_i}{|\boldsymbol{\lambda}|}\theta_i\right)$. In the case of independent states, this corresponds to $-\sum_i \frac{\lambda_i^2}{|\boldsymbol{\lambda}|^2}\text{Var}(\theta_i)$; for iid states, $-\text{Var}(\theta_i)$.

If we were to normalize all biases to be positive, then a positive correlation across states would reduce payoffs in D^{EH}. As we approached perfect correlation, the principal would no longer get a benefit from linking the decisions relative to treating them separately.

When states are independent, the lemma implies that the principal's relative payoff loss from D^{EH} is at worst $\min\left\{\frac{\max_i[\text{Var}(\theta_i)]}{\sum_i \text{Var}(\theta_i)}, \frac{\max_i\left[\lambda_i^2\right]}{|\boldsymbol{\lambda}|^2}\right\}$. So the relative payoff loss would go to 0 as the number of decisions increased if, in the limit, either no state had a disproportionate share of the variance or no decision had a disproportionate share of the bias. The half-space D^{EHS} weakly improves on D^{EH}, by Lemma 4.

When states are iid, the relative payoff loss from D^{EH} is exactly $\frac{1}{N}$, regardless of agent bias $\boldsymbol{\lambda}$. So in an iid problem with many decisions, D^{EHS} approaches first-best payoffs at a rate of at worst $\frac{1}{N}$.

Proposition 3. Consider a sequence of delegation problems with the same iid distribution of states in each problem, with the number of decisions N going to infinity. Then for any corresponding sequence of agent biases, the principal's relative payoff loss from D^{EHS} is at most $\frac{1}{N}$.

When states are iid normal, we can calculate that the relative payoff loss from the optimal (half-space) delegation set D^* goes to 0 at the same rate of $\frac{1}{N}$:

Lemma 6. Fix a sequence of scalar biases $\lambda_1, \lambda_2, ...,$ with $\lambda_i \neq 0$ for some i. Consider a sequence of delegation problems indexed by the number of decisions $N \in \mathbb{N}$, such that in problem N the agent's bias is $(\lambda_1, ..., \lambda_N)$. Let states be iid normal. Then the principal's relative payoff loss from the optimal delegation set D^* goes to 0 at a rate of $\frac{1}{N}$.

I say that the payoff loss goes to 0 at a rate of $f(N)$, with $f(N) \to 0$ as $N \to \infty$, if there exist positive constants $0 < \underline{\eta} < \overline{\eta}$ such that for N large enough, the relative payoff loss when there are N decisions is contained in $(\underline{\eta} f(N), \overline{\eta} f(N))$.

So the payoff loss from the half-space D^{EHS} declines at a rate of at worst $\frac{1}{N}$, for any iid states, and in one particular example the optimal delegation set has payoff loss declining at a rate of exactly $\frac{1}{N}$. In the language of Satterthwaite and Williams (2002), D^{EHS} is worst-case asymptotic optimal.[6,7]

More precisely, loosely following the notation of Satterthwaite and Williams (2002), fix a sequence of agent biases $(\lambda_1, \lambda_2, ...)$ and a set E of finite variance distributions over $\mathbb{R}$. Let a "mechanism" Φ be a function mapping $N \in \mathbb{N}$ and $F \in E$ into a closed subset of $\mathbb{R}^N$. A pair (N, F) should be interpreted as a delegation problem over N decisions in which each state is drawn iid from the distribution F, and the agent's bias is $\boldsymbol{\lambda} = (\lambda_1, ..., \lambda_N)$. Then $\Phi(N, F)$ defines a delegation set over $\mathbb{R}^N$ from which the agent must draw actions. Let Φ^{EHS} be the mechanism which, for any (N, F), maps to the half-space delegation set D^{EHS}.

Let the payoff loss of the delegation set $\Phi(N, F)$ in the relevant delegation problem be denoted $e(\Phi, N, F)$. Let its worst-case payoff loss be defined as $e^{wor}(\Phi, N, E) \equiv \sup_{F \in E} e(\Phi, N, F)$, where the worst-case is taken over distributions. A mechanism Φ is said to be *worst-case asymptotic optimal* if for any other mechanism Φ^* there exists a positive constant η such that

$$e^{wor}(\Phi, N, E) \leq \eta e^{wor}(\Phi^*, N, E)$$

for all $N \in \mathbb{N}$. So a mechanism is worst-case asymptotic optimal if the payoff loss in the worst possible distribution declines at the fastest possible rate as the number of decisions goes to infinity.

Proposition 4. Take any sequence of biases and any set of distributions E which contains a normal distribution. Then the mechanism Φ^{EHS} (inducing the half-space delegation sets D^{EHS}) is worst-case asymptotic optimal.

[6] The hyperplane delegation sets D^{EH} also have payoff losses declining at a rate of $\frac{1}{N}$, but there is no reason for a principal to consider these sets; they are dominated by D^{EHS}.

[7] This result is weaker than the result of worst-case asymptotic optimality in Satterthwaite and Williams (2002), which – in a very different setting – finds a mechanism which is worst-case asymptotic optimal and does not require the principal to input any information about the economic environment. The mechanism I consider, using the delegation sets D^{EHS}, requires the principal to input the agent's bias and also the mean of the distribution.

5.3 Sequential Decisions

Suppose decisions are taken sequentially rather than all at once; the agent is given a delegation set $D \subseteq \mathbb{R}^N$ from which to choose actions, but she must choose action a_i before observing state θ_{i+1}. There is less information available at the time decisions are made. But the half-space delegation sets discussed above continue to align incentives well. Analogs of the results of the above extensions all go through.

Proposition 2 is unchanged in the sequential problem. When the magnitude of the agent's bias increases to infinity, the half-space delegation set D^{EHS} is approximately optimal. The proof just requires slight modifications to account for the difference in the agent's information at the time when she takes actions; I omit this proof from the paper.

It also remains true that, as in Proposition 3, the principal's relative payoff loss from D^{EHS} goes to 0 if we have iid states and increase the number of decisions. But the proof technique of Lemma 5 now gives us a rate of convergence of something like $\frac{\log N}{N}$ rather than $\frac{1}{N}$. This depends on certain regularity conditions on the biases:

Lemma 7. Consider a sequential problem with independent states. The principal's payoff from the hyperplane delegation set D^{EH} is $-\sum_{i=1}^{N} \frac{\lambda_i^2 \text{Var}(\theta_i)}{\sum_{j=i}^{N} \lambda_j^2}$. The principal's payoff from D^{EHS} is weakly higher.

For iid states, the variance is constant. If all biases were equal to the same value λ in an iid problem, the payoff would evaluate to minus the variance times $\lambda^2 \sum_{j=1}^{N} \frac{1}{j}$; this sum increases in N at a rate of $\log N$. This shows that the relative payoff loss from D^{EH} would go to 0 at a rate of $\frac{\log N}{N}$. For other λ_i sequences satisfying reasonable regularity conditions – that $\sum_{i=1}^{N} \lambda_i^2$ increases approximately linearly in N – then the relative payoff loss from D^{EH} goes to 0 at the same rate. The relative payoff loss from D^{EHS} goes to 0 at a weakly faster rate.

Because I do not have results about exactly optimal delegation sets in a sequential environment, I cannot show that the convergence rate of $\frac{\log N}{N}$ is worst-case asymptotically optimal.

5.4 Asymmetrically important decisions

Suppose that certain decisions are agreed on by the principal and the agent to be more important than other ones. For instance, Amador and Bagwell (2011) study the choice of a tariff for a single good in a trade agreement, and argue for a cap on the maximum allowed tariff. But if tariffs are chosen for multiple goods at a time, then we can improve outcomes by linking tariff caps on the different goods through a joint constraint. If markets are all equally important, the current paper would argue for linking them through a simple budget constraint, inducing a half-space delegation set with boundary normal to the agent's biases. How should we modify this half-space when certain markets are larger than others?[8]

Say that the relative significance of decision i is $\gamma_i > 0$. The principal and agent agree about these significances, which are common knowledge at the beginning of time. Consider new objective functions

$$V_P(\boldsymbol{a}|\boldsymbol{\theta}) = -\sum_i \gamma_i \cdot (a_i - \theta_i)^2$$
$$V_A(\boldsymbol{a}|\boldsymbol{\theta}) = -\sum_i \gamma_i \cdot (a_i - \theta_i - \lambda_i)^2$$

Changing utilities to have significance factors gives a mathematically equivalent problem as utilities without significance factors, with an appropriate rescaling of actions, states, and biases. Let $\hat{a}_i = \sqrt{\gamma_i} a_i$, $\hat{\theta}_i = \sqrt{\gamma_i}\theta_i$, and $\hat{\lambda}_i = \sqrt{\gamma_i}\lambda_i$. (Notice that even if states θ_i were iid, the rescaled states $\hat{\theta}_i$ are no longer iid). Then

$$V_P(\boldsymbol{a}|\boldsymbol{\theta}) = -\sum_i (\hat{a}_i - \hat{\theta}_i)^2$$
$$V_A(\boldsymbol{a}|\boldsymbol{\theta}) = -\sum_i (\hat{a}_i - \hat{\theta}_i - \hat{\lambda}_i)^2$$

Under this rescaling, we can consider half-space delegation set $\hat{D}^{HS} = \{\hat{\boldsymbol{a}}| \sum_i \hat{\lambda}_i \hat{a}_i \leq K\}$. The previous extensions suggest that in the absence of solving for fully optimal delegation sets, as a rule of thumb we can be comfortable using delegation sets of this form in the rescaled problem. For instance, with large biases or with many independent states, $\hat{D}^{HS}$ will be approximately optimal when we set $K = \sum_i \mathbb{E}(\hat{\theta}_i)$.

$\hat{D}^{HS}$ corresponds to a delegation set in the original variables of $D^{HS} = \{\boldsymbol{a}| \sum_i \gamma_i \lambda_i a_i \leq K\}$, where K is unchanged. This is a half-space with a different gradient than considered previously – it is normal not to $(\lambda_i)_{i=1}^N$ but to $(\gamma_i\lambda_i)_{i=1}^N$. The "price" of one unit of action a_i is proportional to the significance of decision i as well as the bias on that decision.

[8]Amador and Bagwell actually argue for a cap on the surplus of domestic import-competing firms, measured in dollars, and then point out that this can be implemented by a tariff cap. Surpluses already include a scaling for market size. But given that trade agreements are more naturally written over tariff rates than over implied surpluses, it is interesting to consider how tariff rates themselves should be linked across markets.

6 Conclusion

This paper suggests a general heuristic for a principal who is delegating multiple decisions to a biased agent. Rather than treat the decisions independently, the principal should give the agent a single budget out of which she "purchases" all of her actions. To induce the agent to make the proper trade-offs across decisions, the "price" per unit of action should be proportional to the agent's bias on that decision.

This policy is fully optimal when principal ideal points are iid normal, the agent has a constant bias, the players have quadratic losses about their ideal points, and all decisions are taken at once. And I show that this policy continues to give the principal high payoffs, if not fully optimal, under some relaxed assumptions.

Developing a method to solve for optimal delegation sets under general joint distributions of states, or with different functional forms on payoffs, is a topic for future research. A first step might be to search for regularity conditions on payoffs and distributions, analogous to those explored in the literature on the delegation of a single decision, to guarantee a convex delegation set.

A Mathematical Appendix

A.1 Derivation of D^*

In this appendix I follow the argument and notation from Kovac and Mylovanov (2009) to derive the optimal delegation set in a one-decision problem with state θ distributed according to a standard normal, and agent bias $\lambda > 0$. In this construction I also derive properties of $\overline{k}$, which will be set to be equal to the term β_0 defined below.

Let $g(\theta) \equiv 1 - \Phi(\theta - \lambda) - \lambda\phi(\theta - \lambda)$ for Φ and ϕ the cdf and pdf of a standard normal.

Claim 1.

1. $\lim_{\theta \to -\infty} g(\theta) = 1$
2. $\lim_{\theta \to \infty} g(\theta) = 0$
3. There exists a point $\beta_1 \in \mathbb{R}$ such that $g(\theta) > 0$ for $\theta < \beta_1$; $g(\beta_1) = 0$; and $g(\theta) < 0$ for $\theta > \beta_1$.
4. $g(\theta)$ is decreasing on $\theta < \beta_1$.

Proof. Parts 1 and 2 are straightforward.

To prove parts 3 and 4, recalling the identity that $\phi'(\theta) = -\theta\phi(\theta)$,

$$g'(\theta) = -\phi(\theta - \lambda) - \lambda\phi'(\theta - \lambda) = -\phi(\theta - \lambda)(1 - \lambda(\theta - \lambda))$$

So g is decreasing if $1 - \lambda(\theta - \lambda) > 0$, i.e., $\theta < \frac{1}{\lambda} + \lambda$. And g is increasing if $\theta > \frac{1}{\lambda} + \lambda$. Combining this observation with parts 1 and 2, it must be the case that $g(\theta)$ has a unique intersection with 0 at some $\beta_1 < \frac{1}{\lambda} + \lambda$, that g is positive and decreasing for $\theta < \beta_1$, and g is negative for $\theta > \beta_1$. □

Claim 2. 1. There exists a unique value $\beta_0 \in \mathbb{R}$ such that $\int_{\beta_0}^{\infty} g(\theta)d\theta = 0$, with $\beta_0 < \beta_1$.

2. β_0 is the unique solution to $\mathbb{E}_{\theta\sim\mathcal{N}(0,1)}[\theta|\theta \geq \beta_0 - \lambda] = \beta_0$.
3. β_0 is decreasing in λ, with β_0 going to 0 as $\lambda \to \infty$ and β_0 going to ∞ as $\lambda \to 0^+$.

Proof. 1. The statement $\int_{\beta_0}^{\infty} g(\theta)d\theta = 0$ is equivalent to

$$\int_{\beta_0}^{\beta_1} g(\theta)d\theta = -\int_{\beta_1}^{\infty} g(\theta)d\theta.$$

The function $\int_{X}^{\beta_1} g(\theta)d\theta$ becomes arbitrarily large as X goes to minus infinity (by part 1 of Claim 1); it decreases in X for $X < \beta_1$; and it goes to 0 as X goes to β_1. So there is a unique solution β_0 on the left-hand side, with $\beta_0 < \beta_1$, so long as the right-hand side is finite. Integrating by parts:

$$\begin{aligned} -\int_{\beta_1}^{\infty} g(\theta)d\theta &= -\int_{\beta_1}^{\infty} (1 - \Phi(\theta - \lambda) - \lambda\phi(\theta - \lambda))\, d\theta \\ &= \lambda(1 - \Phi(\beta_1 - \lambda)) + (1 - \Phi(\beta_1 - \lambda))\beta_1 - \int_{\beta_1}^{\infty} \theta\phi(\theta - \lambda)d\theta < \infty \end{aligned}$$

2. To establish the second result, we will again integrate by parts:

$$
\begin{aligned}
& \int_{\beta_0}^{\infty} g(\theta) d\theta = 0 \\
\iff & \int_{\beta_0-\lambda}^{\infty} (1 - \Phi(\theta) - \lambda\phi(\theta)) d\theta = 0 \\
\iff & -\lambda(1 - \Phi(\beta_0 - \lambda)) - (1 - \Phi(\beta_0 - \lambda))\beta_0 - \lambda) + \int_{\beta_0-\lambda}^{\infty} \theta\phi(\theta) d\theta = 0 \\
\iff & \int_{\beta_0-\lambda}^{\infty} \theta\phi(\theta) d\theta = (1 - \Phi(\beta_0 - \lambda))\beta_0 \\
\iff & \frac{\int_{\beta_0-\lambda}^{\infty} \theta\phi(\theta) d\theta}{1 - \Phi(\beta_0 - \lambda)} = \beta_0
\end{aligned}
$$

3. Taking comparative statics of β_0 as a function of λ,

$$
\begin{aligned}
& \int_{\beta_0-\lambda}^{\infty} (1 - \Phi(\theta) - \lambda\phi(\theta)) d\theta = 0 \\
\implies & \left(1 - \frac{d\beta_0}{d\lambda}\right) (1 - \Phi(\beta_0 - \lambda) - \lambda\phi(\beta_0 - \lambda) - \int_{\beta_0-\lambda}^{\infty} \phi(\theta) d\theta = 0 \\
\implies & \frac{d\beta_0}{d\lambda} = 1 - \frac{1 - \Phi(\beta_0 - \lambda)}{1 - \Phi(\beta_0 - \lambda) - \lambda\phi(\beta_0 - \lambda)}
\end{aligned}
$$

where we get the second line by taking the derivative of both sides with respect to λ. Now notice that $\beta_0 < \beta_1$, and so $g(\beta_0) = 1 - \Phi(\beta_0 - \lambda) - \lambda\phi(\beta_0 - \lambda) > 0$. Therefore the fraction on the right-hand side of the final line is greater than 1, hence the right-hand side is negative. This establishes that $\frac{d\beta_0}{d\lambda} < 0$.

Then by the standard properties of the normal distribution, $\mathbb{E}[\theta|\theta \geq X]$ is approximately 0 for X going to minus infinity, corresponding to β_0 going to 0 as $\lambda = \beta_0 - X$ goes to ∞; and $\mathbb{E}[\theta|\theta \geq X]$ is approximately X for X large and positive, corresponding to β_0 going to ∞ as λ goes to 0. □

Claim 3. The function $g(\theta)$ is strictly decreasing at θ when $g(\theta) \in [0, 1]$.

Proof. Immediate from parts 3 and 4 of Claim 1.[9] □

[9] This can also be shown by observing that normal distributions have increasing hazard rates, and applying the observation from footnote 10 of Kovac and Mylovanov (2009).

Claim 4 (Lemma 2, Step 1)**.** For $\theta \sim \mathcal{N}(0,1)$, the optimal delegation set is $(-\infty, \overline{k}]$ for $\overline{k} = \beta_0$ as defined in Claim 2. Moreover, asking the agent to choose an action from this set is optimal even if stochastic mechanisms (in which the agent chooses lotteries over actions) may be allowed.

Proof. By Claim 3, Assumption 1 of Kovac and Mylovanov (2009) holds. This result is therefore an application of Proposition 1 of Kovac and Mylovanov (2009); see that paper as well for details on stochastic mechanisms. □

A.2 Additional Proofs

Proof of Lemma 4.

1. The agent's problem is:

$$\max_{\boldsymbol{a}} - \sum_i (a_i - \theta_i - \lambda_i)^2 \text{ s.t. } \sum_i \lambda_i a_i = K$$

$$\max_{\boldsymbol{a}} - \sum_i (a_i - \theta_i)^2 + \sum_i 2(a_i - \theta_i)\lambda_i - \sum_i \lambda_i^2 \text{ s.t. } \sum_i \lambda_i a_i = K$$

$$\max_{\boldsymbol{a}} - \sum_i (a_i - \theta_i)^2 + 2K - 2\sum_i \theta_i \lambda_i - \sum_i \lambda_i^2 \text{ s.t. } \sum_i \lambda_i a_i = K$$

   And this has identical argmax as the problem

$$\max_{\boldsymbol{a}} - \sum_i (a_i - \theta_i)^2 \text{ s.t. } \sum_i \lambda_i a_i = K$$

2. Let $\boldsymbol{a}^H$ be an agent-optimal choice from D^H, and $\boldsymbol{a}^{HS}$ be an agent-optimal choice from D^{HS}. I seek to show that the principal weakly prefers $\boldsymbol{a}^{HS}$ to $\boldsymbol{a}^H$.

   The agent must weakly prefer $\boldsymbol{a}^{HS}$ to $\boldsymbol{a}^H$ because it is a solution to a more relaxed optimization. Write the agent's objective function as

$$\left[- \sum_i (a_i - \theta_i)^2\right] + \left[2 \sum_i a_i \lambda_i\right] + \left[-2 \sum_i \theta_i \lambda_i - \sum_i \lambda_i^2\right]$$

   From right to left, the terms in the third bracket are independent of $\boldsymbol{a}$. The second bracket is weakly lower for $\boldsymbol{a}^{HS}$ than for $\boldsymbol{a}^H$ (strictly so, if $\boldsymbol{a}^{HS} \notin D^H$). Therefore, the first bracket is weakly higher for $\boldsymbol{a}^{HS}$. That is, the principal weakly prefers $\boldsymbol{a}^{HS}$ to $\boldsymbol{a}^H$.

□

Proof of Proposition 2. By Lemma 4, it suffices to show that there is some strategy under the hyperplane delegation set $D^H = \{\boldsymbol{a} | \sum_i \kappa_i a_i = \sum_i \kappa_i \mathbb{E}[\theta_i]\}$ which achieves approximately optimal payoffs. I show this in two steps.

Step 1: Show that for any delegation set D, there exists some K^D such that the hyperplane delegation set $\{\boldsymbol{a} | \sum_i \kappa_i a_i = K^D\}$ approximately matches the payoffs from D as c gets large.

The agent's payoff is

$$\left[-\sum_i (a_i - \theta_i)^2\right] + 2c\left[\sum_i a_i \kappa_i\right] + \left[-2c\sum_i \theta_i \kappa_i - c^2 \sum_i \kappa_i^2\right]$$

Fix some $\delta > 0$ and some $K \in \mathbb{R}$. Suppose it is that case that there exists a point $\boldsymbol{a} \in D$ such that $\sum_i \kappa_i a_i \geq K$. Then fix $\boldsymbol{a}^I \in D$ (an "inside" action) such that $\sum_i a_i^I \kappa_i \geq K$, and take any $\boldsymbol{a}^O(\boldsymbol{\theta})$ (an "outside" action, not necessarily in D) such that $\sum_i a_i^O \kappa_i \leq K - \delta$. The payoff to the agent from $\boldsymbol{a}^I$ minus the payoff from $\boldsymbol{a}^O(\boldsymbol{\theta})$ is at worst $-L^2 + 2\delta c$.[10] This is positive for $2c\delta > L^2$. So for $L = \sqrt{2c\delta}$, it is the case that for every state within a radius L of $\boldsymbol{a}^I$, the agent chooses an action such that $\sum_i \kappa_i a_i > K - \delta$.

Now we can consider two cases. Case 1 is that $\max_{\boldsymbol{a} \in D} \sum_i \kappa_i a_i$ does not exist – no action vector achieves the maximum. Then taking $c \to \infty$, it holds that the principal's payoffs from D go to minus infinity.[11] Therefore the hyperplane delegation set $\{\boldsymbol{a} | \sum_i \kappa_i a_i = K^D\}$ for any $K^D \in \mathbb{R}$ improves on D for c large enough (where this hyperplane gives a finite payoff which is independent of c – see Lemma 4 part 1).

Case 2 is that the maximum is achieved, in which we can take $K^D = \max_{\boldsymbol{a} \in D} \sum_i \kappa_i a_i$ and fix $\boldsymbol{a}^I$ to be some action such that $\sum_i \kappa_i a_i = K^D$. Then take δ to 0 and c to ∞ in such a way that $L = \sqrt{c\delta}$ goes to infinity, while $L\delta$ goes to 0. So for δ arbitrarily small, for c large enough, an arbitrarily high proportion of states have actions such that $\sum_i \kappa_i a_i \geq K^D - \delta$. For each such action, there is an action with $\sum_i \kappa_i a_i = K^D$ within a distance δ. The action with this sum gives the principal a payoff of at worst $4L\delta + \delta^2$ less than the payoff from the action in D.[12] This goes to 0 as $L\delta$ and δ go to 0. So the principal's worst case payoff loss

[10] The first bracket in the payoff term is at worst $-L^2$ for $\boldsymbol{a}^I$, and at best 0 for $\boldsymbol{a}^O$; the second bracket is $2c$ times the difference in $\sum_i a_i \kappa_i$, which is at worst δ; and the third bracket cancels out.

[11] If almost all states imply actions outside of some bounded region $B \subseteq \mathbb{R}^N$, for B growing to take up the whole space, then the principal's payoff must go to minus infinity because almost all states imply actions very far from the state. And indeed, taking $K \to \sup_{\boldsymbol{a} \in D} \sum_i \kappa_i a_i$, $\delta \to 0$, then $c \to \infty$, we see that this is the case. All actions for which $\sum_i \kappa_i a_i \geq K - \delta$ are eventually outside of a bounded region which grows to take up the space as c grows, and some such action is taken for a set of states with measure approaching one.

[12] The worst case is that the action in D^H is δ greater distance than the action in D from $\boldsymbol{\theta}$, in which case

on an arbitrarily high proportion of states goes to 0, and the payoff contribution from other states goes to 0 (by the assumption of finite variance), hence the payoff from the hyperplane delegation set approximately matches that from D.

Step 2: Out of all hyperplane delegation sets $\{\boldsymbol{a} | \sum_i \kappa_i a_i = K\}$, a value $K = \sum_i \kappa_i \mathbb{E}[\theta_i]$ is optimal. (The agent's play is independent of c – see Lemma 4 part 1).

Given any states, the agent chooses an action equal to the projection of the state onto the relevant hyperplane. This gives the principal an expected payoff equal to minus the distance squared from the state onto the hyperplane. This expected payoff is maximized by choosing a hyperplane which intersects the expected state, i.e., setting $K = \sum_i \kappa_i \mathbb{E}[\theta_i]$. □

Proof of Lemma 5. Given $\boldsymbol{\theta}$, the agent chooses an action $\boldsymbol{a}$ in the hyperplane D closest to $\boldsymbol{\theta}$ (as in Lemma 4 part 1). Letting k be the distance from $\boldsymbol{a}$ to $\boldsymbol{\theta}$, the principal's payoff is k^2. Using dot-product notation, we can solve for k as

$$k = \frac{\boldsymbol{\lambda}}{|\boldsymbol{\lambda}|} \cdot \boldsymbol{\theta} - \frac{\boldsymbol{\lambda}}{|\boldsymbol{\lambda}|} \cdot \mathbb{E}[\boldsymbol{\theta}]$$

And the expectation of this squared is exactly the variance of the relevant sum. □

Proof of Proposition 3. Follows from Lemma 4 part 2 and Lemma 5. □

Proof of Lemma 6. The lifetime payoff from D^* in the problem with N decisions is exactly the payoff from a delegation problem over one decision, in which the state θ has a standard normal distribution, the agent has bias $\lambda = |\boldsymbol{\lambda}|$, and the delegation set is chosen optimally as the optimal half-bounded interval. The payoff from this one decision problem is 0 if $\lambda = 0$, and negative for any other λ. It decreases in λ, going to the no delegation payoff of -1 as λ goes to infinity. So the lifetime payoff from D^* in the N decision problem approaches a negative constant weakly greater than -1. Therefore the relative payoff loss asymptotes to 0 at a rate of $\frac{1}{N}$. □

Proof of Proposition 4. Given these definitions, the result follows immediately. For any distribution F, Proposition 3 states that delegation sets D^{EHS} converge to a payoff loss of 0 at a rate at least as fast as $\frac{1}{N}$. And for the normal distribution, Lemma 6 implies that the optimal delegation sets converge to a payoff loss of 0 at an asymptotic rate of $\frac{1}{N}$.[13] No del-

the payoff difference is $(|\boldsymbol{a} - \boldsymbol{\theta}| + \delta)^2 - |\boldsymbol{a} - \boldsymbol{\theta}|^2 = 2|a - \theta|\delta + \delta^2 \leq 4L\delta + \delta^2$.

[13]The proofs above are for the standard normal distribution, but other normal distributions have the same payoff properties.

egation sets converge faster than the optimal ones. So no delegation sets have a worst-case asymptotic convergence rate of better than $\frac{1}{N}$ over the set of distributions E. $\square$

Proof of Lemma 7. Lemma 4 goes through unchanged in the sequential problem. The agent plays in a hyperplane $D^H = \{\boldsymbol{a} | \sum_i \lambda_i a_i = K\}$ as if she were unbiased (the delegation set is aligned delegation), and the principal's payoff from a corresponding half-space $D^{HS} = \{\boldsymbol{a} | \sum_i \lambda_i a_i \leq K\}$ is weakly higher than that from D^H under any joint distribution of states.

For some hyperplane delegation set of the form of D^H, let K_i be the budget remaining at period i: the agent must choose remaining actions so that $\sum_{j=i}^N \lambda_j a_j = K_i$. I claim that at period i with remaining budget K_i:

(i) Given θ_i, the agent chooses

$$a_i = \theta_i + \frac{\lambda_i K_i - \theta_i \lambda_i^2 - \lambda_i \sum_{j=i+1}^N \lambda_j \mathbb{E}[\theta_j]}{\sum_{j=i}^N \lambda_j^2}$$

(ii) Prior to the realization of θ_i, the principal's expected payoff from actions i through N is

$$-\sum_{j=i}^N \frac{\lambda_j^2 \text{Var}[\theta_j]}{\sum_{s=j}^N \lambda_s^2} - \frac{\left(K_i - \sum_{j=i}^N \lambda_j \mathbb{E}[\theta_j]\right)^2}{\sum_{j=i}^N \lambda_j^2}$$

By the sequential version of Lemma 4, the agent acts to maximize the principal's payoff. Given that, I prove (i) and (ii) by backwards induction:

Base case – show (i) and (ii) for $i = N$:

At period N, the agent must choose $a_i = \frac{K_i}{\lambda_i}$ and this gives the principal a payoff of $-\mathbb{E}\left[(a_i - \theta_i)^2\right] = -\text{Var}[\theta_i] - \left(\frac{K_i}{\lambda_i} - \mathbb{E}[\theta_i]\right)^2$. Plugging in, we can see that this confirms the inductive hypotheses (i) and (ii).

Inductive case – show (i) and (ii) for i, supposing they hold for periods $i+1$ and beyond:

By inductive hypothesis (ii), the principal's payoff (which the agent maximizes) for periods i through N, given K_i and θ_i, is

$$-(a_i - \theta_i)^2 - \sum_{j=i+1}^N \frac{\lambda_j^2 \text{Var}[\theta_j]}{\sum_{s=j}^N \lambda_s^2} - \frac{\left(K_i - \lambda_i a_i - \sum_{j=i+1}^N \lambda_j \mathbb{E}[\theta_j]\right)^2}{\sum_{j=i+1}^N \lambda_j^2}$$

Taking the first order condition of this payoff with respect to a_i confirms (i) for period i. Plugging this optimal action back into the payoff and taking expectation over θ_i

confirms (ii).

Plugging in $K_1 = \sum_{j=1}^{N} \lambda_j \mathbb{E}[\theta_j]$ into (ii) at $i = 1$ establishes the payoff result for D^{EH}. By the sequential analog to Lemma 4 part 2, D^{EHS} gives a weakly higher payoff. □

References

Ricardo Alonso and Niko Matouschek. Optimal delegation. *Review of Economic Studies*, 75: 259–293, 2008.

Manuel Amador and Kyle Bagwell. The theory of optimal delegation with an application to tariff caps. Working paper, 2011.

Manuel Amador, Ivan Werning, and George-Marios Angeletos. Commitment vs. flexibility. *Econometrica*, 74(2):365–396, 2006.

Attila Ambrus and Georgy Egorov. Delegation and nonmonetary incentives. Working paper, 2009.

Susan Athey, Andrew Atkeson, and Patrick J. Kehoe. The optimal degree of discretion in monetary policy. *Econometrica*, 73(5):1431–1475, 2005.

Archishman Chakraborty and Rick Harbaugh. Comparative cheap talk. *The Journal of Economic Theory*, 132:70–94, 2007.

Archishman Chakraborty and Rick Harbaugh. Persuasion by cheap talk. *American Economic Review*, 100(5):2361–2382, 2010.

Peter M. DeMarzo and Michael J. Fishman. Optimal long-term financial contracting. *Review of Financial Studies*, 20(6):2079–2128, 2007.

Peter M. DeMarzo and Yuliy Sannikov. Optimal security design and dynamic capital structure in a continuous-time agency model. *Journal of Finance*, 61(6):2681–2724, 2006.

Alexander Frankel. Aligned delegation. 2010.

Maria Goltsman, Johannes Hörner, Gregory Pavlov, and Francesco Squintani. Arbitration, mediation, and cheap talk. 2007.

Bengt Holmström. *On incentives and control in organizations*. PhD thesis, Stanford University, 1977.

Bengt Holmström. *On the theory of delegation*. In: Bayesian models in economic theory. Elsevier Science Ltd, 1984.

Matthew O. Jackson and Hugo Sonnenschein. Overcoming incentive constraints. *Econometrica*, 75(1):241–257, January 2007.

Eugen Kovac and Tymofiy Mylovanov. Stochastic mechanisms in settings without monetary transfers: the regular case. *Journal of Economic Theory*, 114:1373–1395, 2009.

Andrey Malenko. Optimal design of internal capital markets. Working paper, 2011.

David Martimort and Aggey Semenov. Continuity in mechanism design without transfers. *Economics Letters*, 93:182–189, 2006.

Nahum D. Melumad and Toshiyuki Shibano. Communication in settings with no transfers. *RAND Journal of Economics*, 22(2):173–198, 1991.

Mark A. Satterthwaite and Steven R. Williams. The optimality of a simple market mechanism. *Econometrica*, 70(5), 2002.