

Aligned Delegation

Alexander Frankel*

Chicago Booth

September 2011

Abstract

A principal delegates multiple decisions to an agent, who has private information relevant to each decision. The principal is uncertain about the agent's preferences. I solve for max-min optimal mechanisms – those which maximize the principal's payoff against the worst agent preference type. These mechanisms are characterized by a property I call “aligned delegation.” In an aligned delegation mechanism all agent types play identically, as if they shared the principal's preferences.

Max-min optimal mechanisms may take the simple forms of simultaneous ranking mechanisms, sequential quotas, or budgets. This work motivates the use of these contracts.

1 Introduction

Consider a problem in which a principal (he) delegates a number of decisions to an agent (she). A school requires a teacher to assign grades to all of her students; a firm appoints a manager to determine investment levels in different projects; an organization asks a supervisor to evaluate her employees and make promotion and firing decisions, or give out bonuses. In each of these cases,

*I'd like to thank my advisers Bob Wilson and Andy Skrzypacz for their support, with a special thanks to Bob for the close guidance he has given me throughout this project. I also thank Manuel Amador, Attila Ambrus, Aaron Bodoh-Creed, Renee Bowen, Jeremy Bulow, Matt Elliott, Ben Golub, Bengt Holmstrom, Yair Livne, Michael Ostrovsky, and Ilya Segal for their input, as well as seminar audiences at Stanford and CIMAT. This work was supported by a grant from the SIEPR Dissertation Fellowship program.

the principal relies on the agent because she observes “states of the world” relevant to the principal’s preferences. The teacher knows how well students have done in the class; the manager sees the quality of potential investments; the supervisor has observed the performance of her employees. How should the principal choose a delegation rule that specifies which actions the agent may take?

If the principal and agent had identical preferences, there would be no reason to restrict the agent’s choices. However, preferences may only be partially aligned. For instance, a teacher and school agree that better students should receive higher grades. But the teacher may be biased towards low grades relative to the school’s wishes, or high grades, or something more complicated – failing too many students while giving out too many A’s, say. Giving the agent more freedom lets her make better use of her private information, but it also gives leeway for a biased agent to take advantage of the principal.

I assume that the principal has beliefs about the distribution of states (student performance levels), but has limited information about the agent’s bias – her utility function or “type” mapping states and actions into payoffs. He may only know that her utility satisfies a certain property, or is in some set. The principal seeks a robust mechanism which will work well for any agent type.

Formally, I model this robustness by searching for a *max-min* optimal mechanism. For any set of agent types and for any mechanism the principal suggests, we can find the principal’s minimum expected payoff over all agent types. A max-min optimal mechanism maximizes this worst-type payoff.¹

A “simultaneous” problem is one in which the agent observes all of the underlying states before any actions are taken: a teacher sees all of her students’ test scores before assigning grades. In a simultaneous problem, suppose

¹It might be difficult for the principal to express a prior belief over the distribution of agent types if types come from a highly dimensional set. Even given a well-specified prior, the standard tools of dynamic Bayesian optimal mechanism design – see for example Pavan, Segal, and Toikka (2010) – assume that the agent has one-dimensional private information at every period and that transfer payments may be used. Neither holds in this model. The agent may have many-dimensional private information on her utility function and on the observed states of the world.

that the principal and agent both prefer higher actions in higher states (higher grades for better students), formalized as an increasing difference condition on utility functions. If the principal believes that the agent might have any increasing difference utility, then a *ranking mechanism* is max-min optimal. The agent is only asked to rank states from lowest to highest. The decision with the lowest state is assigned to some predetermined low action, the next higher state is assigned to a higher action, etc. This corresponds to a “strict grading curve” where the top 10% of students get an A and the next 15% get an A-, or a bonus rule which gives \$50,000 to the best-performing employee in a group and \$40,000 to the next one.

In a “sequential” problem, each action is chosen before the next state is observed: the supervisor evaluates one employee and gives her a bonus before evaluating the next employee. Under a stronger assumption than increasing differences, that the principal and agent utilities share a quadratic loss functional form, the max-min optimal mechanism is a *sequential quota*. The agent is given a list of actions to assign to the decisions. Each period she observes the state then chooses an action from the list, without replacement.

Ranking mechanisms and quotas are both special cases of *probability assignment* mechanisms, which I introduce in Section 3. These contracts let the agent take any actions she wants so long as each action is ultimately played an appropriate number of times. The agent is allowed to assign a probability distribution over actions to a decision; the “number of times” corrects for stochastic actions. Theorem 1 establishes that probability assignment mechanisms are max-min optimal if two conditions hold. First, the agent has a *rich* set of possible utilities which contains certain extreme preferences. Second, preferences are *PA-aligned* – i.e., probability assignment mechanisms satisfy the property of *aligned delegation*. A mechanism is aligned delegation if every agent type plays as though she were maximizing the principal’s payoff. Think of the ranking mechanism – any increasing-difference agent type submits honest rankings, just as the principal wants her to do.

Section 4 applies this theorem to show that ranking is max-min optimal in simultaneous problems with increasing difference preferences. Section 5 applies

it to show that quotas are optimal in sequential problems with quadratic loss preferences.

Section 6 shows that in the absence of richness, we may be able to do better than probability assignment by giving the agent more flexibility. Suppose the agent has an unknown constant bias relative to the principal, modeled with a quadratic loss functional form. Then for a simultaneous or sequential problem, a *budget* mechanism is max-min optimal. The agent is allowed to choose any actions which sum to some specified level. This corresponds to a grading curve in which a teacher can give out any grades so long as the class GPA is 3.0, or a bonus pool where a supervisor can divide \$150,000 among her employees as she sees fit. This budget mechanism would do poorly if the agent were inclined towards moderate actions (give every student a B), or towards extreme actions (a lot of A's, a lot of D's). But with a constant bias, it satisfies aligned delegation. The principal and agent may disagree about their desired average action, but for any fixed average they agree about what actions to take.

If the agent may have extreme or moderate preferences, modeled as quadratic loss preferences with an unknown linear bias, then the principal should use a *two-moment* mechanism. The agent chooses actions to fit a predetermined sum and sum-squared – i.e., a fixed mean and variance.

Extensions are considered in Section 7, and Section 8 concludes. All proofs omitted from the body of the paper are found in Appendix B.

Literature Review

A max-min optimality criterion, as opposed to a Bayesian one, is rare in the theory of contracting and mechanism design. One notable early exception is Hurwicz and Shapiro (1978), which shows that a 50% tax may be a max-min optimal sharecropping contract. More recently, Satterthwaite and Williams (2002) justifies double-auctions as worst-case asymptotic optimal in terms of efficiency loss. Other applications of max-min in economics include behavioral analyses of ambiguity aversion (see Gilboa and Schmeidler (1989) for an axiomatization) and macroeconomic work on robust control (see Sargent and

Hansen (2007)). In computer science, algorithms are commonly evaluated by their max-min or worst-case performance; this approach has been applied to auction theory in work reviewed by Hartline and Karlin (2007).

I borrow the basic set-up of my stage game from the literature on the delegation problem, introduced by Holmström (1977). An agent is privately informed about a state of the world which affects both her own preferences over actions and a principal's. The principal "delegates" the decision by specifying a set of actions from which the agent may choose. Actions and states are elements of the real line and contracting is done through restrictions on actions, without transfer payments conditional on actions or on outcomes.

I diverge from the delegation literature by considering multiple decisions and uncertainty over the agent's utility function, while most previous work looks at a single decision and a commonly known agent utility function. See for example Holmström (1977, 1984), Melumad and Shibano (1991), Martimort and Semenov (2006), Alonso and Matouschek (2008), Kovac and Mylovanov (2009), and Amador and Bagwell (2010). These papers develop methods for deriving an optimal one-dimensional delegation set under various assumptions on utility functions and state distributions. A common goal is to find conditions under which interval delegation is optimal.

One paper which looks at uncertain agent preferences in the context of a single decision is Armstrong (1995). Armstrong restricts attention to interval delegation sets which do not vary with the agent's type, and does comparative statics on the endpoints of the interval. Athey et al. (2005) and Amador et al. (2006) consider what are effectively sequential delegation problems over multiple decisions in which the principal and agent share a commonly known stage utility, but the agent has time-inconsistent preferences. They find conditions under which interval delegation is optimal in each period – the agent is allowed to choose any actions below a cutoff level. In sequential problems where the agent has state-independent preferences – she only cares about the actions which are taken – Frankel (2010a) and Malenko (2011) derive forms of quotas and budgets, respectively, as optimal contracts. Frankel (2010a) assumes that the principal does not know agent preferences precisely, and as in the current

paper this prevents the principal from using more flexible contracts.

The issue of eliciting information over time from an agent with a hidden bias has also been addressed in the study of cheap talk, wherein the principal cannot commit to a mechanism. In Sobel (1985), Benabou and Laroque (1992), and Morris (2001) the agent is altruistic with some probability, and otherwise has some specified bias. In equilibrium, agents may shade their reports in early periods in order to earn the principal's trust later.

Finally, this paper is related to a body of literature which shows how to put together multiple independent decisions to improve on the outcome of a single decision. See Jackson and Sonnenschein (2007), Escobar and Toikka (2009), and Cohn (2010) for analyses of allocation problems where each player has private information on her own values, or Chakraborty and Harbaugh (2007), Chakraborty and Harbaugh (2010), and Frankel (2010a) for settings where an informed agent has information on a principal's preferences. These papers construct mechanisms or equilibria which yield high payoffs but (except for Frankel (2010a)) are not necessarily optimal. For instance, Jackson and Sonnenschein (2007) and Chakraborty and Harbaugh (2007) show that quotas and ranking mechanisms may achieve approximately efficient or first-best payoffs when there are many independent and ex ante identical decisions. While the focus of the current paper is on max-min optimality, I extend the results of approximately first-best payoffs to my environment in Appendix A.2.

2 Benchmark Model: Simultaneous Decisions

2.1 Players and Payoffs

A decision problem is comprised of $N < \infty$ decisions, indexed by $i = 1, 2, \dots, N$. For decision i , a state of the world $\theta_i \in \Theta \subseteq \mathbb{R}$ is realized and then an action $a_i \in \mathcal{A} \subseteq \mathbb{R}$ is taken. In the benchmark model, I consider a simultaneous environment in which all states are realized and then all actions are taken. Later sections allow for sequential decisions, where one action is taken before the next state is realized.

There are two players, a principal and an agent. Each player's payoff depends jointly on actions and states, but only the agent observes the states of the world. After state θ_i is realized and action a_i is taken, the principal gets a stage utility for decision i of $U_P(a_i|\theta_i)$ and the agent gets $U_A(a_i|\theta_i)$. The lifetime payoff of each player is the sum of the stage utilities:

$$\text{Principal: } \sum_{i=1}^N U_P(a_i|\theta_i) \quad \text{Agent: } \sum_{i=1}^N U_A(a_i|\theta_i)$$

The sets \mathcal{A} and Θ of actions and states are taken to be compact, i.e., closed and bounded subsets of the real line. The utility functions U_A and U_P are continuous maps from $\mathcal{A} \times \Theta$ into \mathbb{R} ; I call the set of such functions \mathcal{U} .

The principal has a prior belief over the joint distribution of states $(\theta_1, \dots, \theta_N)$, and he knows his own utility function $U_P \in \mathcal{U}$. He is uncertain about the agent's utility function. He only knows a set $\mathcal{U}_A \subseteq \mathcal{U}$ from which her utility is drawn. In my analysis the principal need not have a prior belief about the distribution of U_A over \mathcal{U}_A .

It will be useful to define a notion of equivalence for utility functions:

Definition (Equivalent Utilities). Two stage utility functions U and \tilde{U} in \mathcal{U} are *equivalent* if there exist a positive constant $\zeta \in \mathbb{R}_{++}$ and a function $b: \Theta \rightarrow \mathbb{R}$ such that $U(a|\theta) = \zeta \tilde{U}(a|\theta) + b(\theta)$ for all a and θ .

Equivalent utility functions imply identical preferences over actions.² In later sections when I discuss functional forms of principal or agent utilities, the conclusions should be understood to generalize to the full equivalence classes.

2.2 Timing of the Game

As mentioned above, only the agent observes the states of the world which affect both his and the principal's preferences over actions. The principal's role is to write a mechanism, i.e., a set of rules for the agent. He wants these rules to induce the agent to choose actions which are good for the principal.

²The additive term $b(\theta)$ is exogenous to the chosen actions, while the constant ζ uniformly rescales the action-dependent component of utility.

After the principal decides on a mechanism, the agent observes the states of the world and sends messages according to the given rules. These messages determine the subsequent actions or, for stochastic mechanisms, the distributions over actions. In this benchmark simultaneous environment, the agent observes and sends messages about all states before any actions are taken.

I assume that the agent must participate in any mechanism which the principal designs – there are no “individual rationality” constraints.

The only output of the mechanism is the determination of the actions taken. In particular, there are no transfer payments.³

The only inputs are the agent’s reports. Any outside information regarding the values of states – e.g., the principal’s utility realizations – is noncontractible. This can be thought of as a restriction on the information available to the principal, or as a restriction on the set of mechanisms considered.

Formally, a mechanism $D = (\mathcal{M}_0, \mathcal{M}, M)$ is

- an initial message space \mathcal{M}_0 and an interim message space \mathcal{M} ; and
- a map M from pairs of messages (in $\mathcal{M}_0 \times \mathcal{M}$) into joint distributions over actions (in $\Delta(\mathcal{A}^N)$, where $\Delta(\cdot)$ represents the set of Borel measurable distributions). For simplicity I will describe these maps only by their marginal distributions over actions; by additive separability, the marginals entirely determine payoffs.

The mechanism induces the following (single-player) game for the agent:

1. The agent observes $U_A \in \mathcal{U}_A$.
2. The agent sends initial report $m_0 \in \mathcal{M}_0$.
3. The agent observes states $\boldsymbol{\theta} = (\theta_1, \dots, \theta_N)$.
4. The agent sends interim report $m \in \mathcal{M}$.

³See Krishna and Morgan (2008) for a delegation model with limited-liability monetary payments, or Ambrus and Egorov (2009) and Amador and Bagwell (2010) for models with nonmonetary punishments conditional on actions taken. Frankel (2010a) shows how uncertainty over payoffs in a model with state-independent preferences can make monetary incentives infeasible.

5. Actions $\mathbf{a} = (a_1, \dots, a_N)$ are drawn from the distribution $M(m_0, m)$.⁴

A direct mechanism, for instance, would have $\mathcal{M}_0 = \mathcal{U}_A$ and $\mathcal{M} = \Theta^N$.

An agent's pure reporting strategy σ is an initial message m_0 and a function mapping state vectors into interim messages m . After observing U_A , the agent chooses an optimal (sequentially rational) reporting strategy σ to maximize her expected lifetime utility going forward from each information node. Let Σ^D be the set of all possible reporting strategies, and let $\Sigma^{*D}(U_A) \subseteq \Sigma^D$ be the set of optimal strategies for an agent with utility $U_A \in \mathcal{U}$.⁵

If a principal proposes mechanism D and strategy σ is chosen by the agent, then a player with stage utility $U \in \mathcal{U}$ gets a lifetime expected payoff of

$$\mathbb{E}_{\mathbf{a}, \boldsymbol{\theta}} \left[\sum_i U(a_i | \theta_i) \mid D, \sigma \right]$$

The notation $\mathbb{E}_{\mathbf{a}, \boldsymbol{\theta}}$ signifies that expectation is taken with respect to the exogenous states $\boldsymbol{\theta}$ as well as the actions \mathbf{a} , which – depending on D , $\boldsymbol{\theta}$, and σ – may be stochastic.

The mechanisms as described above do not include all possible indirect mechanisms. However, a revelation principle applies (see, e.g., Myerson (1986)). Any equilibrium of any indirect mechanism without additional informational inputs could be replicated by the truthful equilibrium of an incentive compatible direct mechanism of the form above. The mechanisms I consider are without loss of generality in the sense that they include direct mechanisms.

2.3 Max-Min Optimality

This paper will be primarily concerned with solving for max-min optimal mechanisms – those which maximize the principal's payoff against the worst possible

⁴None of the results of the paper would change if the agent observed information about $\boldsymbol{\theta}$ prior to any reports. See Section 5 for a discussion of the sequential timing, in which θ_i is observed, interim message i is sent, action a_i is taken; and only then is θ_{i+1} observed.

⁵As a matter of notation, I say that U_A is observed and then a reporting strategy is chosen; the reporting strategy does not itself take U_A as an argument. This terminology allows me to talk about different types' playing the same strategy. In more standard notation, a "strategy" would be a map from \mathcal{U}_A into Σ^D .

agent type.

Definition (Max-min Optimality). Say that a mechanism is *max-min optimal* over a set of agent utilities $\mathcal{U}_A \subseteq \mathcal{U}$ if it is an arg max of the following problem:

$$\max_{\text{Mechanisms } D} \left[\inf_{U_A \in \mathcal{U}_A} \left[\max_{\sigma \in \Sigma^{*D}(U_A)} \mathbb{E}_{\theta, \mathbf{a}} \left[\sum_i \gamma_i U_P(a_i | \theta_i) \mid \sigma, D \right] \right] \right]$$

The worst case is taken over utility realizations, not state realizations. In the case of multiple optimal strategies for an agent, I look at the one preferred by the principal – this is the second “max” in the definition.

The max-min problem can be thought of one in which the principal first picks a mechanism D . Given this mechanism, an adversary or “devil” chooses an agent utility type $U_A \in \mathcal{U}_A$ so as to minimize the principal’s expected payoff. Then states are realized, and the agent plays a strategy σ which is optimal for her type U_A .

I have exogenously assumed that money is not used, but in a max-min sense money would not help the principal. Even if the principal knew U_A precisely, he would not be able to use monetary bonuses effectively without knowing the tradeoff of money against action utility. In the extreme cases a righteous teacher would ignore monetary incentives in order to do right by her students; an apathetic teacher would first maximize her bonus, and only then consider student performances.

3 Probability Assignment Mechanisms

In this section I define a class of “probability assignment” (PA) mechanisms and show that they are max-min optimal under two assumptions on the agent’s utility set: \mathcal{U}_A satisfies *richness* and *PA-alignment*. Later sections apply these results to derive simple implementations of PA mechanisms as max-min optimal in economically relevant environments.

3.1 Measure and Probability Assignment

Given a mechanism D , a strategy σ , and a state realization θ , the *induced measure* $\mu_{\sigma, \theta}^D$ is a measure on the set of actions defined by

$$\mu_{\sigma, \theta}^D(B) = \sum_{i=1}^N \text{Prob}[a_i \in B | \sigma, \theta] \quad \text{for each measurable } B \subseteq \mathcal{A}.$$

where, as a matter of notation, $\mu(B)$ denotes the measure placed by μ on the set $B \subseteq \mathcal{A}$.)

If D is a deterministic mechanism, then the induced measure can be thought of as a list telling us the number of times that each action will be played over the course of the game. For stochastic mechanisms, it tells us how many times an action or set of actions will be played in expectation.

Say that a measure on \mathcal{A} is *proper* if it places a mass of N on the full set. In total N actions are taken, so any induced measure $\mu_{\sigma, \theta}^D$ is proper.

A probability assignment (PA) mechanism specifies some proper measure μ , then asks the agent to declare probability distributions from which each action is to be drawn. Any action distributions are allowed so long as the induced measure over all actions – the sum of the distributions – is μ . PA mechanisms can be thought of as stochastic generalizations of quotas. The agent can choose actions as she pleases, as long as each action is played the correct number of times by the end of the game.

Definition (Probability Assignment). A *probability assignment* mechanism $\text{PA}(\mu)$ is a mechanism characterized by a proper measure μ on the set of actions.

There is no time 0 message. The interim message space \mathcal{M} is the set of n -tuples of distributions over actions $(m_1, \dots, m_n) \in \Delta(\mathcal{A})^N$ for which $\sum_i m_i = \mu$. Given message (m_1, \dots, m_n) , action a_i is drawn according to distribution m_i .

The notation $\sum_i m_i$ refers to the measure defined by $(\sum_i m_i)(B) = \sum_i m_i(B)$ for any set B .

By a standard compactness argument, an agent always has some optimal

strategy in a probability assignment mechanism.⁶

3.2 Richness

Richnes is a way of formalizing the notion that the agent may have certain kinds of extreme preferences.

Definition (Richness). Say that a set of agent utilities $\mathcal{U}_A \subseteq \mathcal{U}$ is *rich* if there exists an infinite subset of the natural numbers $\mathcal{N} \subseteq \mathbb{N}$; a function $\psi : \mathcal{A} \times \Theta \times \mathbb{R}_+ \times \mathcal{N} \rightarrow \mathbb{R}$, written as $\psi(a|\theta; \lambda, n)$; and a pair of sign constants $s, t \in \{-1, 1\}$, such that for all $n \in \mathcal{N}$:

- For each $\lambda' > 0$, there exists a function U_A in \mathcal{U}_A with $U_A(a|\theta)$ equivalent to $\psi(a|\theta; \lambda, n) + s \cdot (a + t \cdot \lambda)^{2n}$ for some $\lambda > \lambda'$.
- ψ is of order λ^n : there exists $C > 0$ such that $|\psi(a|\theta; \lambda, n)| \leq C\lambda^n$ for each a and θ , for λ large enough.

Fixing n , as λ goes to infinity the agent with utility $U_A(a|\theta) = \psi(a|\theta; \lambda, n) + s \cdot (a + t\lambda)^{2n}$ cares only about taking low average actions (if $s \cdot t = 1$) or high ones ($s \cdot t = -1$). Then increasing n makes the preferences more concave ($s = -1$) or convex ($s = 1$) in actions, implying extreme preferences over ever higher moments of the induced measure (mean, variance, etc).

If the agent has rich preferences, then for any mechanism the principal proposes there will be extreme agent types which, in the limit, will not condition the number of times an action is played on the state realization: the induced measure will be constant across states.

Lemma 1. Take any mechanism D and any rich set of agent utilities \mathcal{U}_A . There exists some proper measure μ_∞^D and some sequence of types $\langle U_A^j \rangle_{j=1}^\infty$ in \mathcal{U}_A such that for all $\theta \in \Theta^N$ and for all corresponding optimal strategies

⁶Formally, this follows because the set of possible assignments satisfying $\sum_i \mu_i$ is compact with respect to the component-by-component weak convergence of measures, and payoffs are continuous with respect to the same. See footnote 7 for a definition of weak convergence.

$\langle \sigma^j \rangle_{j=1}^\infty$ with $\sigma^j \in \Sigma^{*D}(U_A^j)$, it holds that $\mu_{\sigma^j, \theta}^D$ weakly converges⁷ to μ_∞^D as $j \rightarrow \infty$.

Applied to a probability assignment mechanism, Lemma 1 is trivially true; *any* reports by an agent result in the same induced measure over actions. But in a mechanism which gives more flexibility, Lemma 1 implies that there is an agent type which (for instance) sends messages so as to maximize the average action; conditional on maximizing the average, minimizes the variance; etc. This limiting agent type agent would only use states of the world to decide between different messages which induce the same measure.

Proof Outline of Lemma 1. For formalization, see Appendix B.

Rich utility sets include utility functions with a term ψ of order λ^n which may depend on θ , plus a polynomial $s \cdot (a + t \cdot \lambda)^{2n}$. Expanding out the polynomial gives a constant $s\lambda^{2n}$; constants times $a\lambda^{2n-1}$, $a^2\lambda^{2n-2}$, ..., and $a^{n-1}\lambda^{n+1}$; and terms of order λ^n or lower. For λ large, the agent approximately maximizes lexicographically, looking first at terms in her utility function with a higher order in λ . The highest order terms are state-independent: she maximizes st times the first moment $\sum_i a_i$, then st^2 times the second moment $\sum_i (a_i)^2$, and so forth through the $n - 1^{st}$ moment. Only then does the agent consider the states θ . As we take n and λ to infinity, the agent plays a strategy in which all moments of the induced measure are independent of the realized states. Any measures with identical moments are equal. ■

The definition of richness above provides a sufficient condition on utility sets to imply Lemma 1. But it is by no means a necessary condition. One could find alternate definitions of “richness” which could take the place of this in Lemma 1 (and therefore in Theorem 1, below), and which neither imply nor are implied by this definition. Indeed, I have not defined richness in this manner out of any consideration that the required subsets are economically

⁷ The cumulative mass function of μ at action a , written $\mu((-\infty, a])$, is the measure placed by μ on the set of actions less than or equal to a . A sequence of measures $\langle \mu^j \rangle_{j=1}^\infty$ is said to *weakly converge* to a limiting measure μ if, at all continuity points a of the cumulative mass function of μ , it holds that $\mu^j((-\infty, a]) \xrightarrow{j \rightarrow \infty} \mu((-\infty, a])$.

meaningful. Rather, these subsets are *contained in* sets which I do consider economically meaningful, discussed in later sections.

3.3 Aligned Delegation

I say that a mechanism is aligned delegation if it leads all agent types in \mathcal{U}_A to act exactly in the principal’s best interest.

Definition (Aligned Delegation). Fix a principal’s utility U_P . A mechanism D is *aligned delegation* over \mathcal{U}_A if there exists some “aligned strategy” σ^* which is optimal (i.e., sequentially rational) for every type $U_A \in \mathcal{U}_A$ and (if U_P is not already contained in \mathcal{U}_A) would also be optimal for an agent of type $U_A = U_P$:

$$\exists \sigma^* \text{ s.t. } \sigma^* \in \Sigma^{*D}(U_A) \text{ for all } U_A \in \mathcal{U}_A \cup \{U_P\}.$$

Any time I talk about “the payoff” to the principal of an aligned delegation mechanism, it should be understood to mean the payoff under an aligned strategy: $\mathbb{E}_{\theta, \mathbf{a}} [\sum_i U_P(a_i | \theta_i) | \sigma^*]$. If there are multiple aligned strategies, they will all be payoff equivalent.

Observation (Aligned Delegation for Direct Mechanisms). In a direct mechanism, the agent reports her utility function and then reports the states that she observes. Loosely speaking, an incentive compatible direct mechanism is aligned delegation if and only if the actions taken are independent of the reported utility function. This is precise when $U_P \in \mathcal{U}_A$; otherwise, it holds for an augmented direct mechanism which would also be incentive compatible for an agent of type $U_A = U_P$.

Definition (PA-alignment). Preferences are said to be *PA-aligned* if for all proper measures μ , the probability assignment mechanism $PA(\mu)$ is aligned delegation.

In this paper I treat PA-alignment as a property of preferences, or as a property of the agent’s utility set \mathcal{U}_A , given U_P . I seek conditions on preferences which guarantee the property independently of the other parameters of the

problem. It is trivially the case that the agent's utility is PA-aligned if $\mathcal{U}_A = \{U_P\}$. In later sections I give more interesting families of U_P and \mathcal{U}_A which imply PA-alignment.

In an aligned delegation mechanism, all agent types in \mathcal{U}_A give the same expected payoff to the principal. So we can talk about the payoff to the principal under a given measure without assuming a prior over \mathcal{U}_A .

Lemma 2. When utilities are PA-aligned, there is some optimal measure μ^* which maximizes the principal's expected payoff from $\text{PA}(\mu)$ over proper measures μ .

Call the probability assignment mechanism characterized by the optimal measure the *optimal probability assignment mechanism*. The optimal measure depends on the principal's utility function U_P and on the distribution of states, but not on the agent's utility set \mathcal{U}_A (so long as utility is PA-aligned).

3.4 The Main Result

I can now state and prove the main theorem of the paper.

Theorem 1. Suppose the set of agent utilities is rich and PA-aligned. Then the optimal probability assignment mechanism is max-min optimal.

Proof. Fix an arbitrary mechanism D . I seek to show that the optimal probability assignment mechanism is weakly preferred to D under the max-min criterion.

By Lemma 1, there is some sequence of types $\langle U_A^j \rangle$ and some measure μ_∞^D such that, under every corresponding sequence of optimal strategies $\langle \sigma^j \rangle$ and every vector of states θ , it holds that $\mu_{\sigma^j, \theta}^D$ weakly converges to μ_∞^D .

For any θ , the agent of type U_A^j can replicate the outcome from D under strategy σ^j in a probability assignment mechanism $\text{PA}(\mu_{\sigma^j, \theta}^D)$. So the agent's payoff is weakly higher under $\text{PA}(\mu_{\sigma^j, \theta}^D)$ than under D , conditional on realized states θ . The principal's payoff is also higher for this agent type; by PA-alignment, the agent's choices in $\text{PA}(\mu_{\sigma^j, \theta}^D)$ maximize the principal's payoff.

By continuity of probability assignment payoffs with respect to the measure (see Claim 2 in the proof of Lemma 2, Appendix B), the payoff to the principal from $\text{PA}(\mu_{\sigma^j, \theta}^D)$ approaches that from $\text{PA}(\mu_\infty^D)$ for each θ , and therefore in expectation over θ . So as j goes to infinity, the principal’s expected payoff from D given type U_A^j converges to a value bounded above by the principal’s type-independent payoff from $\text{PA}(\mu_\infty^D)$. In other words, $\text{PA}(\mu_\infty^D)$ weakly max-min dominates D . And the principal’s expected payoff from $\text{PA}(\mu_\infty^D)$ is weakly below that of the optimal probability assignment mechanism, completing the proof. ■

3.5 Discussion

In this section I introduced the concepts of richness and PA-alignment. Rich agent utility sets are “large enough” to include certain extreme functions, while PA-aligned sets are “small enough” that there are no agent types whose play would differ from that of the principal in a probability assignment mechanism.

For a proposed utility set \mathcal{U}_A , the richness condition can be checked directly by confirming that \mathcal{U}_A contains an appropriate subset. The assumption of PA-alignment cannot be verified mechanically, however. It needs to hold for all possible proper measures.

In Section 4, Section 5 (for a sequential problem), and Section 7.1, I provide conditions on utility functions which imply PA-alignment. Furthermore, I show that the max-min optimal probability assignment mechanisms can often be implemented through simple deterministic mechanisms. I discuss relevant examples of rich utility sets as I proceed.

In Section 6 I show that the conclusion of Theorem 1 may fail to hold when the agent’s utility set is not rich – probability assignment mechanisms may no longer be max-min optimal. Other aligned delegation mechanisms which give the agent more freedom may do better.

4 Increasing Differences and Ranking

4.1 PA-Alignment Under Increasing Differences

Definition (Increasing Differences). Say that a utility function $U \in \mathcal{U}$ satisfies *increasing differences* if for all $a^2 > a^1$ in \mathcal{A} and all $\theta^2 > \theta^1$ in Θ ,

$$U(a^2|\theta^2) - U(a^1|\theta^2) \geq U(a^2|\theta^1) - U(a^1|\theta^1)$$

This is a standard condition which implies that a player’s preferred action is increasing in the state of the world – see Topkis (1998), for instance, for applications of increasing differences and similar supermodularity or complementarity conditions to economics. Chakraborty and Harbaugh (2007) consider preferences of this form in a simultaneous cheap talk game over many decisions.

Increasing differences is easy to check when the function U is twice differentiable. In that case, U has increasing differences if and only if it has a nonnegative cross partial derivative: $\frac{\partial^2 U}{\partial a \partial \theta} \geq 0$.

Under increasing differences, I will show that the optimal strategy in a probability assignment mechanism is to assign actions “assortatively” – lower actions get assigned to lower states.

Definition (Assortative Assignments). Given a probability assignment mechanism $\text{PA}(\mu)$ and a state realization θ , say that an assignment (m_1, \dots, m_N) is *assortative* if $\theta_i < \theta_j$ implies that $\max[\text{Supp } m_i] \leq \min[\text{Supp } m_j]$ – their supports are fully ordered.

If all θ ’s are distinct, then there is a unique assortative assignment. Otherwise there may be multiple assortative assignments, all of which are payoff equivalent to both players.

Lemma 3.1 establishes that assortative assignments are optimal in a probability assignment mechanism. Suppose we have a nonassortative assignment. Then there is (some probability mass of) a low action a^1 played in a high state θ^2 and a high action a^2 played in a low state θ^1 . He can weakly improve

this by switching the two actions, increasing payoffs by $U(a^1|\theta^1) + U(a^2|\theta^2) - U(a^1|\theta^2) - U(a^2|\theta^1) -$ greater than or equal to 0 by increasing differences.

If both the principal and agent (i.e., every agent type) have increasing difference utility, then they agree on this optimal strategy: Lemma 3.2 points out that preferences are PA-aligned.

Lemma 3.

1. Fix a probability assignment mechanism and a vector of states. If the agent has increasing difference utility, any assortative assignment is optimal.
2. If the principal and agent have increasing-difference utility, then preferences are PA-aligned.

Lemma 4 (Sufficient conditions for richness). The agent has rich preferences if \mathcal{U}_A is the set of all increasing-difference functions in \mathcal{U} ; or if \mathcal{U}_A contains the increasing-difference functions which are concave in a for any θ , or those which are convex in a .

Suppose the principal has increasing difference utility, and the agent has a rich set of increasing difference utilities. Then Theorem 1 combined with Lemma 3 imply that the optimal probability assignment mechanism is max-min optimal.

4.2 Implementation via Ranking Mechanisms

In Lemma 3 I found that assortative probability assignments were optimal, under increasing difference utility. So the agent doesn't need to report full action distributions m_1, \dots, m_N . She can just report the relative rankings of the N states, and the mechanism can assign the probability mass assortatively on its own. (See Figure 1.)

If the measure defining a probability assignment mechanism is a sum of N unit atoms, then an assortative assignment yields deterministic actions. The lowest state is assigned with probability one to some low action, the next

lowest state is assigned with probability one to some weakly higher action, etc. In other words, we have a ranking mechanism.

Definition (Ranking Mechanism). A *ranking mechanism* is characterized by a list of actions $b^{(1)} \leq b^{(2)} \leq \dots \leq b^{(N)}$ in \mathcal{A} .

There is no time 0 message. The interim message space \mathcal{M} is the set of permutations on $\{1, \dots, N\}$ – bijections from the set into itself. For some reported permutation π , $\pi(i) < \pi(j)$ is interpreted as a report that $\theta_i \leq \theta_j$. Given report π , action a_i at decision i is chosen as $b^{(\pi(i))}$.

For any increasing difference utility type, it is optimal to report honest rankings – ranking mechanisms are aligned delegation. Under the same increasing-difference assumption, Chakraborty and Harbaugh (2007) investigated ranking as a cheap talk protocol rather than a delegation mechanism. If there are many states drawn iid from a known distribution, then knowing the ranking of a state almost reveals its value. So if there are many iid decisions then the principal can take an approximately first-best action at each one.

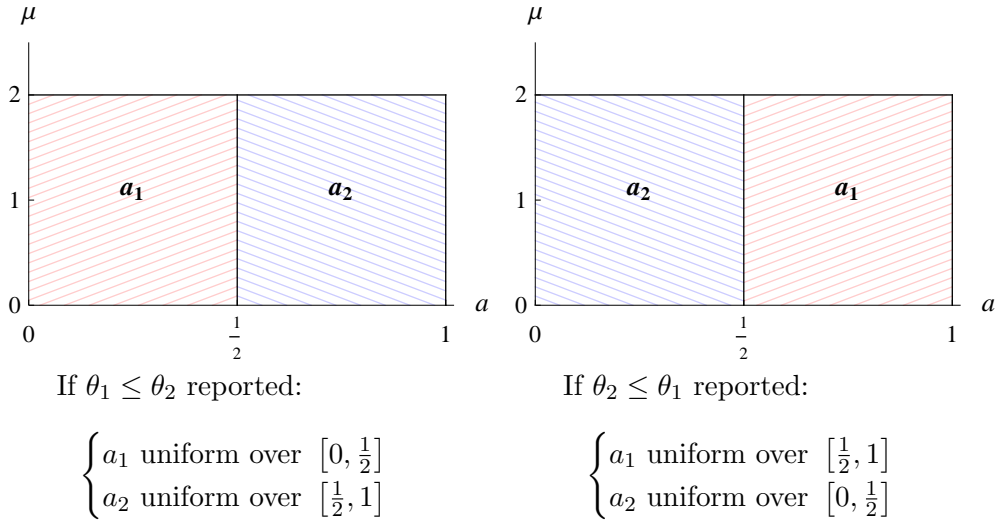
Proposition 1 gives an argument for using ranking mechanisms even when we may not have many iid decisions. Ranking mechanisms are max-min optimal, for any number of decisions and any joint distribution over states, under appropriate assumptions on players’ utility. This follows because the optimal probability assignment mechanism has a measure which is a sum of N unit atoms, so this max-min optimal mechanism can be implemented by ranking. (See Figure 2.)⁸

Proposition 1. If the principal and agent have increasing-difference utility, then the optimal probability assignment mechanism can be implemented by a

⁸Another instance of ranking in economics is in tournament incentive structures. The literature on tournaments (see e.g. Lazear and Rosen (1981)) focuses on the incentives provided to those being evaluated – motivating employees to work hard. I take the qualities of the evaluated to be exogenous, and look at the incentives of the evaluator – the supervisor who observes her employees and pays out bonuses. Malcomson (1984) and Fuchs (2007) have previously pointed out that tournaments may have good incentive properties for firms which prefer ex post to pay employees low bonuses. A firm which gives out bonuses based only employees’ rankings pays the same in aggregate for any report, so can plausibly commit to rank honestly. If an employee’s bonus were instead based on the firm’s nonverifiable report of her individual output, the firm could save money by falsely reporting low outputs.

Figure 1: To implement a PA mechanism under increasing difference preferences, the agent need only report the rankings of states.

Example: Implementing $PA(\mu)$ for μ uniform on $[0,1]$, with $N = 2$.



ranking mechanism. This is max-min optimal if the agent's utility set is rich.


The PA mechanism can be implemented by ranking in the sense that there is a ranking mechanism in which each action a_i is drawn from an identical distribution in the two mechanisms, given optimal play (assortative assignments, honest ranking) by the agent.

Example 1. Let the principal have the increasing-difference utility function $U_P(a|\theta) = -(a - \theta)^2$ and let the action set \mathcal{A} contain Θ . So the principal's preferred action given state θ is $a = \theta$.

Let the agent's utility be taken from a rich subset of the increasing difference functions. So a ranking mechanism assigning the j^{th} lowest reported state to some action $b^{(j)}$ is a max-min optimal mechanism. For this U_P function, the optimal choice of $b^{(j)}$ is the ex ante expected value of the j^{th} lowest state.

If each θ_i is iid uniform over $\Theta = [0, 1]$, the optimal ranking mechanism assigns the j^{th} lowest action to $b^{(j)} = \frac{j}{N+1}$. This mechanism gives the principal

Figure 2: Ranking implements the optimal PA mechanism. ($N = 2$ case).



Take a PA mechanism with any measure μ . Here we have μ uniform on $[0, 1]$. The low state is assigned an action drawn from $[0, \frac{1}{2}]$ and the high state action is drawn from $[\frac{1}{2}, 1]$.

The principal prefers to “consolidate” the leftmost unit of mass into some single point in $[0, \frac{1}{2}]$, and the rightmost unit of mass into a point in $[\frac{1}{2}, 1]$. This induces a (deterministic) ranking mechanism.

a payoff of $-\frac{1}{6(N+1)}$ per period, compared to $-\frac{1}{12}$ from no delegation (taking an uninformed principal’s preferred action of $a_i = \frac{1}{2}$ each period) and 0 from first-best ($a_i = \theta_i$ each period).⁹ Interpreting these numbers, the ranking mechanism gives the principal $\frac{N-1}{N+1}$ of the possible surplus from delegation – his payoff is that proportion of the way from no delegation to first-best.

Confirming the Chakraborty and Harbaugh (2007) result, the principal’s surplus goes to 100% as N gets large.

⁹The j^{th} order statistic (ie, j^{th} lowest number) of N uniformly distributed variables is distributed according to a Beta($j, N + 1 - j$) distribution. This has mean $\frac{j}{N+1}$ and variance $\frac{j(N+1-j)}{(N+1)^2(N+2)}$. The principal’s expected lifetime payoff is minus sum of the variances, which can be calculated to be $-\frac{N}{6(N+1)}$.

5 Sequential Decisions

In this section, I change the timing of the game so that the agent observes states one at a time. The agent sees state θ_{i+1} and takes action a_{i+1} only after action a_i has been taken. A manager chooses today's investment level before learning the profitability of future projects, say. In this sequential environment (as opposed to the earlier simultaneous one) ranking mechanisms are no longer feasible. Final rankings are not known until all but the last action have already been played.

I assume that the principal and agent share a common prior over the distribution of θ in the sequential environment.¹⁰

In a sequential mechanism the agent sends a time-0 message m_0 before the agent observing any states, and in periods $i = 1, \dots, N$ sends a message m_i after the agent observing θ_i . The agent also knows the period- i history (reports and actions from periods 1 through $i - 1$) as well as past states when she reports m_i . After the message m_i is sent, action a_i is drawn from a distribution which depends on the message and the history.

This mechanism form is without loss of generality so long as the principal gets no information about the realizations of past states, or the distributions of future states, over the course of the game. For instance, this rules out a principal's observing his utility realization from decision i and using this to alter the terms offered to the agent at decision $i + 1$.

Probability assignment mechanisms extend straightforwardly to the sequential environment, with the agent choosing action distributions one at a time instead of all at once. The principal specifies a proper measure μ , and the agent's period i message is a distribution $m_i \in \Delta(\mathcal{A})$ from which action a_i is drawn. By the end of the game, the distributions must sum to the measure: $\sum_{i=1}^N m_i = \mu$. This gives a constraint that m_i is less than or equal to the "remaining measure" $\mu - \sum_{j=1}^{i-1} m_j$.¹¹ At the last period, the agent has no

¹⁰What is important is that the agent is at least as well informed as the principal: if they were to share information, the parties would converge on the agent's beliefs. So if the principal and agent had the same utility function, the principal would defer to the agent's choices.

¹¹Under the partial order on measures over \mathcal{A} , measure μ' is said to be greater than or

choice over the action; a_N must be drawn according to $m_N = \mu - \sum_{j=1}^{N-1} m_j$.

With these definitions, Lemma 1 and Theorem 1 go through unchanged in the sequential environment: if the agent has a rich and PA-aligned set of utilities, then the optimal probability assignment mechanism is max-min optimal.¹²

The richness of a utility set isn't affected by the timing of the game. But the PA-aligned sets are different. In the example below, I show that two increasing-difference utility types which play identically in a simultaneous probability alignment mechanism may play differently in a sequential one. Increasing differences no longer implies aligned delegation. In general, it is harder to satisfy PA-alignment in the sequential environment.¹³

Example 2. Let $\Theta = \{0, \frac{1}{2}, 1\}$, with θ_i drawn uniformly from Θ in each period. Let $\mathcal{A} = \{0, 1\}$. Say that the principal and agent are known to have the following utility functions:

		θ		
	U_P	0	$\frac{1}{2}$	1
a	0	10	9	0
	1	0	0	10

		θ		
	U_A	0	$\frac{1}{2}$	1
a	0	10	0	0
	1	0	9	10

Both functions satisfy increasing differences. So in a simultaneous problem, the players agree on an assortative assignment in any probability assignment mechanism. Preferences are PA-aligned.

In a sequential problem, the players no longer agree on a strategy. Let there

equal to measure μ'' if $\mu'(B) \geq \mu''(B)$ for every measurable $B \subseteq \mathcal{A}$. If $\mu' \geq \mu''$, then we can subtract μ'' from μ' : $\mu' - \mu''$ is defined by $(\mu^1 - \mu^2)(B) = \mu^1(B) - \mu^2(B)$.

¹²As in a simultaneous problem, the agent has an optimal strategy under any PA mechanism. And if utilities are PA-aligned, then there is an optimal measure for the principal. In otherwise identical simultaneous and sequential problems, the optimal measures will not necessarily be the same.

¹³That is, suppose we are looking for a condition on utilities which guarantees PA-alignment for any number of decisions and any joint distribution over states. Any such condition guaranteeing sequential PA-alignment also guarantees simultaneous PA-alignment. (Beliefs in a sequential problem can replicate any realized states in a simultaneous problem). The reverse is not true: as the example shows, increasing differences ensures PA-alignment in simultaneous but not sequential problems.

be two periods, and take the probability assignment mechanism specifying that actions $a = 0$ and $a = 1$ must each be played once. If $\theta_1 = 0$, both types would prefer $a_1 = 0$ at period 1 (and $a_2 = 1$). If $\theta_1 = 1$, both types would prefer $a_1 = 1$. But if $\theta_1 = \frac{1}{2}$, type U_P would prefer $a_1 = 0$ while type U_A prefers $a_1 = 1$. We no longer have aligned delegation.

Intuitively, the disagreement is due to the fact that the principal treats the random draw of θ_2 as greater than an observed state $\theta_1 = \frac{1}{2}$, while the agent treats it as less than $\frac{1}{2}$.

To guarantee PA-alignment in a sequential problem, we need a stronger condition than increasing differences. Players must agree on how to rank not just states, but distributions of states. With *quadratic loss* preferences, players rank distributions of states by their expectations.

Definition (Quadratic Loss Preferences). Say that a utility function $U \in \mathcal{U}$ is *quadratic loss* if there exists a weakly increasing continuous function $c : \mathcal{A} \rightarrow \mathbb{R}$ such that $U(a|\theta)$ is equivalent to $-(c(a) - \theta)^2$.

Under quadratic loss preferences, in state θ a player wants an action a which takes $c(a) - \theta$ as close as possible to 0.¹⁴ The optimal action is weakly increasing in the state. And by choosing $c(\cdot)$ appropriately, we can model preferences with any increasing optimal action function.¹⁵ There is enough flexibility in this utility class that the set of all quadratic loss functions is rich.

Lemma 5. The set of quadratic loss functions is a rich subset of the increasing difference functions.

The quadratic loss functions include the most common functional forms in the delegation and cheap talk literature, quadratic loss preferences with a

¹⁴Notice 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 world from better informed agents.

¹⁵The optimal action is strictly increasing as a function of the state anywhere that it maps to the interior of \mathcal{A} , and weakly increasing otherwise. Any such function can be the optimal action function for some quadratic loss utility.

constant or a linear bias. The constant bias preferences have $c(a) = a - \lambda$, corresponding to a utility of $U(a|\theta) = -(a - \theta - \lambda)^2$ and ideal action $a = \theta + \lambda$. The linear bias preferences have $c(a) = \frac{a}{\lambda^{(1)}} - \frac{\lambda^{(0)}}{\lambda^{(1)}}$, giving a utility equivalent to $U(a|\theta) = -(a - \lambda^{(1)}\theta - \lambda^{(0)})^2$ and ideal action of $a = \lambda^{(1)}\theta + \lambda^{(0)}$. I discuss these two functional forms in greater detail in Section 6.

In a sequential probability assignment mechanism, an agent with quadratic loss preferences will play what I call a “sequential-assortative” strategy. Given a current state θ_i , she will assign a_i to some action, or set of actions, from the remaining measure. Remaining actions below the ones chosen are expected to be played in future periods with states less than θ_i , and actions above those chosen are expected to be played in periods with states higher than θ_i .¹⁶ A principal with quadratic loss preferences agrees that this play is optimal, and hence the mechanisms satisfy aligned delegation.

Lemma 6. Let the principal and agent have quadratic loss preferences in a sequential problem. Then preferences are PA-aligned.

Combining Lemma 6 and Theorem 1, suppose that the principal and agent have quadratic loss utilities, and the agent’s utility set is rich. Then the optimal probability assignment mechanism is max-min optimal.

As in the simultaneous problem, it can be shown that the optimal probability assignment has a measure composed of N atoms, and that when the agent plays a sequential-assortative strategy (defined formally in the proof of Lemma 6) actions will be deterministic. This optimal PA mechanism is therefore equivalent to a simpler quota mechanism in which the agent chooses actions directly rather than reporting distributions.

Definition (Sequential Quota). A *sequential quota* mechanism is characterized by a list of actions $b^{(1)} \leq b^{(2)} \leq \dots \leq b^{(N)}$ in \mathcal{A} .

There is no time 0 message. In each period $i \geq 1$, the agent chooses an action $b^{(j)}$ from the list, without replacement. Then the action $a_i = b^{(j)}$ is

¹⁶More precisely, take any action in the support of the remaining measure at $i + 1$. Given the agent’s strategy and the beliefs over future state realizations, we can find the distribution over states in which this action is played. I say that the action is expected to be played in a higher state if the expectation of this distribution is greater than the current state θ_i .

taken.

Proposition 2. If the principal and agent have quadratic loss preferences in a sequential problem, then the optimal probability assignment mechanism can be implemented as a sequential quota. This is max-min optimal if the agent's utility set is rich.

Example 3. Let the principal have the quadratic loss utility function $U_P(a|\theta) = -(a-\theta)^2$ and let the agent's utility be taken from a rich subset of the quadratic loss functions. As in Example 1, let states in each period be uniform over $\Theta = [0, 1]$, and let the action set \mathcal{A} contain $[0, 1]$. The principal's preferred action in state θ is $a = \theta$.

A sequential quota with some action list $b^{(1)} \leq b^{(2)} \leq \dots \leq b^{(N)}$ is a max-min optimal mechanism. For this U_P function, the optimal choice of $b^{(j)}$ is the expected value of the state in which the j^{th} lowest action will be played, given a sequential-assortative strategy. We can solve for this optimal $b^{(j)}$ by backwards induction.

If $N = 1$, the optimal action is $b^{(1)} = \frac{1}{2}$, the expected state.

If $N = 2$, the action which is not played in period 1 will be played in period 2 at an expected state of $\frac{1}{2}$. So if $\theta_1 < \frac{1}{2}$, the agent chooses a_1 as the low action $b^{(1)}$; if $\theta_1 > \frac{1}{2}$, the agent chooses $b^{(2)}$. This means that the low action is played with probability one half at the first period, at an average state of $\frac{1}{4}$; and probability one half at the second period, at an average state of $\frac{1}{2}$. The average state at which $b^{(1)}$ is played is therefore $\frac{3}{8}$, and likewise the average state at which $b^{(2)}$ is played is $\frac{5}{8}$. So the optimal action list is $b^{(1)} = \frac{3}{8}$, $b^{(2)} = \frac{5}{8}$. The principal gets a surplus of 18.75% of the first-best payoff from this quota. In the corresponding simultaneous problem we had action list $\frac{1}{3}$ and $\frac{2}{3}$, with surplus of 33%.

If $N = 3$, the low action $b^{(1)}$ will be played in the first period if $\theta_1 < \frac{3}{8}$; the medium action $b^{(2)}$ will be played if $\frac{3}{8} < \theta_1 < \frac{5}{8}$; and the high action $b^{(3)}$ will be played if $\theta_1 > \frac{5}{8}$. In period 2, the lower remaining action is played if $\theta_2 < \frac{1}{2}$ and the higher one if $\theta_2 > \frac{1}{2}$. In the final period, the last remaining action is played. This gives us optimal actions (equal to the expected state at which the

action is played) of $b^{(1)} = \frac{39}{128} \simeq .305$, $b^{(2)} = \frac{1}{2}$, and $b^{(3)} = \frac{89}{128} \simeq .695$, implying surplus of 30.5%. In the corresponding simultaneous problem we had action list $\frac{1}{4}$, $\frac{1}{2}$, and $\frac{3}{4}$, with surplus of 50%.

In Appendix A.2 I extend a similar result of Jackson and Sonnenschein (2007) to show that the surplus from the optimal sequential quota goes to 100% as N gets large, if states are iid from a known distribution.

6 Preferences without Richness

When the agent has a rich set of preferences in a PA-aligned class, the principal can do no better (in a max-min sense) than to use a probability assignment mechanism – for example, ranking and quota mechanisms. These PA mechanisms let the agent choose actions as she pleases, subject to fixing the number of times that actions are taken.

With less uncertainty about the agent’s preferences, the principal may be able to improve on probability assignment. He can give the agent more flexibility if he knows that the agent will not use her discretion to harm him. (Rich preferences guarantee that some agent type *will* adversely exploit such flexibility). Indeed, I show that under quadratic loss constant bias preferences it is max-min optimal to use aligned delegation *budget mechanisms*. Budgets fix only the mean of actions rather than the entire distribution. With linear rather than constant biases, *two moment mechanisms* which fix the mean and variance are aligned delegation and max-min optimal.

Throughout this section, I assume that the action space is an interval: $\mathcal{A} = [\underline{a}, \bar{a}]$. This assumption does not affect the qualitative results, but it simplifies the definitions of budget and two-moment mechanisms. For instance, it ensures that there will be no need for randomization in a budget mechanism.

6.1 Quadratic Loss Constant Bias preferences

Suppose the principal and the agent each have quadratic loss constant bias (QLCB) preferences. Normalize the principal’s utility to $U_P(a|\theta) = -(a - \theta)^2$,

and let the agent have utility in a set $\mathcal{U}_A \subseteq \{-(a - \theta - \lambda)^2 | \lambda \in \mathbb{R}\}$. The principal wants to match the action a to the state θ , while the agent prefers $a = \theta + \lambda$ for some λ unknown to the principal.

Expanding out the agent's utility function, $U_A(a|\theta) = -(a - \theta - \lambda)^2 = -(a - \theta)^2 + 2\lambda a - 2\theta\lambda - \lambda^2$. So $U_A(a|\theta)$ is equivalent to $U_P(a|\theta) + 2\lambda a$. The agent's problem, given a mechanism, is to choose a strategy which maximizes

$$\mathbb{E} \left[\sum_i U_P(a_i|\theta_i) \right] + 2\lambda \mathbb{E} \left[\sum_i a_i \right] \quad (1)$$

Conditional on any sum of actions, the principal and agent preferences agree. But the agent prefers a higher sum if λ is positive, and a lower sum if λ is negative. As $\lambda \rightarrow \pm\infty$, the agent cares only about the sum of actions.

Definition (Unboundedness). I say that the quadratic loss constant bias utility set \mathcal{U}_A is *unbounded* if there exist types in \mathcal{U}_A with $|\lambda|$ arbitrarily large.

Unboundedness is to QLCB preferences as richness is to more general preferences. Under rich preferences, there exists a type (in the limit) which will choose the same measure of actions for any realized states: for all n , the n^{th} moment of actions $\sum_i \mathbb{E}(a_i)^n$ does not vary over θ . Under unbounded QLCB preferences, there is a limiting type for which the sum of actions – the first moment – does not vary. Other moments might still be conditioned on the states.

Lemma 7. In either a simultaneous or sequential problem, let the agent have unbounded quadratic loss constant bias utilities. Fix any mechanism D . There exists a value K and a sequence of types $\langle U_A^j \rangle_{j=1}^\infty$ in \mathcal{U}_A such that for all states θ and all corresponding sequences of optimal agent strategies $\langle \sigma^j \rangle$, it holds that $\mathbb{E} \left[\sum_i \gamma_i a_i \middle| \sigma^j, \theta \right] \rightarrow K$ as $j \rightarrow \infty$.

A budget mechanism gives the agent complete freedom, subject to fixing the sum of actions at some level: $\sum_i a_i = K$.¹⁷ Say that K is a *proper* budget

¹⁷With an unconnected action space, the mechanism would have to allow for some ran-

if $K \in [N\underline{a}, N\bar{a}]$, i.e., the level is neither so high nor so low that it is impossible to hit the constraint.

Definition (Budget Mechanism). A *budget mechanism* is a mechanism characterized by a proper budget $K \in \mathbb{R}$. There is no time-0 message. The agent’s interim messages are the actions to be taken.

In a simultaneous environment, the agent chooses $(a_1, \dots, a_N) \in \mathcal{A}^N$ such that $\sum_i a_i = K$.

In a sequential environment, in period i the agent chooses a_i such that it is feasible for $\sum_{j=1}^N a_j$ to equal K :

$$a_i \in \left[K - \left(\sum_{j<i} a_j \right) - (N-i)\bar{a}, K - \left(\sum_{j<i} a_j \right) - (N-i)\underline{a} \right]$$

I call $K - \sum_{j<i} a_j$ the “remaining budget.”

By equation (1), if the sum of actions is fixed by the mechanism, the agent plays as if she has no bias and is maximizing the principal’s payoffs. So under quadratic loss constant bias preferences, budget mechanisms are aligned delegation.

Lemma 8. Let the principal and agent have quadratic loss constant bias preferences. Then in either a simultaneous or sequential problem:

1. Any budget mechanism satisfies aligned delegation.
2. There exists an optimal budget mechanism maximizing the principal’s expected payoff over choice of budget level K .

Proposition 3. Let the principal and agent have quadratic loss constant bias preferences, and let the agent’s utility be unbounded. Then in either a simultaneous or sequential problem, the optimal budget mechanism is max-min optimal.

domized actions. For instance, the agent might want action a_i to contribute x to the expected sum, but $x \notin \mathcal{A}$. So to get $\mathbb{E}[a_i] = x$, the mechanism would randomize between choosing a_i as some action higher than x , and some action lower than x . With interval action spaces, even if the agent were free to choose distributions over actions, she would play deterministically.

Given that these mechanisms satisfy aligned delegation, the proof of Proposition 3 is essentially identical to that of Theorem 1. We just replace Lemma 1 (under richness, some extreme type chooses a predetermined measure of actions) with Lemma 7 (under unboundedness, some extreme type chooses a predetermined sum of actions).

Example 4. Let $U_P(a|\theta) = -(a - \theta)^2$ and let the agent's utility be taken from an unbounded subset of the quadratic loss constant bias utilities. Let each θ_i be iid uniformly distributed over $\Theta = [0, 1]$, and let \mathcal{A} be a “big enough” interval containing $[0, 1]$.¹⁸ The optimal budget mechanism in a simultaneous or sequential problem sets $K = \frac{N}{2}$ – the average action should be $\frac{1}{2}$ – and this mechanism is max-min optimal.

Simultaneous Case: After observing states the agent chooses actions so that $a_i - \theta_i$ is constant over i : $a_i = \frac{1}{2} + \theta_i - \frac{1}{N} \sum_{j=1}^N \theta_j$. The principal's expected per-period payoff is $-\frac{1}{12N}$, which corresponds to $\frac{N-1}{N}$ of the possible surplus. If the principal used the optimal ranking mechanism, he would get only $\frac{N-1}{N+1}$ of the surplus (Example 1).

Sequential Case: At period i with remaining budget $K_i = \frac{N}{2} - \sum_{j < i} a_j$, after the agent observes θ_i she chooses action $a_i = \theta_i + \frac{K_i - \theta_i - (N-i)/2}{N-i+1}$. This action choice sets the difference $a_i - \theta_i$ equal to the expected difference $a_j - \theta_j$ for $j > i$.

If $N = 1$, this replicates the no-delegation outcome of $a_1 = 1/2$.

If $N = 2$, the agent chooses action $a_1 = \frac{1}{4} + \frac{\theta_1}{2}$, then action $a_2 = 1 - a_1$.

This gives the principal an expected payoff of $-\frac{1}{16}$ per period, 25% of

¹⁸The players may want to choose actions outside of $[0, 1]$ if given the chance. For instance, say there is a simultaneous problem with 4 decisions, and states are drawn from $[0, 1]$. The mechanism requires $\sum_i a_i = 2$. If it so happens that $\theta_1 = 1$ and $\theta_2 = \theta_3 = \theta_4 = 0$, then given any possible actions in \mathbb{R} , the principal and agent prefer $a_1 = \frac{5}{4}$ and $a_2 = a_3 = a_4 = \frac{1}{4}$. When I assume that the action set is “big enough,” I mean that the agents' preferred choices from \mathbb{R} are always available; the boundaries of \mathcal{A} do not bind.

For $\Theta = [0, 1]$ and $K = \frac{N}{2}$, the action set is “big enough” in a simultaneous problem if \mathcal{A} contains $[-\frac{1}{2} + \frac{1}{N}, \frac{3}{2} - \frac{1}{N}]$, or approximately $[-\frac{1}{2}, \frac{3}{2}]$ for N large. In a sequential problem, \mathcal{A} must contain $[\frac{1}{2} - \sum_{j=2}^N \frac{1}{j}, \frac{1}{2} + \sum_{j=2}^N \frac{1}{j}]$. See Frankel (2010b) for the derivation of these intervals, as well as details on solving for strategies, optimal budget levels, and payoffs.

the possible surplus. The optimal sequential quota mechanism gives only 18.75% (Example 3).

If $N = 3$, the principal gets an expected payoff of about $-.051$ per period from the optimal budget mechanism, 38.9% of the possible surplus. The optimal sequential quota gives 30.5%.

6.2 Quadratic Loss Linear Bias preferences

In this subsection I consider quadratic loss linear bias (QLLB) preferences. The principal’s utility is taken to be $U_P(a|\theta) = -(a - \theta)^2$, and the agent has utility in the set $\mathcal{U}_A \subseteq \{-(a - \lambda^{(1)}\theta - \lambda^{(0)})^2 | \lambda^{(0)} \in \mathbb{R}, \lambda^{(1)} \in \mathbb{R}_{++}\}$. The principal wants to match the action a to the state θ , while the agent prefers $a = \lambda^{(1)}\theta + \lambda^{(0)}$. The principal is uncertain over the two λ parameters. $\lambda^{(0)}$ shifts the agent’s ideal point uniformly up or down, while $\lambda^{(1)}$ is a “sensitivity” parameter which determines how much the ideal point moves when the state goes up by one unit. Melumad and Shibano (1991) provides an early characterization of optimal delegation sets when utilities are of this form, assuming a single decision ($N = 1$) with common knowledge of the agent’s utility.

Expanding the linear bias utility, the agent maximizes

$$\sum_i U_A(a_i|\theta_i) = 2\lambda^{(1)} \sum_i a_i\theta_i + 2\lambda^{(0)} \sum_i a_i - \sum_i (a_i)^2 - \sum_i (\lambda^{(1)}\theta_i + \lambda^{(0)})^2 \quad (2)$$

The last sum is independent of the agent’s choices, in any mechanism. If it so happened that a mechanism fixed $\sum_i a_i = K^{(1)}$ and $\sum_i (a_i)^2 = K^{(2)}$, then the agent would also have no control over the second and third sums. No matter what utility parameters $\lambda^{(0)}$ and $\lambda^{(1)}$ she had, her problem would be to maximize $\sum_i a_i\theta_i$. The agent would play exactly as if she shared the principal’s utility function, with $\lambda^{(0)} = 0$ and $\lambda^{(1)} = 1$. I will call such a mechanism which fixes these sums – equivalently, the mean and variance of actions – a “two-moment mechanism.”¹⁹

¹⁹Veszteg (2005) proposes a multiplayer social choice mechanism in which each agent

If the utility set is unbounded in an appropriate sense, then the principal does not benefit by giving the agent flexibility over the sums $\sum_i a_i$ and $\sum_i (a_i)^2$. Some extreme type will push one of the sums as high or low as possible before considering the states of the world. The principal does best to choose the sums in advance.

Definition (Unboundedness). Suppose that the agent has quadratic loss linear bias utilities, so her utility is of the form $U_A(a|\theta) = -(a - \lambda^{(1)}\theta - \lambda^{(0)})^2$ for some $\lambda^{(0)} \in \mathbb{R}$ and $\lambda^{(1)} \in \mathbb{R}_{++}$. I say that the agent’s utility set is *unbounded* if there is a sequence of types in \mathcal{U}_A with $|\lambda^{(0)}| \rightarrow \infty$ while $\lambda^{(1)} \rightarrow 0$.²⁰

As with QLCB preferences, I say that first and second moments $(K^{(1)}, K^{(2)})$ are *proper* if there exist distributions m_1, \dots, m_N such that $K^{(1)} = \sum_i \mathbb{E}_{a_i \sim m_i} a_i$ and $K^{(2)} = \sum_i \mathbb{E}_{a_i \sim m_i} (a_i)^2$. With $\mathcal{A} = [\underline{a}, \bar{a}]$, this corresponds to $K^{(1)} \in [N\underline{a}, N\bar{a}]$ and

$$K^{(2)} \in \left[(K^{(1)})^2, (K^{(1)})^2 + \frac{(N\bar{a} - K^{(1)})(K^{(1)} - N\underline{a})}{N} \right]$$

For a sequential problem, I say that moments are *proper at period j* if, given m_1 through m_{j-1} , there exist distributions m_j through m_N which bring the sums to the appropriate levels; replace N by $N - j + 1$ in the expressions above. At period $N + 1$, the only proper moments are $(0,0)$.

Definition (Two-Moment Mechanism). A *two-moment mechanism* is a mechanism characterized by proper moments $(K^{(1)}, K^{(2)}) \in \mathbb{R}^2$. There is no time-0

reports her values for each decision, with a similar restriction on the first two moments of the reports. It replicates the asymptotic efficiency properties of the mechanism in Jackson and Sonnenschein (2007), but can be implemented with less information about the ex ante distribution of values.

²⁰The unboundedness condition for the quadratic loss constant bias case was essentially “tight” – we needed the magnitude of the bias to go to infinity to get the max-min result. On the other hand, the richness condition for the general case was merely sufficient – other similar conditions could have guaranteed the same results. This unboundedness condition is in the latter category rather than the former. It will give a single sequence of types for which the agent lexicographically cares about the first moment and then the second, in the limit. Another condition could give a sequence for which the agent cared about the second and then the first, for instance.

message. The agent’s interim messages are the distributions over actions to be taken; action a_i is drawn from a reported distribution m_i .

In a simultaneous environment, the agent reports $(m_1, \dots, m_N) \in \Delta(\mathcal{A})^N$ subject to $\sum_{j=i}^N \mathbb{E}_{a_j \sim m_j} a_j = K^{(1)}$ and $\sum_{j=i}^N \mathbb{E}_{a_j \sim m_j} a_j^2 = K^{(2)}$.

In a sequential environment, in period i the agent reports distribution m_i such that the pair of moments $\left(K^{(1)} - \sum_{j \leq i} \mathbb{E}_{a_j \sim m_j} a_j, K^{(2)} - \sum_{j \leq i} \mathbb{E}_{a_j \sim m_j} a_j^2 \right)$ is proper at period $i + 1$.

The agent may choose not to play deterministically. But the agent “tries” to play deterministically, if such play is possible and if the action set is “big enough”.²¹

Lemma 9. Let the principal and agent have quadratic loss constant bias preferences. Then in either a simultaneous or sequential problem:

1. Any two-moment mechanism satisfies aligned delegation.
2. There exists an optimal two-moment mechanism maximizing the principal’s expected payoff over choice of proper moments $(K^{(1)}, K^{(2)})$.

Proposition 4. Let the principal and agent have quadratic loss linear bias preferences, and let the agent’s utility be unbounded. Then in either a simultaneous or sequential problem, the optimal two-moment mechanism is max-min optimal.

²¹For any proper moments and for any distribution of states, we could suppose that the action set were equal to \mathbb{R} and solve for the optimal strategy. Actions would be deterministic so long as $N \geq 2$. The convex hull of the union of all chosen actions across all state realizations would define some compact interval of \mathbb{R} . If the true action set \mathcal{A} is “big enough” that it contains this compact interval, then the agent will choose these same deterministic actions.

If the action set is an interval but it is “too small,” then examples can be found where the agent chooses stochastic actions even though it is possible to satisfy the constraints with deterministic actions.

7 Discussion and Extensions

7.1 Altruism with Private Costs

All utility functions considered so far have satisfied increasing differences. But Theorem 1 can be applied to other classes of utilities, so long some other condition on preferences guarantees PA-alignment. For instance, utilities are PA-aligned if the agent has what I call “altruistic preferences with private costs”: she maximizes a weighted average of the principal’s payoff (altruism) and some function of actions which do not depend on the states of the world (private costs). This includes the case where the agent’s payoffs are state-independent, and only depend on actions.

Given $U_P \in \mathcal{U}$, the agent has altruistic preferences with private costs if $\mathcal{U}_A \subseteq \{U_A | U_A(a|\theta) \text{ equivalent to } \zeta U_P(a|\theta) + c(a) \text{ for some } \zeta \in \{0, 1\}, c : \mathcal{A} \rightarrow \mathbb{R}\}$. An agent with $\zeta = 1$ places a positive weight on the principal’s utility, while $\zeta = 0$ corresponds to state-independent payoffs. The quadratic loss constant bias preferences are of this form with $U_P(a|\theta) = -(a - \theta)^2$, $\zeta = 1$, and $c(a) = 2\lambda a$ for some $\lambda \in \mathbb{R}$.

This class of preferences guarantees PA-alignment in a simultaneous or a sequential problem – the agent maximizes $\zeta \sum_i U_P(a_i|\theta_i) + \sum_i c(a_i)$, and PA mechanisms fix the latter sum in advance. So the agent’s maximization problem is the same as the principal’s. And for a large enough set of possible cost functions $c(\cdot)$, \mathcal{U}_A is rich. We get richness if c may be any continuous function; any increasing (decreasing) one, corresponding to a positively (negatively) biased agent; any convex (concave) increasing function; etc.

If the agent has a rich set of altruistic with private costs utilities, a probability assignment mechanism will be max-min optimal. But it might not be implementable as a deterministic ranking or quota mechanism if U_P is not of the increasing-difference or quadratic-loss forms.

7.2 General Moment Conditions

Above, I give environments where max-min optimal mechanisms fix the expected sum of actions (budgets); the sum of actions and actions-squared (two-moment mechanisms); and the entire distribution of actions, i.e., the sum of every function of actions (probability assignment). In fact, for any desired “moment conditions” – sums of specified functions of actions – we can reverse-engineer an environment in which the max-min optimal mechanism fixes the set of such moments. See Appendix A.1 for details.

7.3 Discounting

Suppose the principal and agent discount certain decisions. For instance, decisions may be less important if they occur later, affect fewer people, or if they determine investments in exogenously smaller projects. We can model this by calling $\gamma_i > 0$ the significance of decision i , and modifying the principal and agent objectives to be

$$\text{Principal : } \sum_i \gamma_i U_P(a_i|\theta_i) \qquad \text{Agent : } \sum_i \gamma_i U_A(a_i|\theta_i)$$

This framework is covered extensively in the extended working paper, Frankel (2010b).

For the quadratic loss with constant or linear biases, almost nothing changes. Budget and two-moment mechanisms are still max-min optimal. The mechanisms just have to be modified so as to fix the *weighted* sums $\sum_i \gamma_i a_i$ and $\sum_i \gamma_i \cdot (a_i)^2$ rather than the unweighted sums.

Looking at general utility functions, Theorem 1 continues to go through: appropriately modified probability assignment mechanisms are still max-min optimal under richness and PA-alignment. The modified mechanisms constrain action distributions m_i to satisfy $\sum_i \gamma_i m_i = \mu$, for μ a measure of size $\sum_i \gamma_i$ – they are analogous to the discounted quotas studied in Frankel (2010a). Moreover, the PA-aligned sets are unchanged. Increasing difference utilities are PA-aligned in a simultaneous problem, quadratic loss utilities in a sequential

problem, and altruistic with private cost utilities are PA-aligned in either case. (Richness is not affected by the change in objectives).

But there is one big difference when we add discounting to probability assignment mechanisms: we can no longer implement the max-min optimal mechanisms as deterministic ranking or quotas. Indeed, nontrivial probability assignment mechanisms now necessarily require randomized actions.²²

For instance, suppose we have two simultaneous decisions under increasing difference preferences and decision 1 is twice as significant: $\gamma_1 = 2\gamma_2$. The max-min optimal PA mechanism can be implemented by a “ranking” mechanism which specifies three actions – call them low, medium, and high – and asks the agent to rank the two states. If she reports that $\theta_1 \leq \theta_2$, then action a_1 is randomized 50/50 between low and medium, while a_2 is taken deterministically at high. If $\theta_1 \geq \theta_2$, action a_1 is randomized 50/50 between medium and high while a_2 is deterministically low. In essence, we align the agent’s incentives by asking her to trade off action 2 against half of the twice-as-important action 1. For N decisions, we generically randomize over $2^N - 1$ actions.

7.4 Asymptotically First-Best Payoffs

All of the max-min optimal mechanisms in this paper give high payoffs when the number of decisions increases. That is, suppose utilities are in a PA-aligned class and states are drawn iid from a known distribution. Consider a sequence of problems in which all other parameters are fixed, but the number of decisions goes to infinity. Then under an appropriate sequence of probability assignment mechanisms, the principal’s expected payoff per decision approaches that from her full information first-best actions. (If we can do better with alternate mechanisms such as budgets, then those also yield asymptotically first-best payoffs). This extends similar results from Chakraborty and Harbaugh (2007) and Jackson and Sonnenschein (2007); see Appendix A.2 for a formalization. In the working paper Frankel (2010b), I also show that we get asymptotically first-best payoffs even with discounting, so long as no single

²²A trivial PA mechanism would require that all actions be taken at a single predetermined point.

decision is important in the limit.

8 Conclusion

This paper looks at delegation contracts which restrict the actions that an agent may choose. In principle, a contract can be any menu of subsets of the N -dimensional action space. Because of the highly dimensional private information and the absence of transfers, solving for Bayesian optimal mechanisms would be infeasible. But by looking for *max-min* optimal mechanisms, we derive simple and intuitive contracts such as ranking mechanisms, quotas, and budgets. The structure of the contracts depends on the preferences of the principal and the agent.

On the technical side, I solve for these max-min optimal contracts by introducing the concept of aligned delegation. This is form of incentive compatibility for indirect contracts which states that all agent types play exactly as if they are maximizing the principal's payoff. It sounds like an excessively strong condition: in a one decision problem, the only aligned delegation mechanism would give the agent no input at all. But with many decisions, it allows for mechanisms like those above which do make effective use of much the agent's private information.

I believe that the study of aligned delegation mechanisms might have further applications. For instance, consider an environment in which a principal wants to elicit information from many experts. The cheap talk and mechanism design literature focuses on the strategic interactions of experts, trying to find ways to “play experts against each other” to induce full revelation of information. See for example Krishna and Morgan (2001), Battaglini (2002), Ambrus and Takahashi (2008), Mylovanov and Zapechelnjuk (2008), or Ambrus and Lu (2010). These constructions tend to be sensitive to the assumptions that collusion or communication is impossible and that each expert's bias is known in advance, and often also require specific informational structures – e.g., each expert is perfectly informed. But suppose that individual experts are poorly informed, and we want to find a way to pool their information without

necessarily achieving full revelation. Aligned delegation mechanisms sidestep strategic issues entirely. Agents will share information fully and honestly, and reach a consensus. In a ranking mechanism, say, everyone wants to get the final ranking right. The principal can use these mechanisms to pool together the agents' information in a robust way.

References

- Ricardo Alonso and Niko Matouschek. Optimal delegation. *Review of Economic Studies*, 75:259–293, 2008.
- Manuel Amador and Kyle Bagwell. On the optimality of tariff caps. Working paper, 2010.
- 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.
- Attila Ambrus and Shih En Lu. A note on robust equilibria in multi-sender cheap talk. Working paper, 2010.
- Attila Ambrus and Satoru Takahashi. Multi-sender cheap talk with restricted state spaces. *Theoretical Economics*, 3(1):1–27, 2008.
- Mark Armstrong. Delegating decision-making to an agent with unknown preferences. Mimeograph, University College London, 1995.
- Susan Athey, Andrew Atkeson, and Patrick J. Kehoe. The optimal degree of discretion in monetary policy. *Econometrica*, 73(5):1431–1475, 2005.
- Marco Battaglini. Multiple referrals and multidimensional cheap talk. *Econometrica*, 70(4):1379–1401, 2002.
- Roland Benabou and Guy Laroque. Using privileged information to manipulate markets: Insiders, gurus, and credibility. *Quarterly Journal of Economics*, 107(3):921–958, 1992.
- Patrick Billingsley. *Probability and Measure*. 3rd edition, 1995.
- 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):2361–82, 2010.
- Zachary Cohn. A note on linked bargaining. *Journal of Mathematical Economics*, 46:238–247, 2010.
- Juan Escobar and Juuso Toikka. A folk theorem with markovian private information. 2009. Working paper.
- Alexander Frankel. Discounted quotas over actions. Working paper, 2010a.
- Alexander Frankel. Aligned delegation (extended working paper). December 2010b.
- William Fuchs. Contracting with repeated moral hazard and private evaluations. *American Economic Review*, 97(4):1432–1448, 2007.
- Itzhak Gilboa and David Schmeidler. Maximin expected utility with a non-unique prior. *Journal of Mathematical Economics*, 18:141–153, 1989.
- Jason Hartline and Anna Karlin. *Profit Maximization in Mechanism Design*, pages 331–362. In: Algorithmic Mechanism Design. Cambridge University Press, 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.
- Leonid Hurwicz and Leonard Shapiro. Incentive structures maximizing residual gain under incomplete information. *The Bell Journal of Economics*, 9(1):180–191, 1978.
- 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.
- Vijay Krishna and John Morgan. Asymmetric information and legislative rules: some amendments. *American Political Science Review*, 95:435–452, 2001.
- Vijay Krishna and John Morgan. Contracting for information under imperfect commitment. *RAND Journal of Economics*, 39(4):905–925, 2008.
- Edward P. Lazear and Sherwin Rosen. Rank-order tournaments as optimum labor contracts. *Journal of Political Economy*, 89(5):841–864, 1981.

- James M. Malcomson. Work incentives, hierarchy, and internal labor markets. *Journal of Political Economy*, 92(3):486–507, 1984.
- Andrey Malenko. Optimal design of internal capital markets, February 2011. Working Paper.
- 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.
- Stephen Morris. Political correctness. *Journal of Political Economy*, 109(21), 2001.
- Roger B. Myerson. Multistage games with communication. *Econometrica*, 54(2): 323–358, 1986.
- Timofiy Mylovanov and Andriy Zapechelnuk. Contracts for experts with opposing interests. Working paper, 2008.
- Alessandro Pavan, Ilya Segal, and Juuso Toikka. Dynamic mechanism design: Incentive compatibility, profit maximization and information disclosure. 2010.
- Thomas Sargent and Lars Peter Hansen. *Robustness*. Princeton University Press, 2007.
- Mark A. Satterthwaite and Steven R. Williams. The optimality of a simple market mechanism. *Econometrica*, 70(5), 2002.
- Joel Sobel. A theory of credibility. *Review of Economic Studies*, 52(4):557–573, 1985.
- Donald M. Topkis. *Supermodularity and Complementarity*. Princeton University Press, 1998.
- Robert Ferenc Veszteg. Linking decisions with moments. 2005.

APPENDIX

A Elaborations on Extensions

A.1 General Moment Conditions

Take some set of functions $\langle s^{(j)} : \mathcal{A} \rightarrow \mathcal{R} \rangle_{j=1}^J$. I seek assumptions on \mathcal{U}_A such that the max-min optimal mechanism specifies a sequence of real values $\langle K^{(j)} \rangle_{j=1}^J$ and then gives the agent freedom to take any action distributions subject to the constraints that $\mathbb{E} \sum_i s^{(j)}(a_i) = K^{(j)}$ for each j . One way to construct such an environment is to let $U_A(a|\theta) = U_P(a|\theta) + \lambda^{(1)} s^{(1)}(a) + \dots + \lambda^{(J)} s^{(J)}(a)$ for unknown and unbounded constants $\langle \lambda^{(j)} \rangle_{j=1}^J$ in \mathbb{R} .²³

For instance, suppose that \mathcal{U}_A is the set of functions of the form $U_A(a|\theta) = U_P(a|\theta) + c(a)$, with $c(\cdot)$ any polynomial of degree J . Then the max-min optimal mechanism (in a simultaneous or sequential problem) fixes $\sum_i (a_i)^j$ for each $j = 1, \dots, J$.

A.2 Asymptotically First-Best Payoffs

Here, I show that optimal probability assignment mechanisms provide the principal with “high” payoffs (close to the full information first-best) when there are many iid decisions from a known distribution. This holds in a simultaneous or sequential problem, as long as preferences are PA-aligned.

The optimal measure may be difficult to solve for, especially without the utility assumptions which ensure that the measure will be a set of N unit atoms. But we can approximate first-best payoffs with a “naive” measure. The naive measure places a mass on an action proportional to the ex ante probability that the action will be optimal for the principal. If the probability assignment mechanism with the naive measure approximates first-best payoffs, then so too does the one with the optimal measure.

²³This construction generalizes that of the quadratic loss constant bias case; the quadratic loss linear bias case is actually not of this form. These preferences are all in the PA-aligned class of altruistic preferences with private costs.

Let $a_P^*(\theta)$ be some function mapping θ into a principal-optimal action, i.e., $a_P^*(\theta) \in \operatorname{argmax}_a U_P(a|\theta)$. Let the *naive measure* μ^{naive} be a proper measure defined by

$$\mu^{naive}(B) = \sum_i \operatorname{Prob}[a_P^*(\theta_i) \in B] \text{ for } B \subseteq \mathcal{A}$$

I refer to the (simultaneous or sequential) probability assignment mechanism characterized by the naive measure as the naive probability assignment mechanism.

Take a sequence of decision problems indexed by $n = 1, 2, \dots$, with $N = n$ decisions in problem n . Say that this is an *iid sequence of decision problems* if each problem has the same timing (simultaneous or sequential); is over the same \mathcal{A} and Θ ; has the same utilities, U_P and U_A ; and there is a distribution F over Θ such that for each decision problem, all states are drawn iid according to F .

Definition. Fix a sequence of iid decision problems. A sequence of mechanisms $D^{(n)}$ gives *uniformly asymptotically first-best payoffs* if there exists a corresponding sequence of optimal strategies $\sigma^{(n)}(U_A)$ such that for all $\epsilon > 0$ there exists \bar{n} so that if $n > \bar{n}$, then in the n^{th} decision problem (with $N = n$),

$$\frac{\mathbb{E}_{\theta, \mathbf{a}} \sum_{i=1}^N [U_P(a_P^*(\theta_i)|\theta_i) - U_P(a_i|\theta_i) \mid D^{(n)}, \sigma^{(n)}(U_A)]}{N} < \epsilon \text{ for all } U_A \in \mathcal{U}_A.$$

This is uniform with respect to agent types – for a large enough number of decisions, the payoff is within ϵ of first-best for every possible agent utility U_A .

Proposition 5. In either a simultaneous or sequential environment, take any iid sequence of decision problems. If preferences are PA-aligned,²⁴ then the naive probability assignment mechanisms give the principal uniformly asymptotically first-best payoffs.

²⁴That is, preferences are PA-aligned in each separate decision problem with different numbers of decisions. For the PA-aligned utility classes considered in the paper, preferences are PA-aligned for any distribution over states and any number of decisions.

The proof, in Appendix B, follows similar constructions to Jackson and Sonnenschein (2007). I construct a strategy under the naive probability assignment mechanism which gives the principal nearly first-best payoffs when the empirical distribution is close to the theoretical one. As the number of iid decisions increases, a law of large numbers guarantees that these distributions are in fact close. By aligned delegation, the agent’s actual strategy gives the principal at least as high a payoff as does the one I construct.

In the working paper Frankel (2010b), I show that we can extend these asymptotically first-best payoff results to the general discounted environment in which decision i has significance $\gamma_i > 0$. This holds if, as the number of decisions grows, each individual decision becomes vanishingly important. Writing the significance of decision i in the n^{th} problem in the sequence as $\gamma_i^{(n)}$, the sufficient condition is that

$$\lim_{n \rightarrow \infty} \frac{\max_{i \leq n} \gamma_i^{(n)}}{\sum_i \gamma_i^{(n)}} = 0$$

B Omitted Proofs

Proof of Lemma 1. Here I prove Lemma 1 for either a simultaneous or a sequential problem (see Section 5).

For a measure μ over the set of actions \mathcal{A} , for $k = 0, 1, 2, \dots$, define the k^{th} moment of μ as

$$\text{Mom}^k(\mu) \equiv \int_{\mathbb{A}} a^k d\mu(a)$$

In order to show that two measures are identical, it suffices to show that all of their moments are equal. This follows from the compactness of \mathcal{A} and Θ ; see, e.g., Billingsley (1995) Theorem 30.1. And to show that a sequence of measures approaches some limiting measure (in the sense of weak convergence), it suffices to show that each fixed moment converges to the limiting moment; see Billingsley (1995) Theorem 30.2.²⁵

²⁵The cited results are stated for distributions rather than general measures, using the first moment and above; rescaling measures by a factor of $\frac{1}{N}$, this becomes identical.

Consider a public history – a list of all past messages and actions – at some decision node for the agent. At the initial message stage, call this history $h_0 = (\emptyset)$. For the simultaneous problem there are two more subsequent histories, $h_1 = (m_0)$ and $h_{N+1} = (m_0, m, \mathbf{a})$. For the sequential problem there are $N + 1$ further histories $h_1 = (m_0)$, $h_2 = (m_0, m_1, a_1)$, ..., $h_{N+1} = (m_0, m_1, a_1, \dots, m_N, a_N)$. There are no additional reports or actions at history h_{N+1} – at this history, the game is over.

Let $\mu_{\sigma, \theta}^D(h_i)$ be the measure of remaining actions from period i onward, given public history h_i , if past and future states are given by θ .²⁶ For instance, the measure $\mu_{\sigma, \theta}^D(h_0)$ is equal to $\mu_{\sigma, \theta}^D$ and the measure $\mu_{\sigma, \theta}^D(h_{N+1})$ places a mass of 0 on any set. Let N_i be the remaining mass; $N_0 = N_1 = N$, $N_i = N - i + 1$ for $i \leq N + 1$.

For each of the (s, t) cases, starting from any history h_i , I will define the infinite sequence of moments $(\alpha^0(h_i), \alpha^1(h_i), \dots)$ that are “most desirable” for the agent as we take n and then $|\lambda|$ to infinity.

Let $\alpha^0(h_i) = N_i$ in all cases, for any history h_i .

To define $\alpha^k(h_i)$ for $k \geq 1$, first let $Z(k, \epsilon, h_i)$ be the set of k^{th} moments for which all lower moments $l < k$ are within ϵ of $\alpha^l(h_i)$.

$$Z(k, \epsilon, h_i) = \left\{ \begin{array}{l} \text{Mom}^k(\mu_{\sigma, \theta}^D(h_i)) \text{ s.t. } \sigma \in \Sigma^D, \theta \in \Theta^N \text{ consistent with } h_i, \\ \& |\text{Mom}^l(\mu_{\sigma, \theta}^D(h_i)) - \alpha^l(h_i)| < \epsilon \text{ for each } l < k \end{array} \right\}$$

Now define $\alpha^k(h_i)$ inductively, given $\alpha^0(h_i), \dots, \alpha^{k-1}(h_i)$.²⁷

²⁶The notation is slightly redundant – past θ 's are included in both the history and the vector θ .

²⁷I do not include current and past states in the construction of $\alpha^k(h_i)$ even though they may affect an agent's strategy, because they do not affect the set of possible action distributions going forward. The agent can always play as if the states had been something else.

$$\begin{aligned}
\text{If } (s, t) = (-1, 1) : \quad & \alpha^k(h_i) = \lim_{\epsilon \rightarrow 0^+} \inf Z(k, \epsilon, h_i) \\
(-1, -1) : \quad & \alpha^k(h_i) = \begin{cases} \lim_{\epsilon \rightarrow 0^+} \sup Z(k, \epsilon, h_i) & \text{if } k \text{ is odd} \\ \lim_{\epsilon \rightarrow 0^+} \inf Z(k, \epsilon, h_i) & \text{if } k \text{ is even} \end{cases} \\
(1, 1) : \quad & \alpha^k(h_i) = \lim_{\epsilon \rightarrow 0^+} \sup Z(k, \epsilon, h_i) \\
(1, -1) : \quad & \alpha^k(h_i) = \begin{cases} \lim_{\epsilon \rightarrow 0^+} \inf Z(k, \epsilon, h_i) & \text{if } k \text{ is odd} \\ \lim_{\epsilon \rightarrow 0^+} \sup Z(k, \epsilon, h_i) & \text{if } k \text{ is even} \end{cases}
\end{aligned}$$

For each of these richness cases, the list of moments $(\alpha^0(h_i), \alpha^1(h_i), \alpha^2(h_i), \dots)$ uniquely defines a measure over \mathcal{A} . (The moments define *some* measure because they were found as a limiting sequence of the moments of other measures). Let the measure μ_∞^D be the one implied by history h_0 .

Claim 1. For any public history h_i , any $k \geq 0$, and any $\epsilon > 0$, fix an exponent $n > k$ and a vector of states θ . For each of the four (s, t) cases with $U_A = \psi(a|\theta; \lambda, n) + s \cdot (a + t\lambda)^{2n}$, if $|\lambda|$ is large enough then

$$|\text{Mom}^k(\mu_{\sigma, \theta}^D(h_i)) - \alpha^k(h_i)| < \epsilon$$

for any $\sigma \in \Sigma^{*D}(U_A)$.

This holds for a simultaneous or sequential problem.

This claim implies the result. For each of the richness cases, taking n to infinity and then $|\lambda|$ to infinity, we can find a sequence of utility functions $\langle U_A^j \rangle$ for which each fixed k^{th} moment converges to α^k for every state vector and every optimal strategy at the null history h_0 . Therefore the sequence of agent strategies $\langle \sigma^j \rangle$ takes $\mu_{\sigma^j, \theta}^D$ to μ_∞^D .

Proof of Claim 1. I prove this by backwards induction on the period i .

The claim holds for any history h_i with $i = N + 1$ because $\mu_{\sigma, \theta}^D(h_{N+1})$ is always the 0-measure, for which every moment is 0.

Inductive hypothesis: Suppose the claim holds for all histories h_i with $i > i'$. I seek to show that under the inductive hypothesis the claim also holds for any history $h_{i'}$ as well.

At history $h_{i'}$, after observing any states revealed in that period, the agent's expected future payoff given strategy σ can be written as

$$\begin{aligned} \mathbb{E} \sum_{l=i'}^N & \left[\psi(a_l | \theta_l; \lambda, n) + s \binom{2n}{0} a_l^{2n} + s \binom{2n}{1} t \lambda a_l^{2n-1} - \dots + s \binom{2n}{n} t^n \lambda^n a_l^n \right] \\ & + s \mathbb{E} \sum_{k=0}^{n-1} \binom{2n}{2n-k} t^{2n-k} \lambda^{2n-k} \text{Mom}^k (\mu_{\sigma, \theta}^D(h_{i'})) \end{aligned} \quad (3)$$

where the expectations are over future states as well as current and future actions.

Now, take a sequence of utility functions, indexed by j , which each have the same fixed n but for which $\lambda \rightarrow \infty$, and a respective sequence of optimal strategies σ^j . Suppose there is some $k < n$ such that the k^{th} moment $\mathbb{E} \left[\text{Mom}^k (\mu_{\sigma^j, \theta}^D(h_{i'})) \right]$ does not approach $\alpha^k(h_{i'})$; without loss of generality, assume it is bounded away from this value. Well, by construction of $\alpha^k(h_{i'})$ we can construct an alternate strategy σ' for which all of the moments $\mathbb{E} \left[\text{Mom}^l (\mu_{\sigma', \theta}^D(h_{i'})) \right]$ are arbitrarily close to $\alpha^l(h_{i'})$ for each $l < n$. (This strategy σ' may be taken to be state-independent, so that the expectation over future states is irrelevant.)

As λ goes to ∞ , we can find such a σ' which must eventually be strictly preferred to the proposed optimal strategy σ^j . That's because the first sum in (3) is bounded by a constant expression times λ^n , so the difference between σ' and σ^j is as well. But the difference in the second sums goes as at least $|\lambda|^{2n-k > n}$ times the difference in k^{th} moments, which is positive. (Each term in this difference of moments is nonnegative, so the difference in k^{th} moments is a lower bound for the total difference). □ ■

Proof of Lemma 2. The payoff from any mechanism is bounded above by $N \cdot \max_{a, \theta} U_P(a | \theta)$, and so we can find a sequence of proper measures $\langle \mu^n \rangle_{n=1}^\infty$ for

which the payoffs of $\text{PA}(\mu^n)$ approach the supremum over measures of probability assignment payoffs. There is a measure achieving this limiting payoff by the compactness of the set of measures with respect to weak convergence combined with the continuity of payoffs as guaranteed by Claim 2.

Claim 2. Suppose that utilities are PA-aligned. Take a sequence of proper measures $\langle \mu^n \rangle_{n=1}^\infty$ converging weakly to μ . Any player's expected payoff from $\text{PA}(\mu^n)$ approaches that from $\text{PA}(\mu)$.

Proof of Claim 2. As n goes to infinity, the agent can choose distributions under $\text{PA}(\mu^n)$ which weakly approach those from optimal play under $\text{PA}(\mu)$ for each action a_i . So payoffs of $\text{PA}(\mu^n)$ are eventually bounded below by any payoff arbitrarily smaller than that from $\text{PA}(\mu)$ (for the agent and, by aligned delegation, for the principal). By a symmetric argument, the payoff from $\text{PA}(\mu)$ is bounded below by any payoff arbitrarily smaller than that from $\text{PA}(\mu^n)$, for large enough n . So the $\text{PA}(\mu^n)$ payoffs approach those of $\text{PA}(\mu)$. \square ■

Proof of Lemma 3. 1. 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 $\max \text{Supp } m_i > \min \text{Supp } m_j$. Then we can find measures $\nu_i \leq m_i$ and $\nu_j \leq m_j$, each placing a mass $\delta > 0$ on \mathcal{A} , such that the support of ν_i is strictly above the support of ν_j . Consider swapping these measures, replacing the assignment m_i with $m'_i = m_i - \nu_i + \nu_j$ and m_j with $m'_j = m_j - \nu_j + \nu_i$ and holding all other assignments fixed. The payoff change to the agent is

$$\int_{\mathcal{A}} (U(a|\theta_i) - U(a|\theta_j)) d\nu_j(a) - \int_{\mathcal{A}} (U(a|\theta_i) - U(a|\theta_j)) d\nu_i(a)$$

Let $H^\nu : [0, \delta] \rightarrow \mathcal{A}$ be the inverse cumulative mass function of $\nu = \nu_i, \nu_j$, defined (for concreteness) by $H^\nu(x) = \min\{a \in \mathcal{A} | \nu([\underline{a}, a] \cap \mathcal{A}) = x\}$.²⁸

It holds that $H^{\nu_i}(x) > H^{\nu_j}(x)$ for all $x \in (0, \delta)$. We can now rewrite the

²⁸The inverse function is in general uniquely defined everywhere but a measure 0 of points; this construction chooses the lowest possible points whenever there is freedom.

payoff change as

$$\int_0^\delta \left[(U(H^{\nu_j}(x)|\theta_i) - U(H^{\nu_i}(x)|\theta_i)) - (U(H^{\nu_j}(x)|\theta_j) - U(H^{\nu_i}(x)|\theta_j)) \right] dx$$

And this is nonnegative if U satisfies increasing differences, because it is nonnegative for each x .

Now, starting from any non-assortative assignment, we can perform a sequence of such swaps to get to an assortative assignment. Each such swap weakly increases payoffs, so the payoff from this resulting assortative assignment is at least as high as the payoff from the nonassortative one. And all assortative assignments are payoff equivalent, so they must give the agent an optimal payoff.

2. By part 1, both player's payoffs are maximized by an assortative strategy. ■

Proof of Lemma 4. Fix some function $U : \mathcal{A} \times \Theta \rightarrow \mathbb{R}$. If \mathcal{U}_A contains all functions of the form $U(a|\theta) - (a - \lambda)^{2n}$, for $n \rightarrow \infty$ and for λ going to plus or minus infinity, then richness is satisfied with $\psi = U$, independent of λ, n .

If U satisfies strict increasing differences, or concavity, then the function $U(a|\theta) - (a - \lambda)^{2n}$ does as well ($s = -1, t = -1$). For convexity, we can look at a convex U with $s = 1, t = -1$. ■

Proof of Proposition 1. Given $\text{PA}(\mu)$, an assortative strategy assigns the j^{th} through $j + 1^{\text{st}}$ quantiles of measure to the j^{th} lowest realized state. Let F_j^μ be the distribution over which the action for this decision will be drawn:

$$F_j^\mu(a) = \left[\frac{1}{\gamma_i} \left(\mu((-\infty, a]) - \sum_{j \text{ s.t. } \pi(j) < \pi(i)} \gamma_j \right) \right]$$

where $\llbracket y \rrbracket$ is defined as 0 if $y < 0$; y if $y \in [0, 1]$; and 1 if $y > 1$.

Let G_j be the ex ante distribution of the j^{th} lowest state, taking expectation over realizations over θ .

The principal's expected payoff from an assortative strategy in $\text{PA}(\mu)$ can now be written as

$$\sum_j \int_{\Theta} \int_{\mathcal{A}} U_P(a|\theta) dF_j^\mu(a) dG_j(\theta)$$

This is maximized over all possible F_j^μ if F_j^μ is a degenerate distribution placing all probability on some single action in $\text{argmax}_a \int_{\Theta} U_P(a|\theta) dG_j(\theta)$. We can induce this F_j^μ by choosing μ to be a sum of these N degenerate distributions. ■

Proof of Lemma 5. Write U_A as $U_A(a|\theta) = 2c(a)\theta - \theta^2 - (c(a))^2$.

Consider increasing functions of the form $c(a) = (a - \lambda)^n$, for n odd. Then $\psi(a|\theta; \lambda, n) = 2c(a)\theta - \theta^2$ is an n^{th} degree polynomial in λ , and $U_A(a|\theta) = \psi(a|\theta; \lambda, n) - (a - \lambda)^{2n}$. (This has $s = -1$, $t = -1$). ■

Proof of Lemma 6. I first establish a simple equality showing that preferences over actions depend only on expected states:

Claim 3. Take U a quadratic loss utility function, and let a and θ be independent random variables. Then $\mathbb{E}[U(a|\theta)] = \mathbb{E}[U(a|\mathbb{E}[\theta])] - \text{Var}[\theta]$.

Proof of Claim 4.

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

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.

Define $\tilde{\theta}_i : \{1, \dots, N - i + 1\} \times \Theta^{i-1} \rightarrow \mathbb{R}$ by backwards induction. 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 $\tilde{\theta}_i(j; \theta_1, \dots, \theta_{N-1}) = \mathbb{E}_{\theta_i} [R^{(j)}(\langle \theta_i, \tilde{\theta}_{i+1}(1), \dots, \tilde{\theta}_{i+1}(N - i) \rangle)]$. I write $\tilde{\theta}_i(j; \theta_1, \dots, \theta_{N-1})$ as $\tilde{\theta}_i(j)$ if the past states are otherwise implied.

We interpret $\tilde{\theta}_i(j)$ as the expected value of the state in which the j^{th} lowest remaining action will be played, prior to the realization of θ_i . 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 $\tilde{\theta}_{i+1}(1), \dots, \tilde{\theta}_{i+1}(N-i)$. This gives a *sequential-assortative* strategy: action a_i is assigned to quantiles of remaining measure $k-1$ through k , for k such that θ_i is the k^{th} lowest state of itself and the expected future states: $\theta_i = R^{(k)} \left(\left\langle \theta_i, \tilde{\theta}_{i+1}(1), \dots, \tilde{\theta}_{i+1}(N-i) \right\rangle \right)$.²⁹

The following claim shows that a sequential-assortative strategy is optimal for any quadratic loss utility, and hence that the mechanism satisfies aligned delegation.

Claim 4. Let F_j^μ be the distribution over actions corresponding to quantiles $j-1$ through j of measure μ , as in the proof of Proposition 1. Let U be some quadratic loss utility function.

Given a sequential decision problem, there exists C such that for all proper measures μ , a sequential-assortative strategy in $\text{PA}(\mu)$ gives a player with utility U a payoff of $\sum_{j=1}^N \int_{\mathcal{A}} U(a|\tilde{\theta}_1(j)) dF_j^\mu(a) - C$. Any alternative strategy gives a weakly lower payoff.

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

Proof of Claim 4. 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 measure μ_i . I seek to show that a sequential-assortative strategy gives $\sum_{j=1}^{N-i+1} \int_{\mathcal{A}} U(a|\tilde{\theta}_i(j)) dF_j^{\mu_i}(a) - C_i$, for C_i independent of μ_i , and other strategies give weakly less.

For $i = N$, this holds by Claim 3: all strategies give a payoff of $\int_{\mathcal{A}} U(a|\tilde{\theta}_i(1)) dF_j^{\mu_i}(a) - \text{Var}[\theta_N]$, where $\tilde{\theta}_i(1)$ is the expected value of θ_N .

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

²⁹Analogously to the simultaneous definition of assortativity, the minimum of the support of m_i (the distribution assigned to action a_i) is weakly above the j^{th} quantile of remaining measure if $\theta_i > \tilde{\theta}_{i+1}(j)$; the maximum of the support of m_i is weakly below the j^{th} quantile of remaining measure if $\theta_i < \tilde{\theta}_{i+1}(j)$.

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)

$$\sum_{j=1}^{N-i+1} \int_{\mathcal{A}} \left(\begin{cases} U(a|\tilde{\theta}_{i+1}(j)) & \text{if } j < k \\ U(a|\theta_i) & \text{if } j = k \\ U(a|\tilde{\theta}_{i+1}(j-1)) & \text{if } j > k \end{cases} \right) dF_j^{\mu_i}(a) - C_{i+1} \quad (4)$$

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 $F_j^{\mu_i}$ in (4) is $\tilde{\theta}_i(j)$. So applying Claim 3, we get

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

for C_i equal to $\mathbb{E}[C_{i+1}]$ minus the sum of the variance constants.³⁰

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 probability assignment 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} . By Lemma 3, the simultaneous payoff would be maximized by assortative m_i . This corresponds to maximizing the sequential payoff by choosing m_i sequential-assortatively. □ ■

Proof of Proposition 2. Given that the agent will play a sequential-assortative strategy, it suffices to show that the optimal measure in a probability assign-

³⁰ C_{i+1} may depend on the θ_i to the extent that the joint distribution of future states depends on this realization.

ment will be a sum of N unit atoms. From Claim 4 in the proof of Lemma 6, a sequential-assortative strategy in $\text{PA}(\mu)$ gives the principal a payoff of $\sum_{j=1}^N \int_{\mathcal{A}} U_P(a|\tilde{\theta}_1(j)) dF_j^\mu(a) - C$ for some C independent of μ . So the principal's payoff is maximized over measures by having F_j^μ choose a single action in $\text{argmax}_a U_P(a|\tilde{\theta}_1(j))$ with certainty, i.e., choosing μ to be a sum of N unit atoms. ■

Proof of Lemma 7. We can rewrite the agent's utility as

$$\begin{aligned} -(a - \theta - \lambda)^2 &= (-\theta^2 + 2\theta(a - \lambda)) - (a - \lambda)^2 \\ &= \psi(a|\theta; \lambda) - (a - \lambda)^{2n} \text{ for } n = 1, \psi(a|\theta; \lambda) = (-\theta^2 + 2\theta(a - \lambda)) \end{aligned}$$

The proof now follows that of Lemma 1, considering only the first moment $\mathbb{E}[\sum_i \gamma_i a_i | \sigma^*, \theta]$ of the induced measure $\mu_{\sigma^*, \theta}^D$. ■

Proof of Lemma 8.

1. From equation (1), the principal and agent have an identical maximization problem conditional on fixing $\sum_i a_i$.
2. Follows from continuity of the principal's payoffs with respect to the budget (because, by aligned delegation, the agent maximizes the principal's payoff) and compactness of the set of proper budgets. ■

Proof of Proposition 3. If the agent could choose arbitrary action distributions subject to the constraint that $\mathbb{E} \sum_i a_i = K$, then the proof would follow that of Theorem 1 almost identically, observing that the budget mechanism is aligned delegation with QLCB preferences and replacing Lemma 1 with Lemma 7.

It only remains to show that, given the freedom to choose arbitrary distributions, the agent would choose deterministic actions. This is because preferences are concave in actions. Given any state and any proposed nondegenerate distribution m_i at period i , an agent with quadratic loss constant bias utility would prefer choosing action a_i deterministically at the expectation of m_i . Such an action is feasible because the action space is assumed to be a convex interval. If the action space had holes, randomization might be required. ■

Proof of Lemma 9.

1. From equation (2), the principal and agent have an identical maximization problem conditional on fixing $\sum_i a_i$ and $\sum_i (a_i)^2$.
2. Follows from continuity of the principal's payoffs with respect to the moment levels $K^{(1)}$ and $K^{(2)}$ (because, by aligned delegation, the agent maximizes the principal's payoff) and compactness of the set of proper moments. ■

Proof of Proposition 4. Follows the proof of Theorem 1, observing that two moment mechanisms are aligned delegation with QLLB preferences and replacing Lemma 1 with an analog of Lemma 7:

Claim 5. In either a simultaneous or sequential problem, let the agent have unbounded quadratic loss constant bias utilities. Fix any mechanism D . There exist values $K^{(1)}$ and $K^{(2)}$ and a sequence of types $\langle U_A^j \rangle_{j=1}^\infty$ in \mathcal{U}_A such that for all states θ and all corresponding sequences of optimal agent strategies $\langle \sigma^j \rangle$, it holds that

$$\begin{aligned} \mathbb{E} \left[\sum_i \gamma_i a_i \middle| \sigma^j, \theta \right] &\rightarrow K^{(1)} \text{ as } j \rightarrow \infty \\ \mathbb{E} \left[\sum_i \gamma_i (a_i)^2 \middle| \sigma^j, \theta \right] &\rightarrow K^{(2)} \text{ as } j \rightarrow \infty \end{aligned}$$

Proof of Claim 5. The agent's utility $-(a - \lambda^{(1)}\theta - \lambda^{(0)})^2$ is equivalent to $2\lambda^{(1)}\theta a - a^2 + 2\lambda^{(0)}a$. As $|\lambda^{(0)}| \rightarrow \infty$ and $\lambda^{(1)} \rightarrow 0$, this goes to $a^2 + 2\lambda^{(0)}a$; the agent lexicographically maximizes (if $\lambda^{(0)} \rightarrow \infty$) or minimizes (if $\lambda^{(0)} \rightarrow -\infty$) the first moment of the measure $\mathbb{E}[\sum_i a_i]$; then minimizes the second moment of the measure $\mathbb{E}[\sum_i (a_i)^2]$; and only after these are fixed looks at the term $\mathbb{E}[\sum_i a_i \theta_i]$ which is multiplied by $\lambda^{(1)} \simeq 0$ and depends on the states. From here, the argument follows the proof of Lemma 1. □ ■

Proof of Proposition 5. By aligned delegation, it suffices to show that there exists some strategy of the agent for which the principal's weighted expected per-period payoff loss $U_P(a_P^*(\theta_i)|\theta_i) - U_P(a(\theta_i)|\theta)$ is arbitrarily close to 0. An

aligned strategy, which will be optimal for the agent, gives the principal a weakly smaller payoff loss. And under strict aligned delegation, every optimal strategy for an agent is aligned. I propose a strategy which can be played in a sequential problem, and so can be replicated for a simultaneous problem.

For any positive integer L , divide the state space Θ into L “bins” in the following manner: For $1 \leq l < L$, bin l is the set $\Theta^{l,L} = \Theta \cap [\underline{a} + (\bar{a} - \underline{a})\frac{l-1}{L}, \underline{a} + (\bar{a} - \underline{a})\frac{l}{L})$. Bin L is the set $\Theta^{L,L} = \Theta \cap [\underline{a} + (\bar{a} - \underline{a})\frac{L-1}{L}, \bar{a}]$ – closed on the right.

For any l with $\Theta^{l,L}$ nonempty, let $\mathcal{A}^{l,L} = \{a_P^*(\theta) | \theta \in \Theta^{l,L}\}$.

Given a remaining measure $\mu_i = \mu - \sum_{j < i} m_j$ with $\mu_i(\mathcal{A}^{l,L}) \geq 1$, say that the agent “places a_i in bin l (of L)” if she plays actions proportionally from the remaining measure of μ_i on the support $\mathcal{A}^{l,L}$. Formally, she chooses distribution m_i as the measure defined by

$$m_i(B) = \frac{\mu_i(\mathcal{A}^{l,L} \cap B)}{\mu_i(\mathcal{A}^{l,L})} \quad \text{for } B \subseteq \mathcal{A}$$

Tweaking the terminology of Jackson and Sonnenschein (2007), fix L and say that a strategy in the probability assignment mechanism is *approximately truthful* with respect to these bins if, whenever $\theta_i \in \Theta^{l,L}$ and $\mu_i(\Theta^{l,L}) \geq 1$, the agent places action a_i in the “appropriate” bin, bin l . When $\mu_i(\Theta^{l,L}) < 1$ the action cannot be placed in the appropriate bin and so the agent’s choice m_i may be arbitrary, subject to the feasibility constraints.

Given decision problem n , states θ , a number of bins L , and the naive initial measure, let $z^{(n)}(\theta)$ be the last period for which every action from 1 through z has been placed in its appropriate bin under an approximately truthful strategy. By continuity of U_P and compactness of \mathcal{A} and Θ , the result follows from showing that for any fixed L , as $n \rightarrow \infty$, $\mathbb{E}_\theta \frac{z^{(n)}(\theta)}{n} \rightarrow 1$.³¹

Suppose not. Suppose instead that $\mathbb{E}_\theta \frac{z(\theta)}{n}$ has lim inf strictly less than 1.

³¹For any $\epsilon > 0$, we can find L large enough such that for any l with $\Theta^{l,L}$ nonempty, for any $\theta^l \in \Theta^{l,L}$, and for any $a^l \in \mathcal{A}^{l,L}$, it holds that $U_P(a^l | \theta^l) \geq U_P(a_P^*(\theta^l) | \theta^l) - \epsilon$. Moreover, the principal’s worst possible stage utility level $\min\{U_P(a | \theta) | a \in \mathcal{A}, \theta \in \Theta\}$ is a finite value. So if a weighted proportion of actions approaching 1 are placed in their appropriate bins, then the principal’s payoff is close to his optimal payoff.

Since the expression is bounded between 0 and 1, and since there are finitely many bins, this implies that there is some l such that $\Theta^{l,L}$ is realized by $F(\cdot)$ with probability $p^{l,L} > 0$; there is some infinite subsequence of n values; and some sequence $z^{(n)}(\boldsymbol{\theta})$ for which

- $\mathbb{E}_{\boldsymbol{\theta}} \frac{z^{(n)}(\boldsymbol{\theta})}{n}$ is bounded away from 1 – say, is at most $1 - \xi$
- the probability over realizations of $\boldsymbol{\theta}$ that $\frac{\sum_{i=1}^{z^{(n)}(\boldsymbol{\theta})} \chi_{\{\theta_i \in \Theta^{l,L}\}}}{n} > p^{l,L}$ is bounded away from 0

But the weak law of large numbers says that $\frac{\sum_{i=1}^{z^{(n)}(\boldsymbol{\theta})} \chi_{\{\theta_i \in \Theta^{l,L}\}}}{\Gamma(n)}$ approaches its expectation – something at most $(1 - \xi)p^{l,L} < p^{l,L}$ – in probability, contradicting the second bullet.

■