

Aligned Delegation – Online Appendix

Alexander Frankel*

A Omitted Proofs

Proof of Lemma 1. I establish this first for distributional quotas (with $\Phi = \Phi^\infty$), and then generalize to arbitrary moment mechanisms. Distributional quotas fix an aggregate distribution μ – a measure of mass N on the set of actions – and let the agent choose any action distributions m_1, \dots, m_N so that $\sum_i m_i = \mu$. Other moment mechanisms may give additional freedom on $\sum_i m_i$.

Case 1: Φ^∞ -moment mechanisms.

1. For a given state θ_i , stage payoffs $\mathbb{E}_{a_i \sim m_i} U_P(a_i | \theta_i)$ are continuous in the distribution m_i in the sense that if a sequence of distributions $m_i^{(j)}$ weakly converges to m_i , then stage payoffs from $m_i^{(j)}$ converge to m_i . This follows because $U_P(a_i | \theta_i)$ is bounded and continuous in a_i (Billingsley (1995), Theorem 25.8).
2. Write the aligned payoff to the principal from requiring aggregate distribution μ in a Φ^∞ -moment mechanism, conditional on realized states θ , as $V(\mu | \theta)$. The payoff is well-defined because payoffs are continuous in assignments with respect to weak convergence, and the set of assignments consistent with μ is compact (contains its limits with respect to weak convergence).

*University of Chicago, Booth School of Business. 5807 S Woodlawn Ave, Chicago IL. afrankel@uchicago.edu.

Moreover, the aligned payoff $V(\mu|\theta)$ is continuous in μ with respect to weak convergence, because an agent with a nearby aggregate distribution can approximately replicate the optimizing action distribution (m_1, \dots, m_N) .

Therefore the expected payoff over state realizations $\mathbb{E}_\theta V(\mu|\theta)$ is also continuous in μ .

3. Take a sequence of aggregate distributions which approach the (bounded) value $\sup_\mu \mathbb{E}_\theta V(\mu|\theta)$. By Helly's theorem (Billingsley (1995), Theorem 25.9), there exists a subsequence weakly converging to a limiting measure. By continuity of aligned payoffs with respect to the measure, this limiting measure achieves the maximized payoff.

Case 2: Φ -moment mechanisms with finitely many restrictions J .

Let the set of aggregate distributions consistent with the moment restrictions \mathcal{K} be denoted by $\mu(\mathcal{K})$:

$$\mu(\mathcal{K}) \equiv \left\{ \mu \mid \int_{\mathcal{A}} \varphi^{(j)}(a) d\mu(a) = K^{(j)} \text{ for each } j \leq J \right\}$$

The set $\mu(\mathcal{K})$ will be nonempty for each feasible \mathcal{K} , and compact. So for each θ and \mathcal{K} the aligned payoff (maximizing $V(\mu|\theta)$ over $\mu \in \mu(\mathcal{K})$) is well-defined.

Moreover, $\mu(\mathcal{K})$ will be upper-hemicontinuous with respect to \mathcal{K} because the moment functions are bounded and continuous. And the graph of $\mu(\mathcal{K})$ will be convex by linearity of the integration operator: if $\mu' \in \mu(\mathcal{K}')$ and $\mu'' \in \mu(\mathcal{K}'')$, then for each $\alpha \in (0, 1)$ it holds that $\alpha\mu' + (1 - \alpha)\mu'' \in \mu(\alpha\mathcal{K}' + (1 - \alpha)\mathcal{K}'')$.

Convexity of the graph plus closure (due to upper-hemicontinuity) imply lower-hemicontinuity with respect to \mathcal{K} as well. Therefore by Berge's theorem of the maximum, the optimized payoff given \mathcal{K} and θ is continuous in \mathcal{K} . This payoff is likewise continuous in \mathcal{K} taking expectation over θ . The set of

feasible \mathcal{K} is compact, and so an expected payoff-maximizing \mathcal{K} exists.

Case 3: Φ -moment mechanisms with infinitely many moment restrictions.

We can approximate these mechanisms by fixing the first $k \in \mathbb{N}$ moments, and taking k to infinity. For each k -approximation, we can find an expected-payoff maximizing sequence of moments $\mathcal{K}_k = (K_k^{(1)}, \dots, K_k^{(k)})$. As we increase k , we add restrictions and so the payoff decreases to some limiting level.

For each k approximation, we can find some aggregate distribution μ_k consistent with the moments \mathcal{K}_k . By Helly's theorem, there is a subsequence of the μ_k distributions with a limiting distribution; call the limiting aggregate distribution $\hat{\mu}$. This limiting distribution implies limiting moment values $\hat{\mathcal{K}}$ defined by $\hat{K}^{(j)} = \int_{\mathcal{A}} \varphi(a) d\hat{\mu}(a)$. These moment values $\hat{\mathcal{K}}$ achieve the limit of payoffs from $\mathcal{K}_k = (K_k^{(1)}, \dots, K_k^{(k)})$ as $k \rightarrow \infty$. ■

Proof of Lemma 2.

1. See Appendix B, in particular Corollary 1 part 2.
2. I begin by defining the concept of an assortative assignment, and I will then show that it is an aligned strategy for the agent to play assortative assignments.

Any aggregate distribution μ of mass N can be broken down into N distributions of mass 1 representing the lowest unit of measure $\hat{\mu}^{(1)}$, the second lowest unit $\hat{\mu}^{(2)}$, and so forth through the highest unit of measure $\hat{\mu}^{(N)}$. The distributions $\hat{\mu}^{(j)}$ have cdfs¹ given by

$$\hat{\mu}^{(j)}((-\infty, a]) = \llbracket (\mu((-\infty, a]) - (j - 1)) \rrbracket$$

where $\llbracket y \rrbracket$ is defined as 0 if $y < 0$; y if $y \in [0, 1]$; and 1 if $y > 1$. It holds that $\sum_j m^{(j)} = \mu$, and that $j < j'$ implies that $\max \text{Supp } \hat{\mu}^{(j)} \leq \min \text{Supp } \hat{\mu}^{(j')}$.

¹The notation $\hat{\mu}^{(j)}((-\infty, a])$ indicates the cdf evaluated at action a , i.e., the mass placed on actions in $(-\infty, a] \cap \mathcal{A}$. Likewise, $\mu((-\infty, a])$ indicates the cumulative mass function of μ , a measure of mass N .

Suppose that the agent observes states $(\theta_1, \dots, \theta_N)$. We can order the states according to a permutation π on $\{1, \dots, N\}$, so that $\pi(i) < \pi(i')$ implies $\theta_{\pi(i)} \leq \theta_{\pi(i')}$. (State θ_i is the $\pi(i)^{th}$ lowest). In the distributional quota characterized by aggregate distribution μ , I say that an assignment of actions to distributions (m_1, \dots, m_N) is *assortative* if $m_i = \hat{\mu}^{(\pi(i))}$.²

I now show that if players have increasing-difference utilities, then it is an aligned strategy in a distributional quota for the agent to assign actions assortatively.

Fix states $(\theta_1, \dots, \theta_N)$ and consider some non-assortative assignment $m = (m_1, \dots, m_N)$. Find some $\theta_i < \theta_j$ for which $\min \text{Supp } m_i \geq \max \text{Supp } m_j$. Then we can find measures $\bar{\nu} \leq m_i$ and $\underline{\nu} \leq m_j$,³ each placing a mass $\delta > 0$ on \mathcal{A} , such that the support of $\bar{\nu}$ is strictly above the support of $\underline{\nu}$. Consider swapping these measures, replacing the assignment m_i with $m'_i = m_i - \bar{\nu} + \underline{\nu}$ and m_j with $m'_j = m_j - \underline{\nu} + \bar{\nu}$ and holding all other assignments fixed. The payoff change to the agent is

$$\int_{\mathcal{A}} (U_A(a|\theta_i) - U_A(a|\theta_j)) d\underline{\nu}(a) - \int_{\mathcal{A}} (U_A(a|\theta_i) - U_A(a|\theta_j)) d\bar{\nu}(a)$$

By the intermediate value theorem, this is equal to

$$\delta [(U_A(\underline{a}|\theta_i) - U_A(\underline{a}|\theta_j)) + (U_A(\bar{a}|\theta_i) - U_A(\bar{a}|\theta_j))]$$

for some $\underline{a} \leq \bar{a}$, in the supports of $\underline{\nu}$ and $\bar{\nu}$ respectively. By increasing differences, this expression is nonnegative.

Starting from any non-assortative assignment, we can perform a sequence of such swaps to get to an assortative assignment. Each such

²If some states are equal, then there are distinct assortative assignments with respect to each ordering.

³The notation $\nu \leq m$ indicates that the measure of any set is weakly less in ν than m .

swap weakly increases payoffs, so the payoff from this resulting assortative assignment is at least as high as the payoff from the non-assortative one. And all assortative assignments are payoff equivalent, so they must give the agent an optimal payoff. (An optimal assignment exists by the argument in Lemma 1).

The principal also has increasing difference utility, so assortative assignments also maximize his payoff. Hence, this strategy is aligned. ■

Proof of Proposition 1. Given Theorem 1 and Lemma 2, it suffices to show that the aligned-optimal distributional quota has an aggregate distribution μ which is a sum of N degenerate distributions. Therefore this distributional quota induces the same outcomes as a ranking mechanism.

Let G^j be the ex ante distribution of the j^{th} lowest state, taking expectation over all realizations of θ . The principal's payoff from a distribution μ under the aligned (assortative) strategy is

$$\sum_{j=1}^N \left[\int_{\Theta} \int_{\mathcal{A}} U_P(a|\theta) d\hat{\mu}^{(j)} dG^j(\theta) \right]$$

For any distribution G^j , the bracketed expression is maximized over choice of $\hat{\mu}^{(j)}$ by a degenerate distribution placing all probability on some single action $b^{(j)} = \operatorname{argmax}_a \int_{\Theta} U_P(a|\theta) dG^j(\theta)$. (By increasing differences, we can find $b^{(j)}$ values which are increasing in j .) So we can find an aligned-optimal distributional quota by taking μ to be the sum of these N degenerate distributions. Under assortative assignments, the distributional quota characterized by aggregate distribution μ yields equivalent outcomes as the ranking mechanism characterized by action list $b^{(1)}, \dots, b^{(N)}$ under honest ranking. ■

Proof of Lemma 3.

1. See Appendix B, in particular Corollary 1 part 1.

2. Without explicitly solving for players' strategies, it suffices to show that any type of agent faces an identical maximization problem as she would if she shared the principal's utility. Consider a budget mechanism which restricts the expected sum of actions to K . Following the argument in Section II, the agent with bias λ observes states and then chooses action distributions to maximize

$$\begin{aligned} \mathbb{E} \left[\sum_i -(a_i - \theta_i - \lambda)^2 \right] &= \mathbb{E} \left[\sum_i (-(a_i - \theta_i)^2 + 2\lambda a_i - \lambda^2 - 2\lambda\theta_i) \right] \\ &= \mathbb{E} \left[\sum_i -(a_i - \theta_i)^2 \right] + 2\lambda K - \sum_i (\lambda^2 + 2\lambda\theta_i) \end{aligned}$$

where we substituted in the budget constraint $\mathbb{E}[\sum_i(a_i)] = K$. The agent's payoff is equal to the payoff of an agent with $\lambda = 0$, plus terms which are independent of the chosen actions. So the solution to her maximization problem is the same as if $U_A = U_P$. ■

Proof of Proposition 3. Follows from Lemma 3 and Theorem 1. ■

B Sufficient Conditions for Richness

In this section I give some conditions to guarantee that a utility set is Φ^J -rich, where the moment functions Φ^J are defined as $(\varphi^{(j)} \text{ s.t. } \varphi^{(j)}(a) = a^j)_{j \leq J}$.

Lemma 4.

1. Suppose there exist nonnegative integers $r < n$; a nonzero real number α ; a sign constant $s \in \{-1, 1\}$; a continuous function $\psi : \mathcal{A} \times \Theta \times \mathbb{R}_+$, written as $\psi(a|\theta; \lambda)$, of order λ^r ⁴ and a continuous function $\zeta : \Theta \times \mathbb{R}_+$,

⁴The function $\psi(a|\theta; \lambda)$ is of order λ^r if there exists $C > 0$ such that $|\psi(a|\theta; \lambda)| \leq C\lambda^r$ for each a and θ , for λ large enough.

written as $\zeta(\theta; \lambda)$, such that the agent's utility set \mathcal{U}_A contains functions

$$U(a|\theta) = \Psi(a|\theta; \lambda) + \zeta(\theta; \lambda) + \alpha \cdot (a + s\lambda)^{n+r+1}$$

for a sequence of real numbers $\lambda > 0$ going to infinity. Then \mathcal{U}_A is Φ^n rich.

2. If the above holds for arbitrarily large n , then \mathcal{U}_A is Φ^∞ rich.⁵

Proof.

1. Fix an arbitrary mechanism D . For an agent with preferences of the form above, we can rewrite the agent's payoff $\sum_i \mathbb{E}[U_A(a_i|\theta_i)]$ (from some strategy, starting at some point of the game) as

$$\sum_i \mathbb{E}[\Psi(a_i|\theta_i; \lambda) + \zeta(\theta_i; \lambda)] + \alpha \sum_{j=0}^{n+r+1} \lambda^{n+r+1-j} s^{n+r+1-j} \binom{n+r+1}{j} \sum_i \mathbb{E}[a_i^j]$$

where $\sum_i \mathbb{E}[a_i^j] = \sum_i \mathbb{E}[\varphi^{(j)}(a_i)]$ is the j^{th} moment value of Φ^J for $J \geq j$.

As λ goes to infinity, the agent puts increasing weight on the terms of higher order in λ . There are state-independent terms of order λ^{r+1} through λ^{n+r+1} , along with terms of order up to λ^r – including ψ – which may depend on the state. (The ζ term may always be ignored, because it does not depend on actions and so does not affect the agent's preferences).

The highest order term, of order λ^{n+r+1} , is the sum $\sum_i a_i^0 = N$ times a constant in a_i ; this sum is constant over action choices and so is irrelevant.

⁵The other parameters of the problem – $\psi, \zeta, s, r, \alpha$ – need not be the same across n .

The next term, of order λ^{n+r} , is a positive constant times $\text{sign}(\alpha) \cdot s^{n+r}$ times $\sum_i \mathbb{E}[a_i]$. So for large λ , the agent approximately maximizes $\text{sign}(\alpha) \cdot s^{n+r} \cdot \sum_i \mathbb{E}[a_i]$. Let $\bar{K}^{(1)}$ be defined as

$$\bar{K}^{(1)} = \begin{cases} \max_{\text{messages in } D} \sum_i \mathbb{E}[a_i] & \text{if } \text{sign}(\alpha) \cdot s^{n+r} = 1 \\ \min_{\text{messages in } D} \sum_i \mathbb{E}[a_i] & \text{if } \text{sign}(\alpha) \cdot s^{n+r} = -1 \end{cases}$$

Any strategy of the agent takes $\sum_i \mathbb{E}[a_i]$ close to $\bar{K}^{(1)}$, at all realizations of θ . Initial messages which do not allow for the possibility of values close to this are dominated by ones which do, for large enough λ . And conditional on such an initial message, under any state realization the agent will never choose an interim message which does not push λ close to this extreme.

The term of next highest order, λ^{n+r-1} , is a positive constant times $\text{sign}(\alpha) \cdot s^{n+r-1}$ times $\sum_i \mathbb{E}[a_i^2]$. So as $\lambda \rightarrow \infty$, conditional on setting $\mathbb{E}[a_i]$ close to $\bar{K}^{(1)}$, the agent then tries to maximize $\text{sign}(\alpha) \cdot s^{n+r-1} \cdot \sum_i \mathbb{E}[a_i]$. We proceed lexicographically – the agent first approximately fixes $\sum_i \mathbb{E}[a_i]$ at the appropriate extremal value, then $\sum_i \mathbb{E}[a_i^2]$, then $\sum_i \mathbb{E}[a_i^3]$, and so forth through $\sum_i \mathbb{E}[a_i^n]$, which has order $\lambda^{n+r+1-n} = \lambda^{r+1}$. Only after all of these moments are set at their extremal values does the agent consider the state of the world.

Formally, for each $j \leq n$, let $\bar{K}^{(j)}$ be defined inductively as

$$\bar{K}^{(j)} = \begin{cases} \max_{\text{messages in } D} \left\{ \sum_i \mathbb{E}[a_i^j] \mid \mathbb{E}[a_i^l] = \bar{K}^{(l)} \text{ for each } l < j \right\} & \text{if } \text{sign}(\alpha) \cdot s^{n+r-j} = 1 \\ \min_{\text{messages in } D} \left\{ \sum_i \mathbb{E}[a_i^j] \mid \mathbb{E}[a_i^l] = \bar{K}^{(l)} \text{ for each } l < j \right\} & \text{if } \text{sign}(\alpha) \cdot s^{n+r-j} = -1 \end{cases}$$

Any strategy of the agent takes $\sum_i \mathbb{E}[a_i] = \sum_i \mathbb{E}[\varphi^{(j)}(a_i)]$ close to $\bar{K}^{(j)}$ under any realization of states, for $j \leq n$.

Note that I take minimums and maximums rather than infimums and

supremums because I have assumed compactness of the space of assignments. I consider messages unconditional on states, rather than strategies conditional on states, because we can always choose a state-independent strategy to find extremal moments.

2. If such an infinite sequence of n 's exists, then I can find some subsequence of which induces identical values of $\bar{K}^{(j)}$ for each n and each $j \leq n$. (There are two possible signs of α , and two possible signs of s ; find some subsequence in which α and s always have the same sign). Therefore, for any ϵ and any n , I can find a utility function in \mathcal{U}_A which fixes each of the $j \leq n$ moments of the aggregate distribution to within ϵ of $\bar{K}^{(j)}$. ■

The lemma can now be applied to establish several economically relevant classes of utilities as rich with respect to appropriate moment functions.

Corollary 1.

1. *Any set of quadratic loss constant bias utilities is Φ^1 rich so long as the possible biases are unbounded (negative or positive).*
2. *The subset of increasing difference functions which are concave in actions is Φ^∞ -rich.*
3. *The set of generalized quadratic loss functions is Φ^∞ -rich. These are functions of the form $U_A(a|\theta) = -(c(a) - \theta)^2$ for some weakly increasing, continuous function $c(a)$.*

Any set containing a Φ -rich subset is itself also Φ -rich, so either part 2 or part 3 imply that the set of all increasing-difference functions is Φ^∞ -rich.

Proof.

1. Taking $\lambda > 0$, quadratic loss constant bias utilities are of the form $U_A(a|\theta) = -(a - \theta + s\lambda)^2$ for some sign constant s : positive biases

correspond to $s = 1$, negative to $s = -1$. This can be rewritten as $2a\theta - 2\theta\lambda - (a + s\lambda)^2$. The utility function is of the form in the proposition for $n = 1$, $r = 1$, $\alpha = -1$, $\psi(a|\theta) = 2a\theta$, and $\zeta(\theta; \lambda) = -2\theta\lambda$. So the utilities are Φ^1 -rich as long as there exist utilities with $\lambda \rightarrow \infty$.

2. Take $\Psi(a|\theta)$ any function which is concave in actions. Then $U_A(a|\theta) = \psi(a|\theta) + \alpha(a - s\lambda)^{2n}$ is concave in actions, for $\alpha = -1$ and for any s, λ, n .
3. Each $U(a|\theta)$ can be written as $-(c(a) - \theta)^2 = 2c(a)\theta - \theta^2 - (c(a))^2$. Consider increasing functions of the form $c(a) = (a - s\lambda)^n$, for n odd. Then $U(a|\theta) = 2(a - s\lambda)^n\theta - \theta^2 - (a - s\lambda)^{2n}$. This can be written as $\Psi(a|\theta; \lambda, n) - (a - s\lambda)^{2n}$ for $\Psi(a|\theta; \lambda, n) = 2(a - s\lambda)^n\theta - \theta^2$, which is indeed of order λ^n in λ and n . ■

Finally, I seek to show that the quadratic loss linear bias functions of the form $U_A(a|\theta) = -(a - \lambda^{(1)}\theta - \lambda^{(0)})^2$ are Φ^2 -rich, with $\lambda^{(0)} \in \mathbb{R}$ and $\lambda^{(1)} \in \mathbb{R}^+$. Although the argument for richness is similar to that in Lemma 4 above, this result is not a special case of the lemma.

Lemma 5. *The set of quadratic loss linear bias preferences is Φ^2 -rich.*

Proof. The agent's utility $-(a - \lambda^{(1)}\theta - \lambda^{(0)})^2$ can be written as

$$2\lambda^{(1)}\theta a - a^2 + 2\lambda^{(0)}a - [(\lambda^{(0)} + \lambda^{(1)}\theta)^2]$$

where the bracketed expression is independent of actions taken.

We can take $\lambda^{(0)} \rightarrow \infty$ and $\lambda^{(1)} \rightarrow 0^+$ in this class of utility functions, which – ignoring the action-independent component – takes the agent's utility to $-a^2 + 2\lambda^{(0)}a$. So the agent lexicographically maximizes $\sum_i \mathbb{E}[a_i] = \sum_i \mathbb{E}[\varphi^{(1)}]$, and then minimizes $\sum_i \mathbb{E}[a_i^2] = \sum_i \mathbb{E}[\varphi^{(1)}]$, as in the proof of Lemma 4.

This corresponds to

$$\begin{aligned}\bar{K}^{(1)} &= \max_{\text{messages}} \sum_i \mathbb{E}[a_i] \\ \bar{K}^{(2)} &= \min_{\text{messages}} \left\{ \sum_i \mathbb{E}[a_i^2] \middle| \sum_i \mathbb{E}[a_i] = \bar{K}^{(1)} \right\}.\end{aligned}$$

(Other limits than $\lambda^{(0)} \rightarrow \infty$ and $\lambda^{(1)} \rightarrow 0^+$ yield alternative values of \bar{K} .) ■

C Quadratic Loss Linear Biases

Under *quadratic loss linear bias* preferences, the principal has utility $U_P(a|\theta) = -(a-\theta)^2$ and the agent has utility $-(a - \lambda^{(1)}\theta - \lambda^{(0)})^2$ for some unknown parameters $\lambda^{(0)} \in \mathbb{R}$ and $\lambda^{(1)} \in \mathbb{R}_+$. This means that while the principal prefers $a_i = \theta_i$, the agent's ideal point is $a_i = \lambda^{(1)}\theta_i + \lambda^{(0)}$. The parameter $\lambda^{(1)}$ is the agent's relative sensitivity to state changes. A value above 1 corresponds to an extreme bias – when the state changes by a little, the agent's ideal point move a lot – and a value below 1 corresponds to a bias towards moderate actions which don't vary strongly in the state. This preference class has been used to model moderate- or extreme-biased agents in the delegation papers such as Melumad and Shibano (1991) and Alonso and Matouschek (2008).

Proposition 3. *If the players have quadratic loss linear bias preferences and the agent may have any linear biases, then the aligned-optimal Φ^2 -moment mechanism is max-min optimal.*

Proof. Lemma 5 in Appendix B confirms that the set of all quadratic loss linear bias preferences is Φ^2 -rich. So by Theorem 1, we need only show that Φ^2 -moment mechanisms satisfy aligned delegation under these preferences.

Expanding out the agent's payoff, in a Φ^2 mechanism where $\sum_i \mathbb{E}a_i =$

$K^{(1)}$ and $\sum_i \mathbb{E}(a_i)^2 = K^{(2)}$, the agent chooses an assignment to maximize

$$\begin{aligned} & \mathbb{E} \left[\sum_i -(a_i - \lambda^{(1)}\theta_i - \lambda^{(0)})^2 \right] \\ &= \mathbb{E} \left[2\lambda^{(1)} \sum_i a_i\theta_i + 2\lambda^{(0)} \sum_i a_i - \sum_i (a_i)^2 - \sum_i (\lambda^{(1)}\theta_i + \lambda^{(0)})^2 \right] \\ &= 2\lambda^{(1)} \mathbb{E} \left[\sum_i a_i\theta_i \right] + 2\lambda^{(0)} K^{(1)} - K^{(2)} - \sum_i (\lambda^{(1)}\theta_i + \lambda^{(0)})^2 \end{aligned}$$

So no matter what her preference parameters $\lambda^{(1)}$ and $\lambda^{(0)}$, the agent plays so as to maximize $\mathbb{E}[\sum_i a_i\theta_i]$. All agents play identically, as if they shared the principal's utility with $\lambda^{(0)} = 0$ and $\lambda^{(1)} = 1$. ■

D Generalization to Sequential Decisions

In the sequential environment, the principal and agent begin the game with a common prior over the joint distribution of states. The principal gives the agent a delegation mechanism, and then agent observes her utility function and sends an initial message. Then in each period $i = 1, \dots, N$ the agent observes θ_i before sending an interim message. After the interim message is sent, action a_i is drawn from a distribution which depends on the history of all messages (periods 1 through i) and actions (1 through $i - 1$). Once the action is taken, we move to period $i + 1$. The mechanism itself is an initial message space, a set of N interim message spaces, and a set of functions mapping histories into action distributions.

The mechanism form I consider is without loss of generality so long as the principal only learns about the realizations of past states, or the distributions of future states, from the agent's reports. This rules out a principal's observing his utility realization at decision i and using this information to alter the terms offered to the agent at decision $i + 1$. In an environment where the principal had more information, the mechanisms I consider would

still be implementable but might not be optimal.

In this sequential environment the common prior assumption is no longer simply for convenience. In deriving max-min optimal mechanisms, my proofs rely on the principal’s not wanting to overrule the agent when their utilities are appropriately aligned. But without common priors, even an agent who shared the principal’s utility function might play a mechanism differently than the principal wanted. For instance, a player facing a quota over actions might choose a high action today if she thought tomorrow’s state would be low, and a low action if she thought tomorrow’s state would be high. I assume common priors to guarantee that the principal would constrain an agent only to correct for perceived biases in preferences, not beliefs.⁶

D.1 Generalized Quadratic Loss Preferences

Example 3 illustrates that in sequential problems, increasing differences does not guarantee aligned delegation of distributional quotas. In the example, both players want to take high actions in high states and low actions in low states. But they disagree about whether $\theta = \frac{1}{2}$ is “high” or “low” relative to a uniformly drawn state over $\Theta = \{0, \frac{1}{2}, 1\}$. The principal’s payoff from $\theta = \frac{1}{2}$ is close to that from $\theta = 0$, and so he considers it to be $\theta = \frac{1}{2}$ to be lower than a uniform draw. The agent disagrees, considering $\theta = \frac{1}{2}$ to be close to 1 and therefore above a uniform draw.

To guarantee agreement in a distributional quota, players must agree on how to rank not just states but distributions of states. We get this when payoffs of the players are of the *generalized quadratic loss* functional form: $U(a|\theta) = -(c(a) - \theta)^2$ for some weakly increasing function $c(a)$.⁷ Under

⁶It would also be sufficient for the agent to be strictly better informed at the start of play, so that the principal would defer to the agent’s choices if utilities were identical.

⁷Notice that the “quadratic losses” are with respect to perturbations of the state, not of the action. So for any distribution of beliefs over the current state, a decisionmaker’s preferences over actions depend only on the expected state. This property makes these utilities natural for problems in which a principal elicits information on the state of the

these preferences, players rank distributions of states by their expectations. All players would agree that an observed state $\theta = \frac{1}{2}$ would be above any distribution with mean below $\frac{1}{2}$, and below any distribution with mean above $\frac{1}{2}$. This statement that preferences over actions depend only on expected states can be formalized in the following manner:

Lemma 6. *Take U a generalized quadratic loss utility function, and let a and θ be independent random variables. Then $\mathbb{E}[U(a|\theta)] = \mathbb{E}[U(a|\mathbb{E}[\theta])] - \text{Var}[\theta]$.*

(Proofs are deferred to the end of this section).

The class of generalized quadratic loss utilities includes the quadratic loss constant bias preferences as well as (up to an affine transformation) the quadratic loss linear biases.⁸ By choosing $c(\cdot)$ appropriately, we can model preferences with ideal actions equal to any increasing function of the state, not just linear functions. (The optimal action is always weakly increasing in the state because $c(\cdot)$ is increasing; these preferences satisfy increasing differences.)

In a sequential distributional quota, an agent with quadratic loss preferences will play what I call a “sequential-assortative” strategy. She tries to play low actions in low states and high actions in high states, subject to the information she has available at any time. The strategy is most clearly illustrated through an example.

Example 4. *Suppose that in each of N sequential periods, the state θ is drawn uniformly from $\Theta = [0, 1]$. The agent has some generalized quadratic loss utility, and is given a sequential quota requiring her to take each of N actions one time each.*

In the last period, the agent has a single action remaining, and plays this action independently of the state θ_N . The average state in which this action

world from better informed agents.

⁸A linear bias utility $-(a - \lambda^{(1)}\theta - \lambda^{(0)})^2$ can be written as $-(\lambda^{(1)})^2 (c(a) - \theta)^2$ for $c(a) = \frac{a}{\lambda^{(1)}} - \frac{\lambda^{(0)}}{\lambda^{(1)}}$, and the leading coefficient $(\lambda^{(1)})^2$ is irrelevant to preferences over actions.

is played is $\frac{1}{2}$.

At period $N-1$, the agent has two actions remaining. Under the sequential-assortative strategy, she plays the higher action if the current state θ_{N-1} is above $\frac{1}{2}$, and the lower action if θ_{N-1} is below $\frac{1}{2}$. Notice that the lower of the two actions is played with 50% probability in period $N-1$, at an average state of $\frac{1}{4}$; and with 50% probability in period N , at an average state of $\frac{1}{2}$. So the average state in which the action is played is $.5 \cdot \frac{1}{4} + .5 \cdot \frac{1}{2} = \frac{3}{8}$. Likewise, the higher action is played at an average state of $.5 \cdot \frac{3}{4} + .5 \cdot \frac{1}{2} = \frac{5}{8}$.

At period $N-2$, there are three actions remaining. The agent plays the lowest one if $\theta_{N-2} < \frac{3}{8}$; the middle one if θ_{N-2} is in $(\frac{3}{8}, \frac{5}{8})$; and the highest action if $\theta_{N-2} > \frac{5}{8}$. This leads to the lowest action's being played at an average state of $\frac{39}{128}$; the middle action at $\frac{1}{2}$; and the highest action at $\frac{89}{128}$.⁹

We can continue to backwards induct to find the agent's strategy at each period $i \leq N$.

In any sequential distributional quota, a principal and agent with generalized quadratic loss preferences will agree to play a sequential-assortative strategy, so the mechanism satisfies aligned delegation. Moreover, this preference class is Φ^∞ -rich and so the principal cannot profitably relax any of the restrictions of the sequential distributional quota.

Lemma 7.

1. *The set of generalized quadratic loss utilities is Φ^∞ -rich.*
2. *If the principal and agent have generalized quadratic loss preferences, then sequential distributional quotas (Φ^∞ -moment mechanisms) satisfy aligned delegation.*

⁹Calculating the expected state of the lowest action, there is a $\frac{3}{8}$ chance the action is taken in the current period at an expected state of $\frac{3}{16}$; there is a $\frac{5}{8}$ chance it is not taken today, in which case it will be the lower of the remaining two actions and will be taken at an expected state previously calculated to be $\frac{3}{8}$. We can calculate the other values similarly.

Just as in the simultaneous problem, the optimal aggregate distribution specifies that each of N actions are each to be taken once. So the aligned-optimal sequential distributional quota can be implemented as a (deterministic) sequential quota. The agent chooses actions without replacement from a list rather than reporting a distribution m_i from which each action is to be drawn.

Definition 1. *A sequential quota is characterized by a list of N actions, $b^{(1)} \leq b^{(2)} \leq \dots \leq b^{(N)}$ in \mathcal{A} . At the beginning period $i \geq 1$, there are $N - i + 1$ actions remaining on the list. After observing θ_i the agent chooses one action to be played, and then deletes this action from the list.*

Proposition 4. *Consider a sequential delegation problem. If the principal and agent have generalized quadratic loss preferences and the agent may have any generalized quadratic loss utility function, then the optimal sequential quota is max-min optimal.*

If the principal's utility is $U_P(a|\theta) = -(a - \theta)^2$ and if the set of actions \mathcal{A} contains the set of states Θ , the optimal sequential quota sets each action $b^{(j)}$ equal to the expected state in which the action is played. Consider the setting of Example 4 above. If $N = 1$, the optimal action list is $b^{(1)} = \frac{1}{2}$. If $N = 2$, the optimal action list is $b^{(1)} = \frac{3}{8}$ and $b^{(2)} = \frac{5}{8}$. If $N = 3$, the optimal action list is $b^{(1)} = \frac{39}{128} \simeq .305$, $b^{(2)} = \frac{1}{2}$, and $b^{(3)} = \frac{89}{128} \simeq .695$.

Compare this to an otherwise identical simultaneous problem. The optimal ranking mechanism would specify the same action $b^{(1)} = \frac{1}{2}$ if $N = 1$. But for $N = 2$, the principal would choose $b^{(1)} = \frac{1}{3}$ and $b^{(2)} = \frac{2}{3}$. For $N = 3$, he would choose $b^{(1)} = \frac{1}{4}$, $b^{(2)} = \frac{1}{2}$, and $b^{(3)} = \frac{3}{4}$. Relative to the simultaneous actions, the sequential actions are pinched in towards $\frac{1}{2}$.

Proof of Lemma 6.

$$\begin{aligned} \mathbb{E}[-(c(a) - \theta)^2] &= \mathbb{E} [(-c(a)^2 + \mathbb{E}[\theta]c(a) - \mathbb{E}[\theta]^2) + \mathbb{E}[\theta]^2 - \mathbb{E}[\theta^2]] \\ &= \mathbb{E} [-(c(a) - \mathbb{E}[\theta])^2] - \text{Var}[\theta]. \end{aligned}$$

■

Proof of Lemma 7. 1. See Appendix B, in particular Corollary 1 part 3.

2. I first introduce some notation to allow us to define the sequential-assortative strategy, which I will show is optimal. Under this strategy, the agent will sort the current state with respect to certain expected values of the future states.

For a finite list of real numbers L (possibly with duplicates), let $R^{(i)}(L)$ be the i^{th} lowest element of L . So $R^{(1)}(L)$ is the minimum of L , $R^{(2)}(L)$ is the value of the second lowest element, et cetera.

Now define a function $\tilde{\theta}_i : \{1, \dots, N - i + 1\} \times \Theta^{i-1} \rightarrow \mathbb{R}$ by backwards induction. We will interpret $\tilde{\theta}_i(j)$ as *the expected value of the state in which the j^{th} lowest remaining action will be played*. The expectation is taken with respect to the common prior over the joint distribution of states at the beginning of period i , conditional on the observations of θ_1 through θ_{i-1} . For $i = N$, let $\tilde{\theta}_N(1; \theta_1, \dots, \theta_{N-1}) = \mathbb{E}_{\theta_N} [\theta_N | \theta_1, \dots, \theta_{N-1}]$. For $i < N$, given the function $\tilde{\theta}_{i+1}$, let

$$\tilde{\theta}_i(j; \theta_1, \dots, \theta_{N-1}) = \mathbb{E}_{\theta_i} \left[R^{(j)} \left(\left\langle \theta_i, \tilde{\theta}_{i+1}(1), \dots, \tilde{\theta}_{i+1}(N - i) \right\rangle \right) \right].$$

I write $\tilde{\theta}_i(j; \theta_1, \dots, \theta_{N-1})$ as $\tilde{\theta}_i(j)$ if the past states are otherwise implied.

By the quadratic loss utility function, where preferred actions depend only on expected states, the agent will want to assign action a_i assortatively *as if* true future states were known to be $\tilde{\theta}_{i+1}(1), \dots, \tilde{\theta}_{i+1}(N - i)$. At the start of period i , there is a *remaining distribution* $\mu_i = \mu - (m_1 + \dots + m_{i-1})$ which is the sum of the distributions of actions which have yet to be played. The agent chooses m_i as in an assortative strategy in a ranking mechanism over $N - i + 1$ actions, with aggregate distribution μ_i and state realizations $\theta_i, \tilde{\theta}_{i+1}(1), \dots, \tilde{\theta}_{i+1}(N - i)$. This defines the *sequential-assortative* strategy.

Claim 1. Fix some generalized quadratic loss utility function U and some sequential decision problem. There exists $C \geq 0$ such that for all measures μ of mass N on \mathcal{A} , a sequential-assortative strategy in a distributional quota with aggregate distribution μ gives a player with utility U a payoff of $\sum_{j=1}^N \int_{\mathcal{A}} U(a|\tilde{\theta}_1(j))d\hat{\mu}^{(j)}(a) - C$. Any alternative strategy gives a weakly lower payoff. Here the notation $\hat{\mu}^{(j)}$ indicates the distribution on \mathcal{A} corresponding to the j^{th} lowest unit from μ , as in the proof of Proposition 1.

The C term corresponds to the quadratic payoff loss due to the variance of states away from their expectations, as in Lemma 6. It depends on the distribution of states, but not on the chosen measure.

This claim completes the proof of the lemma; under generalized quadratic loss preferences, sequential-assortative strategies are optimal for both the principal and the agent.

Proof of Claim 1. I will prove this by backwards induction on the number of periods remaining. Consider period i , prior to the realization of θ_i , with remaining distribution μ_i . I seek to show that a sequential-assortative strategy gives $\sum_{j=1}^{N-i+1} \int_{\mathcal{A}} U(a|\tilde{\theta}_i(j))d\hat{\mu}_i^{(j)}(a) - C_i$, for C_i independent of μ_i , and other strategies give weakly less.

For $i = N$, this holds by Lemma 6. There is a remaining distribution μ_N of mass 1, and all strategies give a payoff of $\int_{\mathcal{A}} U(a|\tilde{\theta}_N(1))d\hat{\mu}_N^{(1)}(a) - \text{Var}[\theta_N]$, where $\tilde{\theta}_N(1)$ is the expected value of θ_N and $\hat{\mu}_N^{(1)} = \hat{\mu}_N$.

Suppose the claim holds for $i + 1$; I want to show that it holds for i as well.

Given some θ_i ranked k^{th} lowest of the expected future states, the payoff

of a sequential-assortative strategy is (by the inductive hypothesis)

$$\sum_{j=1}^{N-i+1} \int_{\mathcal{A}} \left(\begin{cases} U(a|\tilde{\theta}_{i+1}(j)) & \text{if } j < k \\ U(a|\theta_i) & \text{if } j = k \\ U(a|\tilde{\theta}_{i+1}(j-1)) & \text{if } j > k \end{cases} \right) d\hat{\mu}_i^{(j)}(a) - C_{i+1}. \quad (1)$$

The sum of integrals is exactly just the payoff of a simultaneous probability assignment mechanism of an assortative assignment over $N-i+1$ states, given measure μ_i and states $\theta_i, \tilde{\theta}_{i+1}(1), \dots, \tilde{\theta}_{i+1}(N-i)$. Then the C_{i+1} term lowers payoffs due to uncertainty over future states. Taking expectation over θ_i , the expected value of the state which is integrated over $\hat{\mu}_i^{(j)}$ in (1) is $\tilde{\theta}_i(j)$. So applying Lemma 6, (1) becomes

$$\sum_{j=1}^{N-i+1} \int_{\mathcal{A}} U(a|\tilde{\theta}_i(j)) d\hat{\mu}_i^{(j)}(a) - C_i$$

for C_i equal to $\mathbb{E}[C_{i+1}]$ minus the sum of the variance constants.¹⁰

Finally, I seek to show that the payoff of a sequential-nonassortative strategy is weakly less than this. By the inductive hypothesis, given any state θ_i and any assignment m_i in period i , it is optimal to revert to a sequential-assortative strategy at $i+1$. This gives a payoff from current and future periods equal to that from simultaneous distributional quota with measure μ_i and states $(\theta_i, \tilde{\theta}_{i+1}(1), \dots, \tilde{\theta}_{i+1}(N-i))$, if the agent assigns m_i to state θ_i and assigns the rest of the probability mass assortatively; minus the constant C_{i+1} . The simultaneous payoff would be maximized by assortative m_i . This corresponds to maximizing the sequential payoff by choosing m_i sequential-assortatively. \square

¹⁰ C_{i+1} may depend on the θ_i to the extent that the joint distribution of future states depends on this realization.

Proof of Proposition 5. By the sequential analog of Theorem ?? combined with Lemma 7, it suffices to show that the aligned-optimal distributional quota can be implemented by a sequential quota.

By Claim 1 in the proof of Lemma 7 part 2, the principal’s payoff from a sequential distributional quota with aggregate distribution μ can be expressed as the payoff in a simultaneous distributional quota with aggregate distribution μ (under some particular distribution of states) minus a constant independent of μ . And by Proposition 1, we know that there is an optimal aggregate distribution which is a sum of N degenerate distributions – actions taken with certainty. Under sequential-assortative play, each of these actions will be played with certainty in one period. So the outcome is identical to a sequential quota. ■

D.2 Quadratic Loss Constant Bias Preferences

Under quadratic loss constant bias preferences, the argument that budget mechanisms (sequential or otherwise) satisfy aligned delegation goes through unchanged. Without explicitly solving for the agent’s strategy, any agent’s maximization problem is identical to that of an unbiased agent. Her lifetime payoff going forward from any point is equal to the principal’s payoff, plus some terms that don’t depend on her choices. So sequential budgets continue to satisfy aligned delegation.

Definition 2. *A sequential budget is characterized by a number $K \in \mathbb{R}$. At each interim period, the agent chooses an action or distributions of actions. These choices are required to satisfy $\sum_i \mathbb{E}[a_i] = K$.*

Proposition 5. *Consider a sequential delegation problem. If the players have quadratic loss constant bias utilities and the agent may have any bias, then the optimal sequential budget is max-min optimal.*

Proof. Follows from the argument of the proof of Proposition 3. ■

E Nonseparable Preferences

In this section I show that ranking and budget mechanisms can be max-min optimal even if the players have non-additively separable preferences.

E.1 Preferences over the distribution of actions

Suppose that a teacher and school care about each student's grade conditional on his or her performance, as before, but have an additional preference over the distribution of grades itself. Third parties may judge the school and/or the teacher poorly if the distribution of grades is far from the norm, even if all students deserve their grades individually. We can model this with payoffs of the following form:

$$\begin{aligned} \text{Principal: } & \left(\sum_i U_P(a_i|\theta_i) \right) + W_P(\mathbf{a}) \\ \text{Agent: } & \left(\sum_i U_A(a_i|\theta_i) \right) + W_A(\mathbf{a}) \end{aligned}$$

where W_P and W_A are symmetric across decisions, i.e., invariant to permutations of actions. For instance, the W terms can put arbitrary penalties on distributions with low variance (not enough grade dispersion); distributions with high means (looks like grade inflation); distributions with too many grades below a threshold (parents complain); etc.

In a ranking mechanism (a strict grading curve), the value of $W_A(\mathbf{a})$ would be fixed in advance. So an agent with increasing-difference utility U_A would rank honestly, exactly as with separable preferences. If the principal utility U_P also satisfied increasing differences, then the ranking mechanism would satisfy aligned delegation. And under some minor regularity conditions, ranking mechanisms continue to be max-min optimal when the agent may have any increasing-difference utility.

Proposition 6. *Fix an increasing difference principal utility U_P , and symmetric payoff functions W_P and W_A over the vector of actions. If the agent may have any increasing difference utility function U_A , and if the ex ante probability that $\theta_i = \theta_j$ is zero for any $i \neq j$, then a ranking mechanism is max-min optimal.*

The result holds for any W_A , and therefore also holds if the principal is uncertain over W_A .

The formal argument for the max-min optimality of ranking mechanisms in the body of the paper (following Theorem 1) does not quite go through. I previously showed that (i) given an arbitrary delegation mechanism D , we can find some agent type for which the principal prefers a specific distributional quota; (ii) all types play this distributional quota identically, so it is a max-min improvement on D ; and (iii) a ranking mechanism implements the aligned-optimal distributional quota, which improves on the previous quota for all types. In this new context, steps (ii) and (iii) fail. Not all types play a distributional quota identically: the principal or agent's preferred play might no longer be assortative, depending on the W function.¹¹ For some agent types, a ranking mechanism may sometimes give worse payoffs than an alternate distributional quota. However, we can recover the result as long as we can still show that there exists an agent type for which the principal-optimal ranking mechanism (in which all types do play identically) improves on the arbitrary mechanism D .

Sketching out an informal proof, consider agent types for whom the payoffs from W_A are very unimportant relative to those from U_A .¹² These agent

¹¹In fact, such quotas should now be generalized to allow for the choice of *joint* distributions over actions, since marginal distributions no longer determine payoffs. Suppose that there are $N = 2$ decisions, two possible actions 0 or 1, and W gives a large payoff if $a_1 = a_2$. Then given a distributional quota that action 0 and 1 each must be taken once, the player might prefer a 50/50 draw over $a_1 = a_2 = 0$ and $a_1 = a_2 = 1$ instead of assigning the low state to action 0 and the high state to action 1.

¹²In other words, fix W_A , and fix some function $U : \mathcal{A} \times \Theta \rightarrow \mathbb{R}$ satisfying strict increasing differences. Now take U_A equal to $\alpha \cdot U$ for $\alpha \rightarrow \infty$. This agent plays approximately as

types will play an arbitrary mechanism as if they had preferences from the body of the paper, at least outside of the zero-probability event that two states are equal, and thus we know that there are extreme such agent types against whom ranking mechanisms are unimprovable.¹³ Because all agents play identically in this ranking mechanism, the mechanism is worst-case optimal.

E.2 Preferences over the sum of actions

A more structured form of nonseparable preferences occurs when a firm or manager cares about each investment conditional on a project's productivity, and also about the total amount of money invested. In particular, suppose that the player faces increasing marginal costs of capital. When there is a large investment in project 1, the firm or agent wants to keep the investment low for project 2 because capital has grown expensive. Using the quadratic loss constant bias utilities, we can model preferences as:

$$\begin{aligned} \text{Principal: } & - \sum_i (a_i - \theta_i)^2 + W_P \left(\sum_i a_i \right) \\ \text{Agent: } & - \sum_i (a_i - \theta_i - \lambda_i)^2 + W_A \left(\sum_i a_i \right) \end{aligned}$$

where W_P and W_A are weakly concave, implying convex costs.

All agent types will play a budget mechanism identically as long as actions will be deterministic and thus $\sum_i a_i$ is fixed in advance at the budget level. And concave W functions punish randomization even more than quadratic losses already do (which build in a cost of variance). So we can guarantee

though she is maximizing $\sum_i U(a_i|\theta_i)$, ignoring W_A .

¹³When two states are exactly equal, the agent will resolve indifferences in favor of maximizing W_A , no matter how unimportant W_A is relative to the U_A terms. If W_P and W_A are aligned, this could lead the agent to play in a manner that improves the principal's payoffs relative to strict ranking; if misaligned, it could hurt the principal's payoffs.

that players would always choose deterministic actions as long as the action set \mathcal{A} is a connected interval. (With disconnected action sets, randomization may be unavoidable when trying to play an action between two possibilities.)

Proposition 7. *Suppose the principal and agent have quadratic loss constant bias utilities with concave payoff functions W_P and W_A over the sum of actions. If the agent may have any bias λ , and if the action set is an interval, then a budget mechanism is max-min optimal.*