Randomization and the Robustness of Linear Contracts

Ashwin Kambhampati* Juuso Toikka† Rakesh Vohra‡

June 11, 2024

Abstract

We consider a principal-agent model with moral hazard, bilateral risk-neutrality, and limited liability, where the principal knows only some of the agent's feasible actions. The principal evaluates contracts by their payoff guarantee over all possible sets of feasible actions. We show that the guarantee from any randomization over contracts can be weakly improved by randomizing over the same or smaller number of linear contracts. Thus, random linear contracts are robustly optimal as a class. We derive the optimal guarantee and characterize a maximizing sequence of random linear contracts attaining it in the limit as the size of their support tends to infinity; these contracts do not require commitment to randomization. The gain from randomization can be arbitrarily large.

1 Introduction

Recent years have witnessed a surge in contract-theoretic research on the optimal design of incentives when the designer does not know all the details of the contracting environment, or is otherwise reluctant to commit to a single probabilistic model of the situation. For moral hazard problems, an influential paper by Carroll (2015) offers a model where the principal is unaware of all the possible actions the agent can take and evaluates contracts based on their guaranteed payoff over possible action sets. Assuming limited liability on the part of the agent and risk neutrality of both, Carroll shows that the contract that maximizes the principal's guarantee is a linear one. This is an attractive result as it predicts a commonly observed contract form without the need to resort to specific assumptions about how actions

*Department of Economics, United States Naval Academy. Email: kambhamp@usna.edu.

†The Wharton School of Business, University of Pennsylvania. Email: toikka@wharton.upenn.edu.

‡Department of Economics & Department of Electrical and Systems Engineering, University of Pennsylvania. Email: rvohra@seas.upenn.edu.

map to outputs and costs, or, indeed, without even specifying a complete model of the situation. Subsequent work has extended the approach to other environments with moral hazard and established the robust optimality of linear contracts therein.[1]

A contract's payoff guarantee is, however, a worst-case criterion, which means that randomizing over contracts could allow the principal to secure a higher payoff. This was recently confirmed by Kambhampati (2023). He shows in the Carroll model that randomizing over two linear contracts strictly improves the principal's guarantee compared to the optimal deterministic contract whenever the latter is non-trivial (i.e., when the zero contract is not the uniquely optimal deterministic contract). Kambhampati's result leaves open the form of the optimal random contract. Most concerningly, it raises the possibility that the linearity result, and the intuition about robust optimality of interest alignment it formalizes, are driven by the restriction to deterministic contracts.

Motivated by this observation, this paper revisits the robust moral hazard problem allowing for random contracts. Specifically, we assume that the principal can commit to any (finitely-supported) randomization over which contract to give to the agent in an effort to hedge against the Knightian uncertainty regarding the agent's action set. Our main finding is that linear contracts are still robustly optimal in the sense that any random contract is weakly dominated by a randomization over the same number, or fewer, linear contracts. This implies that random linear contracts are an optimal class of contracts given any upper bound on the size of the support. (The result for deterministic contracts obtains by taking the support to be a singleton.) We also derive the optimal guarantee over all random contracts, and characterize a sequence of random linear contracts attaining it in the limit as the size of their support tends to infinity. The principal is indifferent among all linear contracts in the support of these contracts, and thus commitment to the randomization is not needed in order to (asymptotically) attain the optimal guarantee.

To quantify the importance of allowing for random contracts, we define the gain from randomization as the ratio of the optimal guarantee to the guarantee from the best deterministic contract. An example demonstrates that this gain can be arbitrarily large.

A heuristic explanation for the optimality of random linear contracts is as follows. Our proof shows that any random contract can be weakly improved by appropriately linearizing each contract in its support. The alignment of the principal's and the agent's interests achieved by linear contracts thus provides a robustly optimal way to motivate the agent, whether the contract is chosen randomly or not. In this sense, the intuition from the deter-

[1] These include common agency (Marku, Ocampo Diaz, and Tondji, 2024), moral hazard in teams (Dai and Toikka, 2022), contracting with an organization (Walton and Carroll, 2022), and settings where the principal is also subject to moral hazard (Carroll and Bolte, 2023).

ministic case carries over. The randomization then hedges against the uncertainty regarding the agent's set of feasible actions as it prevents Nature from targeting the action set to any particular contract.

More formally, our proof is based on explicitly formulating a random contract's guarantee as a linear program for adversarial Nature. This problem is in general a multidimensional mechanisms design problem, where the agent's type is the realized contract. Nature designs for each type an action subject to usual incentive compatibility constraints and a participation constraint stemming from the presence of actions known to the principal. Our proof proceeds by using the dual to Nature's problem to identify a random linear improvement contract. Applied to the deterministic case in Section 3, the approach yields a short proof of Carroll's (2015) result, with linear programming duality used in place of his separating hyperplane argument.

That randomization may provide a hedge against Knightian uncertainty, or ambiguity, was suggested by Raiffa (1961) in response to Ellsberg (1961). In the present context, the question ultimately comes down to the principal's attitude toward uncertainty; decision theory provides axiomatizations for both extremes and even for intermediate cases (see, e.g., Saito, 2015; Ke and Zhang, 2020). If we adopt the interpretation that the environment is chosen by adversarial Nature, then the issue can be phrased as a question of timing or observability. Randomization has value if Nature can condition its choice on the random contract, but not on its realization. If instead Nature can react to the realized contract, then randomization has no value and the analysis reduces to that in Carroll (2015).

In practice, if the uncertainty regards the realization of some physical process such as the choice of agent or the determination of feasible production technologies, then it seems reasonable to assume that the resolution of a lottery over contracts just prior to one is presented to the agent should have no effect on these phenomena, and randomization should thus have value. Interestingly, this is the position (implicitly) adopted in much of the recent literature on robust mechanism design as it is common to allow for random mechanisms, which means that the designer has access to lotteries that are resolved only after the environment has been determined (see, e.g., Brooks and Du, 2024, and the references therein). Allowing for randomization in max-min problems is also standard in computer science.

This work is related to the literature seeking foundations for linear and other commonly observed contracts following Holmström and Milgrom (1987), who concluded their fully Bayesian analysis of the optimality of linear intertemporal incentives by suggesting that the reason for the popularity of linear schemes might be their great robustness, and that the case for it could perhaps be made more effectively with a non-Bayesian model. We refer the reader to Carroll (2015) for a review of this literature. Of the more recent works on non-Bayesian models, most closely related are those by Dai and Toikka (2022), Walton and Carroll (2022), Carroll and Bolte (2023), and Marku, Ocampo Diaz, and Tondji (2024), who show the optimality of deterministic linear contracts in different moral hazard environments. Related models are studied also by Antic (2021), Antic and Georgiadis (2023), Rosenthal (2023), Burkett and Rosenthal (2023), and Kambhampati (2024), who find worst-case optimal deterministic contracts different from linear contracts. More broadly, our work belongs to the literature on the design of worst-case optimal contracts; see Carroll (2019) for a recent survey.

After having presented early versions of this work at several seminars and conferences, we learned of independent contemporaneous work by Peng and Tang (2024), who consider the same problem. Our approaches are different. Peng and Tang conjecture the form of the optimal contract in the space of all distributions over contracts, and then verify its optimality. In contrast, we give an improvement argument showing that random linear contracts are optimal even given any bound on the size of the contracts' support. We then derive the form of the optimal contract for any (sufficiently large) grid. Both papers give the same formula for the optimal guarantee, and the limit of our maximizing sequence of contracts in Proposition 7 corresponds to Peng and Tang's optimal contract. Peng and Tang's proof is elegant, but it is difficult to distill an intuition from it as to why linearity is optimal. As discussed above, our proof allows one to separate the role of the contract form and that of the randomization.

The rest of this paper is organized as follows. Section 2 sets up the model. Section 3 revisits the deterministic case. Section 4 presents our main result on the optimality of random linear contracts. Section 5 derives the optimal guarantee and the maximizing sequence. Section 6 concludes. The Appendix contains four proofs omitted from the main text.

2 Model

We consider the problem of a principal designing a contract to motivate an agent subject to moral hazard, with the principal facing non-quantifiable uncertainty over the agent's set of feasible actions.

Let $Y := \{y_1, \ldots, y_n\} \subset \mathbb{R}_+$ be the finite set of possible output levels, labeled in increasing order so that $0 \equiv y_1 < \cdots < y_n$. It will be convenient to view the possible output levels as a vector $y = (y_1, \ldots, y_n) \in \mathbb{R}^n_+$ and introduce the index set $I := \{1, \ldots, n\}$.

An (unobservable) action for the agent is a pair $(\pi, c) \in \Delta(Y) \times \mathbb{R}_+$, where $\pi = (\pi_1, \ldots, \pi_n)$ is the output distribution and c is the associated cost to the agent.[2] A *technology*

[2] As is standard, we denote the set of all probabilities on a finite set B by $\Delta(B)$. More generally, given

is a nonempty finite set of actions $A \subset \Delta(Y) \times \mathbb{R}_+$.

The agent is protected by limited liability, requiring payments to him to be non-negative. A *(deterministic) contract* is thus a non-negative n-vector $w = (w_1, \ldots, w_n) \in \mathbb{R}^n_+$, where w_i is the payment from the principal to the agent if output y_i is realized.

Both parties are assumed risk-neutral. Thus, if the agent plays the action (π, c) given contract w, then the principal's expected payoff is $\pi(y - w) = \sum_{i \in I} \pi_i(y_i - w_i)$, whereas the agent gets $\pi w - c = \sum_{i \in I} \pi_i w_i - c$.[3]

The principal does not know the actual technology available to the agent. She is only aware of some technology A^0, referred to as the *known technology*, and views any technology $A \supseteq A^0$ as possible. Faced with this uncertainty, the principal evaluates contracts based on their guaranteed performance over all technologies that contain the known one. To state this formally, given a contract w and a technology A, denote the agent's maximum payoff by

$$U(w|A) := \max_{(\pi,c) \in A} \pi w - c. \tag{2.1}$$

In the special case of the known technology A^0, we write

$$U^0(w) := U(w|A^0). \tag{2.2}$$

We note for future reference that $U^0(w)$ is continuous in the contract w.

The corresponding payoff to the principal is

$$V(w|A) := \min_{(\pi,c) \in A} \pi(y - w) \quad \text{s.t.} \quad \pi w - c = U(w|A). \tag{2.3}$$

Note, in particular, that if the agent has multiple maximizers, then ties are broken against the principal. The principal's *guarantee* from the contract w is then

$$V(w) := \inf_{A \supseteq A^0} V(w|A). \tag{2.4}$$

We assume that the principal can hedge against the uncertainty over the agent's technology by randomizing over contracts. Formally, a *random contract* is a finitely supported probability measure p on $\mathbb{R}^n_+$. Let $\Delta(\mathbb{R}^n_+)$ denote the space of all random contracts. Given a random contract p, the agent observes the realized contract w before choosing an action.

any set B, we use $\Delta(B)$ to denote the set of all finitely-supported probability measures on B.

[3]Given two vectors v and u in $\mathbb{R}^d$, we write $uv := \sum_{i=1}^d u_i v_i$ for their usual inner product.

We can thus extend the principal's payoff in (2.3) to random contracts by setting

$$V(p|A) := \sum_{w \in \text{supp}(p)} p(w) V(w|A), \tag{2.5}$$

where supp(p) is the support of p, a finite set. Similarly, we extend the guarantee (2.4) to random contracts by setting

$$V(p) := \inf_{A \supseteq A^0} V(p|A). \tag{2.6}$$

The principal's *optimal guarantee* is the supremum of the guarantee (2.6) over all random contracts, or $\sup_{p \in \Delta(\mathbb{R}^n_+)} V(p)$. A random contract is *optimal* if it attains this optimal guarantee. In contrast, the *optimal deterministic guarantee* is the supremum of (2.4) over all contracts, or equivalently, the supremum of (2.6) over all degenerate random contracts (i.e., contracts whose support is a singleton). A contract w is an *optimal deterministic contract* if it attains the optimal deterministic guarantee.

We assume throughout that the known technology A^0 contains a productive action, i.e., there exists $(\pi, c) \in A^0$ such that $\pi y - c > 0$. This non-triviality assumption ensures that the optimal guarantee is positive.

Some remarks are in order regarding the formulation of the problem:

1. Because we assume that the agent observes the realized contract before choosing an action, the agent need not know the underlying random contract, or even be aware that the contract was generated via randomization. Indeed, because of bilateral risk-neutrality, giving the agent a random contract would not be useful: Given any action (π, c), the principal's and the agent's payoffs from a randomized contract p would be $\pi(y - \sum_{w \in \text{supp}(p)} p(w)w)$ and $\pi \sum_{w \in \text{supp}(p)} p(w)w - c$, and thus we could equivalently use the deterministic contract $\tilde{w}$ with $\tilde{w}_i = \sum_{w \in \text{supp}(p)} p(w)w_i$ for all $i \in I$.

2. Our contracting environment is the same as in Carroll (2015), save for the following minor differences. First, we take the set of outputs, Y, to be finite (rather than just compact) to minimize technicalities. This is inessential for the economic logic of the results, but it allows for simple duality-based proofs. (Any compact continuous output set can of course be approximated with a finite grid.) Second, we also assume a technology to be a finite set; this is only used to simplify the handling of a corner case in the proofs of Propositions 6 and 7. Third, we assume adversarial tie-breaking in (2.3), whereas Carroll broke ties in the principal's favor. Adversarial tie-breaking could be argued to better capture the spirit of the robustness exercise, and it simplifies the analysis as the infimum in (2.6) becomes a minimum. However, both assumptions lead

to the same optimal guarantee, and, as we will see, the existence of optimal contracts is similarly unaffected by tie-breaking except for an uninteresting corner case.

3. We define a random contract as a finitely supported probability over deterministic contracts. This assumption is without loss of generality in the sense that any randomization actually implementable on a computer would have to be finite.

3 The Deterministic Case

The duality approach we adopt to prove the optimality of random linear contracts also yields a short proof of the optimality of linear contracts within the class of deterministic contracts. It is instructive to see the argument first in this simpler case.

The basic idea is to formulate the guarantee (2.4) of a contract w as a linear program. To this end, note that, given any technology $A \supseteq A^0$, only the action chosen by the agent matters for the principal's payoff $V(w|A)$ in (2.3).[4] It thus suffices to take the infimum in (2.4) over technologies that add at most one new action to the known technology A^0. It follows that the contract's guarantee is characterized by the following linear program:

$$V(w) = \min_{\pi,c} \pi(y-w) \tag{3.1}$$

$$\text{s.t.} \quad \pi w - c \geq U^0(w), \tag{3.2}$$

$$\sum_{i=1}^{n} \pi_i = 1, \tag{3.3}$$

$$c, \pi_i \geq 0 \quad \forall i \in I. \tag{3.4}$$

That is, Nature designs an action (π, c) to minimize the principal's profit, with constraint (3.2) ensuring that (π, c) is a best response for the agent since, by (2.2), $U^0(w)$ is the agent's maximum payoff from the known technology A^0. Problem (3.1–3.4) is feasible because taking (π, c) to be an action in A^0 that attains $U^0(w)$ satisfies all constraints.

A contract w is *linear* if $w = \alpha y$ for some slope $\alpha \in [0,1]$. With slight abuse of notation, we will identify a linear contract with its slope α. We can now show that any contract can be weakly improved upon by some linear contract.

Proposition 1. *For every contract w, there is a linear contract α such that $V(w) \leq V(\alpha)$.*

Proof. Fix a contract w. The guarantee $V(w)$ is clearly bounded from below by $-\max_i w_i$.

[4] I.e., if (π^*, c^*) attains the minimum in (2.3) given w and $A \supseteq A^0$, then $V(w|A) = V(w|A^0 \cup \{(\pi^*, c^*)\})$.

By strong duality, the following dual to problem (3.1–3.4) is then feasible and bounded:

$$V(w) = \max_{\lambda,\mu} \lambda U^0(w) + \mu \tag{3.5}$$

$$\text{s.t.} \quad \lambda w_i + \mu \leq y_i - w_i \quad \forall i \in I, \tag{3.6}$$

$$\lambda \geq 0. \tag{3.7}$$

Let (λ^*, μ^*) be an optimal solution. Evaluating (3.6) at $i = 1$ gives $\mu^* \leq -(1+\lambda^*)w_1 \leq 0$ as $y_1 = 0$. Define a new contract w' as follows. For each i, choose w'_i to satisfy (3.6) as an equality, i.e., let

$$w'_i := \frac{y_i - \mu^*}{\lambda^* + 1}. \tag{3.8}$$

By inspection, w' is an affine function of output. As the original contract w also satisfies (3.6), we have $w' \geq w \geq 0$, and thus w' is a well-defined contract. Moreover, we have

$$V(w') \geq \lambda^* U^0(w') + \mu^* \geq \lambda^* U^0(w) + \mu^* = V(w), \tag{3.9}$$

where the first inequality follows because (λ^*, μ^*) is by construction of w' still feasible when w is replaced with w' in the dual, and the second inequality follows because $w' \geq w$ and $U^0(w)$ is weakly increasing in w by inspection of (2.2). We conclude that the original contract w is weakly dominated by the affine contract w'.

To get a linear contract, we may drop the lump-sum payment $-\mu^*/(\lambda^* + 1) \geq 0$ from w' by letting $w'' := y/(\lambda^*+1)$. This leaves the agent's choice from every technology unchanged, but weakly increases the principal's payoff. Thus, $V(w'') \geq V(w') \geq V(w)$. □

A reader familiar with Carroll's (2015) proof will see the parallels between the arguments, with duality here replacing the separating hyperplane theorem. Given the equivalence of linear programming duality and the separating hyperplane theorem, at a deeper level the proofs are the same. However, adopting a linear programming perspective seems to allow a somewhat more concise argument.[5]

[5] The proof of Proposition 1 extends mutatis mutandis to any compact set $Y \subset \mathbb{R}_+$ such that $\min Y = 0$, with a contract then defined to be a continuous function $w : Y \to \mathbb{R}_+$. The primal problem (3.1–3.4) then becomes one of choosing a regular nonnegative Borel measure π on Y and a cost $c \in \mathbb{R}$ to solve

$$V(w) = \min \int_Y (y - w(y))\, d\pi(y) \ \text{ s.t. } \int_Y w(y)\, d\pi(y) - c \geq U^0(w), \ \int_Y 1\, d\pi(y) = 1, \ \text{and } c \geq 0.$$

The dual problem (3.5–3.7) in turn becomes one of choosing numbers $\lambda \in \mathbb{R}$ and $\mu \in \mathbb{R}$ to

$$\max \ \lambda U^0(w) + \mu \ \text{ s.t. } \lambda w(y) + \mu \leq y - w(y) \ \forall y \in Y, \ \text{and } \lambda \geq 0.$$

As the zero contract has $V(0) \geq 0$, we may assume that $V(w) > 0$. By weak duality, this implies that the

Proposition 1 implies that linear contracts are optimal, in the following sense:

Corollary 1. *The optimal deterministic guarantee satisfies* $\sup_{w\in\mathbb{R}^n_+} V(w) = \sup_{\alpha\in[0,1]} V(\alpha)$.

Proof. Clearly $\bar{v} := \sup_{w\in\mathbb{R}^n_+} V(w) \geq \sup_{\alpha\in[0,1]} V(\alpha)$. For the converse, let w_n be a sequence such that $V(w_n) \to \bar{v}$. By Proposition 1, there is a sequence of linear contracts α_n such that $V(w_n) \leq V(\alpha_n) \leq \sup_{m\geq n} V(\alpha_m)$ for all n. Hence, $\bar{v} \leq \limsup_n V(\alpha_n) \leq \sup_{\alpha\in[0,1]} V(\alpha)$. $\square$

Having established the optimality of linear contracts as a class, it only remains to characterize the optimal deterministic guarantee and the optimal linear contracts. The remaining steps are specific to the deterministic case, and so we develop them here only to the extent needed for the subsequent analysis. They also serve to illustrate that adversarial and principal-optimal tie-breaking yield essentially the same results.

To this end, consider Nature's problem (3.1–3.4) given a linear contract with slope α:

$$V(\alpha) = \min_{\pi,c} (1-\alpha)\pi y \quad \text{s.t.} \quad \alpha\pi y - c \geq U^0(\alpha), \quad \sum_{i=1}^{n} \pi_i = 1, \quad \text{and} \quad c, \pi_i \geq 0 \quad \forall i \in I.$$

It is clearly optimal to set $c = 0$ and choose π such that $\alpha\pi y = U^0(\alpha)^+$, where our convention is to write $b^+ = \max\{b, 0\}$ for any $b \in \mathbb{R}$. If $\alpha = 0$, then it is optimal to put $\pi y = U^0(0)^+ = 0$, and thus the guarantee from the zero contract is zero. (This is of course immediate given adversarial tie-breaking.) On the other hand, the guarantee for any positive slope $\alpha > 0$ is $V(\alpha) = (1-\alpha)U^0(\alpha)^+/\alpha$. Recalling the definition of U^0 from (2.2), we thus have

$$V(\alpha) = \begin{cases} 0 & \text{if } \alpha = 0, \\ (1-\alpha)\max_{(\pi,c)\in A^0}\left(\pi y - \frac{c}{\alpha}\right)^+ & \text{if } \alpha \in (0,1]. \end{cases} \tag{3.10}$$

By Corollary 1, the optimal deterministic guarantee is given by the supremum of the function $V : [0,1] \to \mathbb{R}_+$ defined by (3.10). The non-triviality assumption implies that $V(\alpha) > 0$ for α sufficiently close to 1, and thus the optimal guarantee is positive. It can also readily be verified to coincide with Carroll's (2015) guarantee for principal-optimal tie-breaking.[6] In what follows, we will use the fact that the optimal deterministic guarantee can also be written as

$$\sup_{w\in\mathbb{R}^n_+} V(w) = \max_{\alpha\in[0,1]} \frac{1-\alpha}{\alpha} U^0(\alpha), \tag{3.11}$$

---

dual is also bounded. Moreover, if we let $\lambda' = 1$ and $\mu' = \min\{y - 2w(y) : y \in Y\} - 1$, then (λ', μ') satisfies all dual constraints as strict inequalities. Therefore, strong duality holds for this pair of semi-infinite linear programs (see, e.g., Lai and Wu, 1992, Theorem 2.1). The rest of the argument now proceeds as above, with the affine improvement contract constructed analogously to (3.8) given some optimal dual solution.

[6] Principal-optimal tie-breaking leads in general to a weakly higher guarantee from the zero contract than adversarial tie-breaking, but the guarantees are the same for any positive slope. As the guarantee under principal-optimal tie-breaking is continuous in α, this implies that the optimal guarantees are the same.

provided that $\frac{(1-0)}{0}U^0(0)$ is taken to mean the limit as $\alpha \to 0$.

An optimal deterministic contract, which is linear, can be found by simply maximizing the guarantee (3.10). However, there is a small wrinkle as V, which is continuous on $(0,1]$, may have a downward jump in the limit as $\alpha \to 0$. This leads to the following characterization.[7]

Proposition 2. *A linear contract α is an optimal deterministic contract if and only if it maximizes the function $V : [0,1] \to \mathbb{R}_+$ defined by* (3.10)*. If no such maximizer exists, then there does not exist an optimal deterministic contract. In this case, the optimal deterministic guarantee is attained in the limit of any sequence of linear contracts $\alpha_n > 0$ such that $\alpha_n \to 0$.*

Proof. The first claim characterizing optimal linear contracts follows by Corollary 1 and the construction of V. That there is no optimal deterministic contract when there is no linear optimal contract follows by Proposition 1. Finally, the only possible point of discontinuity of V is 0, and hence if a maximizer does not exist, then $\sup_{\alpha\in[0,1]} V(\alpha) = \lim_{\alpha\to 0^+} V(\alpha)$. $\square$

While adversarial tie-breaking generates an existence problem relative to the case of principal-optimal tie-breaking where an optimal (linear) contract always exists, the issue is arguably minor: It only arises in the corner case where under principal-optimal tie-breaking, the zero contract is the uniquely optimal linear contract. This case seems uninteresting from the perspective of optimal incentive provision as none is required. It is incidentally also the only case where randomization does not improve the principal's payoff—see Corollary 3. A simple sufficient condition to rule it out and to ensure the existence of an optimal contract is for any known action generating a positive surplus to have a non-zero cost.

4 Optimality of Random Linear Contracts

We now turn to our main contribution—showing that the class of random linear contracts is optimal in the space of random contracts.

We say that a random contract $p \in \Delta(\mathbb{R}^n_+)$ is *linear*, or a *random linear contract*, if every contract in the support of p is linear, i.e., if for all $w \in \mathrm{supp}(p)$, there exists a slope $\alpha \in [0,1]$ such that $w = \alpha y$. Whenever convenient, we identify each (deterministic) linear contract with its slope, and then identify the set of all random linear contracts with $\Delta([0,1])$, the space of finitely supported probability measures on the unit interval. With slight abuse of notation, we continue to use V to denote the principal's guarantee over any of these spaces.

The following results generalize Proposition 1 and Corollary 1 to random contracts.

[7] It is easy to show that all optimal deterministic contracts are linear under Carroll's (2015) full support condition by verifying that then certain inequalities in the proof of Proposition 1 are strict. As this argument is similar to Carroll's, we omit it in the interest of space.

Proposition 3. *For every random contract p, there is a random linear contract q such that*

$$V(p) \leq V(q) \text{ and } |\operatorname{supp}(p)| \geq |\operatorname{supp}(q)|.$$

Corollary 2. *The optimal guarantee satisfies* $\sup_{p\in\Delta(\mathbb{R}^n_+)} V(p) = \sup_{p\in\Delta([0,1])} V(p)$.

Proof. Same as Corollary 1, mutatis mutandis. □

We note that because of the conclusion regarding supports, Proposition 3 nests Proposition 1 as a special case by taking p to be a degenerate random contract. More generally, Proposition 3 implies that random linear contracts remain optimal given any upper bound on the size of a contract's support (e.g., because of complexity considerations).

In order to establish Proposition 3, the first step is to formulate the guarantee (2.6) as a linear program. Let p be a random contract. Enumerate the contracts in the support of p so that $\operatorname{supp}(p) = \{w^1, \dots, w^k\}$ and denote the index set by $T := \{1, \dots, k\}$. We will write $p_t := p(w^t)$ for the probability of contract t. In searching for the worst-case technology against p, it suffices to consider technologies $A \supseteq A^0$ that have at most k new actions, one for each possible realized contract.[8] The guarantee $V(p)$ is thus characterized by the following finite linear program, which we refer to as the primal problem:

$$V(p) = \min_{\{(\pi^t, c_t)\}} \sum_{t\in T} p_t \pi^t (y - w^t) \tag{4.1}$$

$$\text{s.t.} \quad \pi^t w^t - c_t \geq \pi^s w^t - c_s \qquad \forall t, s \in T : t \neq s, \qquad (\kappa_{ts}) \tag{4.2}$$

$$\pi^t w^t - c_t \geq U^0(w^t) \qquad \forall t \in T, \qquad (\lambda_t) \tag{4.3}$$

$$\sum_{i\in I} \pi^t_i = 1 \qquad \forall t \in T, \qquad (\mu_t) \tag{4.4}$$

$$c_t, \pi^t_i \geq 0 \qquad \forall i \in I,\ t \in T. \tag{4.5}$$

That is, Nature designs, for each contract w^t in the support of p, an action (π^t, c_t) the agent will take if contract w^t is realized. This problem can be interpreted as a multi-dimensional quasilinear mechanism design problem where the agent's type t corresponds to a contract $w^t \in \mathbb{R}^n_+$, the allocation is an output distribution $\pi \in \Delta(Y)$, and the transfer is a

[8] To see this, given any technology $A \supseteq A^0$, let $a^t := (\pi^t, c_t)$ denote the action the agent chooses from A if contract w^t is realized. (That is, a^t attains $V(w^t|A)$.) It is straightforward to verify that, for all $t \in T$, we have $V(w^t|A) = V(w^t|A^0 \cup \{a^1, \dots, a^k\})$. This implies that

$$V(p|A) = \sum_{t\in T} p_t V(w^t|A) = \sum_{t\in T} p_t V(w^t|A^0 \cup \{a^1, \dots, a^k\}) = V(p|A^0 \cup \{a^1, \dots, a^k\}).$$

Therefore, for any $A \supseteq A^0$, there exists a technology $A' \supseteq A^0$ with $|A' \setminus A^0| \leq k$ such that $V(p|A') = V(p|A)$.

cost c, constrained to be nonnegative. The constraint (4.3) is a participation constraint that ensures that each type t prefers its own action (π^t, c_t) to any of the known actions in A^0. (Note that the utility from this outside option depends on the type t.) The new constraint (4.2), absent from Nature's problem studied in the deterministic case, is an incentive compatibility constraint that ensures that each type t prefers its own action (π^t, c_t) to the action (π^s, c_s) of every other type s.

Problem (4.1–4.5) is feasible because taking each (π^t, c_t) to be an action in A^0 that attains $U^0(w^t)$ satisfies all constraints. (Put differently, Nature can always just give the agent the known technology A^0.) Moreover, its value is clearly bounded from below by $-\max_{i,t} w_i^t$, and thus an optimal solution exists.

As in the deterministic case, we will study the dual to Nature's problem. Strong duality implies that the following dual is feasible and bounded, with $V(p)$ the optimal value:

$$V(p) = \max_{\kappa,\lambda,\mu} \sum_{t \in T} \lambda_t U^0(w^t) + \sum_{t \in T} \mu_t \tag{4.6}$$

$$\text{s.t.} \quad \lambda_t w_i^t + \mu_t + \sum_{s \neq t} \kappa_{ts} w_i^t - \sum_{s \neq t} \kappa_{st} w_i^s \leq p_t(y_i - w_i^t) \qquad \forall i \in I, t \in T, \qquad (\pi_i^t) \tag{4.7}$$

$$\lambda_t + \sum_{s \neq t} \kappa_{ts} - \sum_{s \neq t} \kappa_{st} \geq 0 \qquad \forall t \in T, \qquad (c_t) \tag{4.8}$$

$$\kappa_{ts}, \lambda_t \geq 0 \qquad \forall t, s \in T : t \neq s. \tag{4.9}$$

We will use *dual solution* to refer to any vector $(\kappa, \lambda, \mu) \in \mathbb{R}^{T(T-1)} \times \mathbb{R}^T \times \mathbb{R}^T$ satisfying constraints (4.8) and (4.9), and use the qualifiers *feasible* or *optimal* to indicate, respectively, that the dual solution is feasible or optimal in (4.6–4.9) for a given random contract p.

In order to make duality between problems (4.1–4.5) and (4.6–4.9) easier to verify, each non-trivial primal constraint in (4.1–4.5) has the associated dual variable displayed next to it in parenthesis. Similarly, the associated primal variable is displayed next to each non-trivial dual constraint in (4.6–4.9). It may also be instructive to compare the above dual problem to the dual (3.5–3.7) from the deterministic case; by inspection, the problems coincide if the random contract p is degenerate so that T is a singleton.

The general idea in the proof of Proposition 3 is analogous to the proof of Proposition 1 in the deterministic case: Given a random contract p, we take an optimal dual solution $(\kappa^*, \lambda^*, \mu^*)$ and use constraint (4.7) to construct, for each contract w^t in the support of p, an affine contract that dominates it from above. It will then be shown that a randomization over these affine contracts improves on p, and that a further improvement is obtained by a random linear contract.

We will need a preliminary result about the structure of optimal dual solutions. Given a dual solution (κ, λ, μ), let $G(\kappa)$ be a directed graph with vertex set $T = \{1, \ldots, k\}$ and an arc directed from t to s whenever $\kappa_{ts} > 0$. Note that if (κ, λ, μ) is an optimal solution to (4.6–4.9), then by complementary slackness, each arc in $G(\kappa)$ corresponds to a binding incentive compatibility constraint in (4.2) in the primal problem.

Lemma 1. *There is an optimal dual solution $(\kappa^*, \lambda^*, \mu^*)$ for which the graph $G(\kappa^*)$ is acyclic.*

The proof of Lemma 1 is lengthy, so we only sketch the proof for a special case to illustrate the idea and relegate the full argument to the Appendix.

Proof of Lemma 1 for two equi-probable contracts. Suppose $\mathrm{supp}(p) = \{w^1, w^2\}$ and $p_1 = p_2 = 1/2$. Consider a perturbed version of the primal problem (4.1–4.5) where the incentive constraints in (4.2) are relaxed by some $\varepsilon > 0$:

$$\pi^1 w^1 - c_1 \geq \pi^2 w^1 - c_2 - \varepsilon \quad \text{and} \quad \pi^2 w^2 - c_2 \geq \pi^1 w^2 - c_1 - \varepsilon. \tag{4.10}$$

As ε only enters the right-hand sides of these constraints, the perturbed dual is otherwise identical to (4.6–4.9) but the objective in (4.6) has the additional term $-(\kappa_{12} + \kappa_{21})\varepsilon$.

Suppose toward contradiction that $(\kappa^*, \lambda^*, \mu^*)$ is an optimal solution to the perturbed dual and $G(\kappa^*)$ has a cycle. That is, $\kappa^*_{12} > 0$ and $\kappa^*_{21} > 0$. By complementary slackness, both constraints in (4.10) hold as equalities at an optimal solution $\{(\pi^1, c_1), (\pi^2, c_2)\}$ to the perturbed primal. We observe that then type 1 strictly prefers (π^2, c_2) to (π^1, c_1), and conversely for type 2.

Construct a new primal solution by swapping the actions, i.e., let $(\hat{\pi}^1, \hat{c}_1) := (\pi^2, c_2)$ and $(\hat{\pi}^2, \hat{c}_2) := (\pi^1, c_1)$. This new solution is clearly feasible: It trivially satisfies the probability and non-negativity constraints (4.4) and (4.5). And because each type is now strictly better off, $\{(\hat{\pi}^1, \hat{c}_1), (\hat{\pi}^2, \hat{c}_2)\}$ satisfies both the incentive compatibility constraints (4.10) and the participation constraints (4.3). Moreover, the new solution attains a strictly lower primal objective value as total surplus is unchanged, but more of it now goes to the agent:

$$\begin{aligned} 0.5[\hat{\pi}^1(y - w^1) + \hat{\pi}^2(y - w^2)] &= 0.5[\pi^2(y - w^1) + \pi^1(y - w^2)] \\ &= 0.5[\pi^2 y - \pi^1 w^1 - c_2 + c_1 - \varepsilon + \pi^1 y - p^2 w^2 - c_1 + c_2 - \varepsilon] \\ &= 0.5[\pi^1(y - w^1) + \pi^2(y - w^2)] - \varepsilon, \end{aligned}$$

where the second equality follows because the old solution $\{(\pi^1, c_1), (\pi^2, c_2)\}$ satisfies (4.10) with equality. Thus, the old solution is in fact not optimal, a contradiction.

By the above argument, every optimal solution to the perturbed dual has an acyclic graph. Upper hemi-continuity of optimal dual solutions in the perturbation ε then allows us to establish the existence of an optimal solution with an acyclic graph to the unperturbed dual via a limit argument. See the Appendix for the details and the general case. □

The reason we resort to the perturbation argument to show Lemma 1 is that in the original problem, the incentive compatibility constraints that hold with equality at the optimum may well form a cycle. For example, this is the case if contracts 1 and 2 are pooled to the same action. In contrast, such a cycle is impossible in the perturbed problem as illustrated by the above proof sketch. By complementary slackness, the graph $G(\kappa^*)$ is thus always acyclic in the perturbed problem. As we send the perturbation parameter ε to zero, we uncover the existence of an acyclic graph for the original problem as well. In the case where the optimal primal solution pools contracts 1 and 2 to the same action, this implies that even if both incentive compatibility constraints then hold with equality, there is an optimal dual solution where at most one of them has a positive shadow price.

The next two lemmas establish the existence of affine contracts that dominate the contracts in the support of p. In showing this, we will view the dual constraint (4.7) as a system of linear inequalities in the contracts $\{w^1, \dots, w^k\}$, in the following sense.

Definition 1. Given a dual solution (κ, λ, μ) and probabilities $(p_1, \dots, p_k) \in [0,1]^k$, the *$w$-system* is the system of linear inequalities in $(w^1, \dots, w^k) \in \mathbb{R}^{kn}_+$ defined by (4.7).

Given a dual solution (κ, λ, μ), define the set of *in-neighbors* of $t \in T$ in the graph $G(\kappa)$ as $N(t) := \{s \in T : \kappa_{st} > 0\}$, and the set of *out-neighbors* as $O(t) := \{s \in T : \kappa_{ts} > 0\}$. We say that contract $w \in \mathbb{R}^n_+$ is *positive affine* if $w_i = \alpha y_i + \beta$ for all $i \in I$ for some $\alpha \in (0,1]$ and $\beta \geq 0$ independent of i. (Because $y_1 = 0$, every affine contract has $\beta \geq 0$, so being positive affine is a restriction on α.)

Lemma 2. *Let $(w^1, \dots, w^k) \in \mathbb{R}^{kn}_+$ be feasible in the w-system given (κ, λ, μ) and $(p_1, \dots, p_k)$. If all in-neighbors of a contract $t \in T$ are positive affine, i.e., if the contract w^s in $(w^1, \dots, w^k)$ is positive affine for all $s \in N(t)$, then there is a positive affine contract x^t satisfying the following properties:*

1. *$x^t \geq w^t$,*
2. *(x^t, w^{-t}) is feasible in the w-system.*[9]

Proof. Let $t \in T$ be a contract whose every in-neighbor $s \in N(t)$ is a positive affine contract. The (i,t)-instance of constraint (4.7) then takes the form

$$\lambda_t w_i^t + \mu_t + \sum_{s \in O(t)} \kappa_{ts} w_i^t - \sum_{s \in N(t)} \kappa_{st}(\alpha_s y_i + \beta_s) \leq p_t(y_i - w_i^t),$$

[9] We use the standard shorthand $(x^t, w^{-t}) := (w^1, \dots, w^{t-1}, x^t, w^{t+1}, \dots, w^k)$.

or, equivalently,

$$w_i^t \leq \frac{p_t + \sum_{s\in N(t)} \kappa_{st}\alpha_s}{\lambda_t + p_t + \sum_{s\in O(t)} \kappa_{ts}} y_i + \frac{\sum_{s\in N(t)} \kappa_{st}\beta_s - \mu_t}{\lambda_t + p_t + \sum_{s\in O(t)} \kappa_{ts}} =: x_i^t. \tag{4.11}$$

As i ranges over I, the right-hand side defines a vector x^t. By construction, $x^t \geq w^t \geq 0$, and thus x^t is a contract. It is affine by inspection. To verify that x^t is positive affine, note that

$$0 < \frac{p_t + \sum_{s\in N(t)} \kappa_{st}\alpha_s}{\lambda_t + p_t + \sum_{s\in O(t)} \kappa_{ts}} \leq \frac{p_t + \sum_{s\in N(t)} \kappa_{st}}{\lambda_t + p_t + \sum_{s\in O(t)} \kappa_{ts}} \leq \frac{p_t}{p_t} = 1,$$

where the first inequality follows because $p_t, \alpha_s > 0$ and $\kappa_{ts}, \kappa_{st}, \lambda_t \geq 0$, the second follows because $\alpha_s \leq 1$ as every in-neighbor s is positive affine, and the third inequality follows because the dual solution (κ, λ, μ) satisfies (4.8) by definition.

For property 2, note that (x^t, w^{-t}) satisfies the (i,t)-instance of (4.7) with equality for all $i \in I$ by construction of x^t. Moreover, because $x^t \geq w^t$ and $\kappa_{st} \geq 0$ for all $s \neq t$, replacing w^t with x^t weakly relaxes all other instances of (4.7). □

Lemma 3. *Let $(w^1, \ldots, w^k) \in \mathbb{R}_+^{kn}$ be feasible in the w-system given (κ, λ, μ) and $(p_1, \ldots, p_k)$. If the graph $G(\kappa)$ is acyclic, then there exists a vector of positive affine contracts $(x^1, \ldots, x^k)$ that is feasible in the w-system and satisfies $x^t \geq w^t$ for every $t \in T$.*

Proof. Let $(w^1, \ldots, w^k) \in \mathbb{R}_+^{kn}$ be feasible in the w-system and suppose $G(\kappa)$ is acyclic. Partition T recursively as follows:

- Let $L_0 := \{t \in T : N(t) = \emptyset\}$.
- For $\ell \geq 1$, let $L_\ell := \{t \in T \setminus \cup_{m<\ell} L_m : N(t) \subseteq \bigcup_{m<\ell} L_m\}$.

That is, the zero layer L_0 consists of all contracts that have no in-neighbors in the graph $G(\kappa)$. For $\ell > 0$, the ℓ-th layer L_ℓ consists of all contracts not contained in any of the lower layers L_m, $m < \ell$, but whose in-neighbors are in these layers. Because $G(\kappa)$ is acyclic, it is straightforward to verify that there exists $\bar{\ell} \in \mathbb{N}_0$ such that L_ℓ is nonempty if and only if $0 \leq \ell \leq \bar{\ell}$, and that $\{L_1, \ldots, L_{\bar{\ell}}\}$ is a partition of T.

By relabeling if necessary, we can assume without loss of generality that the contracts $(w^1, \ldots, w^k)$ are labeled monotonically so that contracts with higher indices are in higher layers, in the following sense: If $t < t'$, $t \in L_\ell$, and $t' \in L_{\ell'}$, then $\ell \leq \ell'$.

We now construct the positive affine contracts $(x^1, \ldots, x^k) \in \mathbb{R}_+^{kn}$ by induction on t.

Base case: Because the contracts are labeled monotonically, we have $1 \in L_0$ and thus contract 1 has no in-neighbors. Moreover, $(w^1, \ldots, w^k)$ is feasible in the w-system given

(κ, λ, μ) and $(p_1, \ldots, p_k)$ by assumption. Thus, by Lemma 2, there is a positive affine contract $x^1 \geq w^1$ such that $(x^1, w^2, \ldots, w^k)$ is feasible in the w-system.

Induction step: Let $1 < t \leq k$. Suppose there exist positive affine contracts $x^1, \ldots, x^{t-1}$ with $x^s \geq w^s$ for all $1 \leq s \leq t-1$, and $(x^1, \ldots, x^{t-1}, w^t, \ldots, w^k)$ is feasible in the w-system. Let L_ℓ be the layer containing contract t. By definition of L_ℓ, we have $N(t) \subseteq \cup_{m<\ell} L_m$. The induction hypothesis and monotonicity of labeling imply that t's in-neighbors are positive affine. (Note that this subsumes the special case $\ell = 0$, in which case $N(t) \subseteq \cup_{m<0} L_m = \emptyset$ and the conclusion holds vacuously.) Thus, Lemma 2 again gives us a positive affine contract $x^t \geq w^t$ such that $(x^1, \ldots, x^t, w^{t+1}, \ldots, w^k)$ is feasible in the w-system. $\square$

With the above results in hand, we are now ready to prove Proposition 3.

Proof of Proposition 3. Let p be a random contract with $\operatorname{supp}(p) = \{w^1, \ldots, w^k\}$ and let $(\kappa^*, \lambda^*, \mu^*)$ be an optimal solution to the dual (4.6–4.9) such that $G(\kappa^*)$ is acyclic, the existence of which is ensured by Lemma 1.

Because $G(\kappa^*)$ is acyclic and $(w^1, \ldots, w^k)$ is obviously feasible in the w-system given $(\kappa^*, \lambda^*, \mu^*)$ and $(p_1, \ldots, p_k)$, Lemma 3 gives us a vector of positive affine contracts $(x^1, \ldots, x^k)$ that is feasible in the w-system and satisfies $x^t \geq w^t$ for all $t \in T$.

Consider now the dual (4.6–4.9) where we replace the contracts $(w^1, \ldots, w^k)$ with the positive affine contracts $(x^1, \ldots, x^k)$ while keeping the probabilities $(p_1, \ldots, p_k)$ unchanged:

$$
\begin{aligned}
V(p_1, \ldots, p_k; x^1, \ldots, x^k) := \max_{\kappa,\lambda,\mu} & \sum_{t \in T} \lambda_t U^0(x^t) + \sum_{t \in T} \mu_t \\
\text{s.t.} \quad & \lambda_t x_i^t + \mu_t + \sum_{s \neq t} \kappa_{ts} x_i^t - \sum_{s \neq t} \kappa_{st} x_i^s \leq p_t(y_i - x_i^t) && \forall i \in I, t \in T, \\
& \lambda_t + \sum_{s \neq t} \kappa_{ts} - \sum_{s \neq t} \kappa_{st} \geq 0 && \forall t \in T, \\
& \kappa_{ts}, \lambda_t \geq 0 && \forall t, s \in T : t \neq s.
\end{aligned}
$$

Because $(x^1, \ldots, x^k)$ is feasible in the w-system given $(\kappa^*, \lambda^*, \mu^*)$ and $(p_1, \ldots, p_k)$, the dual solution $(\kappa^*, \lambda^*, \mu^*)$ is feasible in the above problem. Therefore,

$$
V(p_1, \ldots, p_k; x^1, \ldots, x^k) \geq \sum_{t \in T} \lambda_t^* U^0(x^t) + \sum_{t \in T} \mu_t^* \geq \sum_{t \in T} \lambda_t^* U^0(w^t) + \sum_{t \in T} \mu_t^* = V(p), \quad (4.12)
$$

where the second inequality follows because $x^t \geq w^t$ for all t and U^0 is weakly increasing in the contract (in the pointwise order) by inspection of (2.2).

Moreover, strong duality implies that $V(p_1, \ldots, p_k; x^1, \ldots, x^k)$ is the optimal value of the primal problem (4.1–4.5), which using the positive affine form $x_i^t = \alpha_t y_i + \beta_t$ $(i \in I)$ of each contract t now becomes

$$V(p_1,\dots,p_k;x^1\dots,x^k)=\min_{\{(\pi^t,c_t)\}}\sum_{t\in T}p_t(1-\alpha_t)\pi^t y-\sum_{t\in T}p_t\beta_t \tag{4.13}$$

$$\text{s.t.}\quad \alpha_t\pi^t y-c_t\geq\alpha_t\pi^s y-c_s \qquad \forall t,s\in T:t\neq s, \tag{4.14}$$

$$\alpha_t\pi^t y-c_t\geq U^0(\alpha_t) \qquad \forall t\in T, \tag{4.15}$$

$$\sum_{i\in I}\pi_i^t=1 \qquad \forall t\in T, \tag{4.16}$$

$$c_t,\pi_i^t\geq 0 \qquad \forall i\in I,\, t\in T. \tag{4.17}$$

Note that the feasible set is independent of the constants $\beta_t\geq 0$, $t\in T$. (They enter both sides of the incentive compatibility and participation constraints (4.14) and (4.15), and cancel out.) By inspection of (4.13–4.17), replacing each x^t with the linear contract α_t gives

$$V(p_1,\dots,p_k;\alpha_1,\dots,\alpha_k)=V(p_1,\dots,p_k;x^1,\dots,x^k)+\sum_{t\in T}p_t\beta_t\geq V(p_1,\dots,p_k;x^1,\dots,x^k),$$

which together with (4.12) implies $V(p_1,\dots,p_k;\alpha_1,\dots,\alpha_k)\geq V(p)$.

To conclude the proof, define the random linear contract q by setting $q(\alpha)=\sum_{t:\alpha_t=\alpha}p_t$ for all $\alpha\in[0,1]$, with the sum over the empty set equal to zero, as usual.[10] It is then immediate that $|\mathrm{supp}(q)|\leq|\mathrm{supp}(p)|$. Furthermore, $V(q)$ is the optimal value of the problem (4.13–4.17) with $\beta_t\equiv 0$ for all t and with the additional constraint that $(\pi^t,c_t)=(\pi^s,c_s)$ for all $s,t\in T$ such that $\alpha_s=\alpha_t$. Therefore, $V(q)\geq V(p_1,\dots,p_k;\alpha_1,\dots,\alpha_k)\geq V(p)$. □

Proposition 3 shows that every random contract p is weakly dominated by a random linear contract. The proof shows, essentially, that p can be weakly improved by linearizing each contract in its support while keeping the probabilities fixed, thus generalizing the proof of Proposition 1 from the deterministic case. The new challenge in the random case is the presence of the incentive compatibility constraints (4.2), which create interlinkages between contracts as the action (π^t,c_t) Nature targets to contract w^t will be available to the agent also when any other contract is realized. (This is why randomization is useful in the first place.) Because the primal (4.1–4.5) is a multi-dimensional screening problem, the structure of binding incentive compatibility constraints (i.e., those with a positive shadow price) is not clear a priori. Our proof deals with this by showing that there nevertheless is an optimal solution under which the binding constraints do not form a cycle. This is enough to then use the dual to Nature's problem to construct the improvement contract.

[10] If the list $\alpha_1,\dots,\alpha_k$ includes multiple copies of some slope, then the tuple $(p_1\dots,p_k;\alpha_1,\dots,\alpha_k)$ is itself not a random linear contract according to our definition as it is not a finitely supported probability on $[0,1]$.

5 Optimal Randomization

Having established the optimality of random linear contracts as a class, we now characterize optimal random linear contracts by means of a maximizing sequence of such contracts that attains the optimal guarantee. In the process, we obtain a formula for the optimal guarantee. We also show that the maximizing sequence converges to a distribution over linear contracts that puts positive density on an interval of slopes. While it is obviously impossible to exactly implement this limiting "continuous random linear contract" in practice, it nevertheless provides a concise way to summarize the maximizing sequence.

5.1 Guarantee of a Random Linear Contract

We start by specializing Nature's problem to random linear contracts. This brings about a considerable simplification as it renders the mechanism design problem one-dimensional.

We passed through Nature's problem against a randomization over linear contracts in the proof of Proposition 3. Specifically, setting $\beta_t \equiv 0$ for all $t \in T$ in problem (4.13–4.17) gives the primal problem for a random linear contract p with $\operatorname{supp}(p) = \{\alpha_1, \ldots, \alpha_k\}$. By inspection, the problem depends on each output distribution π^t only via the expected output $e_t := \pi^t y \in [0, y_n]$. With this change of variables, we can write the primal problem as follows:

$$V(p) = \min_{\{(e_t, c_t)\}} \sum_{t \in T} p_t (1 - \alpha_t) e_t \tag{5.1}$$

$$\text{s.t.} \quad \alpha_t e_t - c_t \geq \alpha_t e_s - c_s \qquad \forall t, s \in T : t \neq s, \tag{5.2}$$

$$\alpha_t e_t - c_t \geq U^0(\alpha_t) \qquad \forall t \in T, \tag{5.3}$$

$$e_t \leq y_n \qquad \forall t \in T, \tag{5.4}$$

$$e_t, c_t \geq 0 \qquad \forall t \in T. \tag{5.5}$$

The agent's type is now a slope α, the allocation is an expected output $e \in [0, y_n]$, the "transfer" is a cost $c \geq 0$, and the outside option in (5.3) is still type-dependent.

It is a standard observation that problem (5.1–5.5) is valid even if $\operatorname{supp}(p) \subsetneq \{\alpha_1, \ldots, \alpha_k\}$ so that $p_t = 0$ for some t. This is because Nature can satisfy constraints (5.2–5.5) for any type $\alpha_t \notin \operatorname{supp}(p)$ without any impact on the objective (5.1) by simply setting

$$(e_t, c_t) \in \arg\max\{\alpha_t e - c : (e, c) \in A^0 \cup \{(e_s, c_s)\}_{s \,:\, \alpha_s \in \operatorname{supp}(p)}\}.$$

The next result allows for this possibility in anticipation of the analysis that follows.

Lemma 4. *Let p be a random linear contract such that* $\operatorname{supp}(p) \subseteq \{\alpha_1, \ldots, \alpha_k\}$ *for some*

$\alpha_1 < \cdots < \alpha_k$. *Write* $p_t = p(\alpha_t)$ *for all* $t \in T = \{1, \ldots, k\}$ *and* $\alpha_0 \equiv 0$. *Then the guarantee* $V(p)$ *is the optimal value of the following program:*

$$
\begin{aligned}
& \min_{\{e_t\}} \sum_{t\in T} p_t(1-\alpha_t)e_t && && (5.6)\\
\text{s.t.} \quad & \sum_{s=1}^{t}(\alpha_s - \alpha_{s-1})e_s \geq U^0(\alpha_t) && \forall t \in T, \quad (\lambda_t) && (5.7)\\
& e_{t+1} - e_t \geq 0 && \forall t \in T \setminus \{k\}, \quad (\theta_t) && (5.8)\\
& e_t \geq 0 && \forall t \in T. && (5.9)
\end{aligned}
$$

The proof consists of variations of standard constraint-simplification arguments, so we relegate it to the appendix. Substantively, the key step is noting that upward adjacent incentive constraints in (5.2) can be taken to hold with equality without loss of optimality, i.e., $\alpha_t e_t - c_t = \alpha_t e_{t+1} - c_{t+1}$ for all $t < k$. This allows us to recursively eliminate the costs from the problem up to c_1, which can then be set to zero without loss of optimality.

Lemma 4 has the following dual characterization, which will play an important role.

Lemma 5. *Let* p *be a random contract satisfying the assumptions of Lemma 4. Then*

$$
\begin{aligned}
& V(p) = \max_{\lambda,\theta} \sum_{t\in T} \lambda_t U^0(\alpha_t) && && (5.10)\\
\text{s.t.} \quad & \sum_{s=t}^{k} \lambda_s(\alpha_t - \alpha_{t-1}) + \theta_{t-1} - \theta_t \leq p_t(1-\alpha_t) && \forall t \in T, \quad (e_t) && (5.11)\\
& \lambda_t \geq 0 && \forall t \in T, && (5.12)\\
& \theta_t \geq 0 && \forall t \in T \setminus \{k\}, && (5.13)
\end{aligned}
$$

where $\alpha_0 \equiv 0$, $\theta_0 \equiv 0$, *and* $\theta_k \equiv 0$.

Proof. Program (5.10–5.13) is the dual of (5.6–5.9), so the claim follows by strong duality. □

5.2 Grid-Based Contracts

A natural and tractable class of random linear contracts consists of those restricted to using slopes from a fixed, finite grid. Formally, we define a *grid* to be a set $\Gamma \subset [0, 1]$ of the form $\Gamma = \{\alpha_1, \ldots, \alpha_k\}$ for some number $k \in \mathbb{N}_{>0}$ and slopes $0 < \alpha_1 < \cdots < \alpha_k \equiv 1$ defined by $\alpha_t := t/k$ for $t = 0, \ldots, k$. (Thus, $\alpha_{t+1} - \alpha_t = k^{-1}$.) Note that the grid is uniquely determined by its size $|\Gamma| = k$.

We denote the set of random linear contracts based on the grid Γ by

$$\Delta(\Gamma) := \{p \in \Delta([0,1]) : \operatorname{supp}(p) \subseteq \Gamma\}.$$

We then define the maximal guarantee from Γ-based contracts as

$$V(\Gamma) := \max_{p \in \Delta(\Gamma)} V(p).$$

(We shall see later that the maximum exists.) Note that $V(\Gamma) \leq V(\Gamma')$ whenever $\Gamma \subset \Gamma'$.

The following result verifies that, as one would expect, the guarantee of any random linear contract can be approximated arbitrarily well with grid-based contracts.

Lemma 6. *Let p be a random linear contract and let $\varepsilon > 0$. Then there exists a grid size $\bar{k} \in \mathbb{N}$ such that $V(\Gamma) \geq V(p) - \varepsilon$ whenever $|\Gamma| \geq \bar{k}$.*

See the appendix for the proof, which relies on continuity properties of linear programs.

The above approximation lemma implies that we can obtain a maximizing sequence of random linear contracts by considering optimal grid-based contracts for ever-increasing grids.

Proposition 4. *Let $\Gamma_1, \Gamma_2, \ldots$ be a sequence of grids such that $|\Gamma_m| \to \infty$ as $m \to \infty$. Then*

$$\lim_{m\to\infty} V(\Gamma_m) = \sup_{p \in \Delta(\mathbb{R}^n_+)} V(p).$$

Proof. By Corollary 2, we may take the supremum on the right over random linear contracts. Moreover, we clearly have $\limsup_{m\to\infty} V(\Gamma_m) \leq \sup_{p\in\Delta([0,1])} V(p)$, and thus it suffices to show that $\liminf_{m\to\infty} V(\Gamma_m) \geq \sup_{p\in\Delta([0,1])} V(p)$.

Fix $\varepsilon > 0$. Take a sequence of random linear contracts $p_n \in \Delta([0,1])$ such that $V(p_n) \to \sup_{p\in\Delta([0,1])} V(p)$. By Lemma 6, for every n, there exists m_n such that

$$V(p_n) - \frac{\varepsilon}{n} \leq \inf_{m \geq m_n} V(\Gamma_m).$$

By redefining the cutoffs recursively via $\hat{m}_1 := m_1$ and $\hat{m}_n := \max\{\hat{m}_1, \ldots, \hat{m}_{n-1}\} + 1$ for $n > 1$ if necessary, we may assume that $m_1 < m_2 < \cdots$. Then $\inf_{m\geq m_n} V(\Gamma_m)$ is nondecreasing in n and thus it converges to $\liminf_m V(\Gamma_m)$ as $n \to \infty$. Hence,

$$\sup_{p\in\Delta([0,1])} V(p) = \lim_{n\to\infty} \left(V(p_n) - \frac{\varepsilon}{n}\right) \leq \lim_{n\to\infty} \inf_{m\geq m_n} V(\Gamma_m) = \liminf_{m\to\infty} V(\Gamma_m). \qquad \square$$

Motivated by Proposition 4, we characterize optimal grid-based contracts. In order to rule out pathological cases, we say that a grid Γ is *eligible* if there exists $\alpha \in \Gamma$ such that

$\alpha < 1$ and $U^0(\alpha) > 0$. By assumption, the known technology A^0 is non-trivial and thus it contains an action (π, c) such that $U^0(1) = \pi y - c > 0$. Because $U^0(\alpha)$ is continuous in α, it follows that every sufficiently large grid is eligible.

Proposition 5. *Suppose Γ is an eligible grid of size k. Let*

$$t^* \in \arg\max_{t\in T} \frac{U^0(\frac{t}{k})}{\sum_{s=1}^{t}\frac{1}{k-s}}. \tag{5.14}$$

Then $t^* < k$ and the maximal guarantee from Γ-based contracts is

$$V(\Gamma) = \frac{U^0(\frac{t^*}{k})}{\sum_{s=1}^{t^*}\frac{1}{k-s}} > 0, \tag{5.15}$$

and it is attained by the random linear contract $p \in \Delta(\Gamma)$ defined by

$$p_t = p(\tfrac{t}{k}) = \begin{cases} \frac{1}{k-t}\Big(\sum_{s=1}^{t^*}\frac{1}{k-s}\Big)^{-1} & \textit{if } 1 \le t \le t^*, \\ 0 & \textit{if } t^* < t \le k. \end{cases} \tag{5.16}$$

Proof. The definition of $V(\Gamma)$ and Lemma 5 imply that

$$V(\Gamma) = \max_{p\in\Delta(\Gamma)} V(p) = \max_{p\in\Delta(\Gamma)} \max_{\lambda,\theta} \Big\{ \sum_{t\in T} \lambda_t U^0(\alpha_t) \text{ s.t. } (5.11–5.13)\Big\}. \tag{5.17}$$

Writing the max-max problem on the right-hand side as a single maximization problem and using the facts that $\alpha_t = \frac{t}{k}$ and $\alpha_t - \alpha_{t-1} = k^{-1}$ for all t gives the following program:

$$V(\Gamma) = \max_{p,\lambda,\theta} \sum_{t\in T} \lambda_t U^0(\tfrac{t}{k}) \tag{5.18}$$

$$\text{s.t.} \quad k^{-1}\sum_{s=t}^{k}\lambda_s + \theta_{t-1} - \theta_t - p_t(1-\tfrac{t}{k}) \le 0 \qquad \forall t \in T, \qquad (e_t) \tag{5.19}$$

$$\sum_{t\in T} p_t = 1, \qquad (\delta) \tag{5.20}$$

$$p_t, \lambda_t, \theta_t \ge 0 \qquad \forall t \in T, \tag{5.21}$$

where $\theta_0 \equiv 0$ and $\theta_k \equiv 0$.

By strong duality, $V(\Gamma)$ is also the value to the following dual:

$$
\begin{aligned}
& V(\Gamma) = \min_{\{e_t\},\delta} \ \delta && && (5.22)\\
\text{s.t.} \quad & \delta - (1-\tfrac{t}{k})e_t \geq 0 && \forall t \in T, \quad (p_t) && (5.23)\\
& k^{-1}\sum_{s=1}^{t} e_s \geq U^0(\tfrac{t}{k}) && \forall t \in T, \quad (\lambda_t) && (5.24)\\
& e_{t+1} - e_t \geq 0 && \forall t \in T \setminus \{k\}, \quad (\theta_t) && (5.25)\\
& e_t \geq 0 && \forall t \in T. && (5.26)
\end{aligned}
$$

Define t^* as in (5.14). Because the grid Γ is eligible, we have $t^* < k$ and the expression on the right-hand side of (5.15) is strictly positive.[11] We will show that the following constitute optimal solutions to programs (5.22–5.26) and (5.18–5.21):

$$
\delta = \frac{U^0(\frac{t^*}{k})}{\sum_{s=1}^{t^*}\frac{1}{k-s}}, \quad e_t = \delta(1-\tfrac{t}{k})^{-1} \quad \forall t \in T \setminus \{k\}, \quad e_k = \max\{e_{k-1}, y_n\}, \tag{5.27}
$$

and

$$
\theta_t = 0 \ \ \forall t \in T \setminus \{k\}, \quad \lambda_t = \begin{cases} \left(\sum_{s=1}^{t^*}\frac{1}{k-s}\right)^{-1} & \text{if } t = t^*, \\ 0 & \text{if } t \neq t^*, \end{cases} \quad p_t = \begin{cases} \frac{\lambda_{t^*}}{k-t} & \text{if } 1 \leq t \leq t^*, \\ 0 & \text{if } t^* < t \leq k. \end{cases} \tag{5.28}
$$

We first check that the objectives (5.18) and (5.22) coincide under the above solutions:

$$
\delta = \frac{U^0(\frac{t^*}{k})}{\sum_{s=1}^{t^*}\frac{1}{k-s}} = \lambda_{t^*}U^0(\tfrac{t^*}{k}) = \sum_{t\in T}\lambda_t U^0(\tfrac{t}{k}).
$$

Therefore, the solutions are optimal by weak duality, if we show them to be feasible.

Note that (5.27) and (5.28) clearly satisfy the non-negativity constraints (5.21) and (5.26) as well as the monotonicity constraint (5.25). We verify the other constraints in turn.

Consider first constraint (5.19). If $1 \leq t \leq t^*$, we have

$$
k^{-1}\sum_{s=t}^{k}\lambda_t + \theta_{t-1} - \theta_t - p_t(1-\tfrac{t}{k}) = \frac{\lambda_{t^*}}{k} - \frac{\lambda_{t^*}}{k-t}\frac{k-t}{k} = 0.
$$

If $t^* < t \leq k$, then $k^{-1}\sum_{s=t}^{k}\lambda_t + \theta_{t-1} - \theta_t - p_t(1-\frac{t}{k}) = 0$. Thus, (5.19) holds for all $t \in T$.

[11] To see this, adopt the usual conventions $a/\infty = 0$ for $a \geq 0$ and $a/0 = \infty$ for $a > 0$. Evaluating the objective in (5.14) at $t = k$ then gives $U^0(1)/\infty = 0$, whereas eligibility implies that the objective is positive for some $t < k$.

We then verify (5.20):

$$\delta\left[\sum_{t\in T} p_t - 1\right] = \delta\left[\sum_{t=1}^{t^*} \frac{\lambda_{t^*}}{k-t} - 1\right] = \delta\left[\frac{\sum_{t=1}^{t^*}\frac{1}{k-t}}{\sum_{s=1}^{t^*}\frac{1}{k-s}} - 1\right] = 0.$$

Moving to the dual, consider (5.23). We have $\delta - (1-\frac{t}{k})e_t = 0$ for all $t < k$, whereas for $t = k$, we get $\delta - (1-\frac{k}{k})e_k = \delta > 0$ verifying the constraint.

Consider then (5.24). Observe that for all $t < k$, the left-hand side is

$$\frac{1}{k}\sum_{s=1}^{t} e_s = \frac{1}{k}\sum_{s=1}^{t}\frac{\delta}{1-\frac{s}{k}} = \delta\sum_{s=1}^{t}\frac{1}{k-s} = \frac{U^0(\frac{t^*}{k})}{\sum_{s=1}^{t^*}\frac{1}{k-s}}\sum_{s=1}^{t}\frac{1}{k-s} \geq U^0(\tfrac{t}{k}),$$

where the inequality follows by definition of t^* in (5.14). Thus, the constraint holds for all $t < k$, with equality for $t = t^*$. It remains to verify the constraint for $t = k$. We have

$$\frac{1}{k}\sum_{s=1}^{k} e_s \geq \frac{1}{k}\left(\sum_{s=1}^{k-1}\frac{\delta}{1-\frac{s}{k}} + y_n\right) \geq U^0(\tfrac{k-1}{k}) + \frac{y_n}{k} \geq U^0(\tfrac{k}{k}),$$

where the last inequality follows because—letting (e_k^0, c_k^0) denote the optimal choice of type k from A^0—we have

$$U^0(\tfrac{k}{k}) - U^0(\tfrac{k-1}{k}) \leq e_k^0 - c_k^0 - \left(\frac{k-1}{k}e_k^0 - c_k^0\right) = \left(1 - \frac{k-1}{k}\right)e_k^0 \leq \frac{y_n}{k}.$$

We conclude that (5.24) holds for $t = k$ as well.

Therefore, the solutions (5.27) and (5.28) are feasible and hence optimal. The maximal guarantee (5.15) now follows from (5.27) because $V(\Gamma) = \delta > 0$ by (5.22). The contract p defined by (5.16) attains it by the optimality of (5.28). □

It is worth noting that the principal is indifferent among all linear contracts in the support of the random linear contract p in (5.16). To see this, note that (5.27) implies that $(1-\frac{t}{k})e_t = \delta$ for all $1 \leq t \leq t^*$, because $t^* < k$. Importantly, this means that implementing p does not require the principal to be able to commit to the randomization.

5.3 The Optimal Guarantee and the Gain from Randomization

We are now in position to characterize the optimal guarantee over all random contracts. Notice that the quotient

$$\frac{U^0(\alpha)}{-\ln(1-\alpha)}$$

is well-defined and continuous in α on $(0,1)$. By convention, we extend it to all of $[0,1]$ by its left and right limits at the end points.[12] This simplifies notation as the quotient is then essentially a continuous function on $[0,1]$, with the possible exception that it may take on the value $-\infty$ at $\alpha = 0$. In particular, it attains its maximum on $[0,1]$, and this maximum is the optimal guarantee.

Proposition 6. *The optimal guarantee is*

$$\sup_{p\in\Delta(\mathbb{R}^n_+)} V(p) = \max_{\alpha\in[0,1]} \frac{U^0(\alpha)}{-\ln(1-\alpha)}.$$

Proof. Let $\bar{v} := \sup_{p\in\Delta(\mathbb{R}^n_+)} V(p)$. Because A^0 is non-trivial, the optimal deterministic guarantee is positive, and thus $\bar{v} > 0$ a fortiori.

Denote the grid of size k by Γ_k. Throughout the proof, we assume without loss of generality that k is large enough so that Γ_k is eligible. Then Propositions 4 and 5 imply that

$$0 < \bar{v} = \lim_{k\to\infty} V(\Gamma_k) = \lim_{k\to\infty} \frac{U^0(\frac{t^*_k}{k})}{\sum_{s=1}^{t^*_k} \frac{1}{k-s}}, \tag{5.29}$$

where we have labeled the maximizer t^*_k by k to highlight its dependence on the grid size. In particular, the limit on the right exists.

Let $\alpha^*_k := t^*_k/k$. Because $1 \le t^*_k < k$, we have $\alpha^*_k \in (0,1)$. By extracting a subsequence if necessary, we may assume that α^*_k converges to some $\alpha^* \in [0,1]$ as $k\to\infty$.

The numerator on the right-hand side of (5.29) is continuous by continuity of $U^0(\alpha)$ in α. To evaluate the denominator, we prove the following result in the appendix.

Lemma 7. *Let (t_k) be a sequence in $\mathbb{N}$ such that $1 \le t_k < k$ for all k and $\frac{t_k}{k} \to \alpha \in [0,1]$. Then*

$$\lim_{k\to\infty} \sum_{s=1}^{t_k} \frac{1}{k-s} = \begin{cases} -\ln(1-\alpha) & \text{if } \alpha \in [0,1), \\ +\infty & \text{if } \alpha = 1. \end{cases}$$

We can now show that $\alpha^* \in [0,1)$. Suppose to the contrary that $\alpha^* = 1$. Then using

[12] That is, let

$$\frac{U^0(0)}{-\ln(1-0)} := \lim_{\alpha\to 0^+} \frac{U^0(\alpha)}{-\ln(1-\alpha)} \in [-\infty,\infty) \quad \text{and} \quad \frac{U^0(1)}{-\ln(1-1)} := \lim_{\alpha\to 1^-} \frac{U^0(\alpha)}{-\ln(1-\alpha)} = 0,$$

where the first limit is bounded from above because $U^0(0) \le 0$. However, it may be positive. For example, if $A^0 = \{(\pi, 0)\}$ and $\pi y > 0$, then $\frac{U^0(0)}{-\ln(1-0)} = \lim_{\alpha\to 0^+} \frac{\alpha\pi y}{-\ln(1-\alpha)} = \pi y > 0$.

continuity of U^0 and Lemma 7 (with $t_k = t_k^*$ and $\alpha^* = 1$) yields a contradiction in (5.29):

$$0 < \bar{v} = \lim_{k\to\infty} \frac{U^0(\frac{t_k^*}{k})}{\sum_{s=1}^{t_k^*} \frac{1}{k-s}} \left(= \frac{U^0(1)}{+\infty}\right) = 0.$$

We will then show that $\bar{v} \geq \max_{\alpha\in[0,1]} \frac{U^0(\alpha)}{-\ln(1-\alpha)}$. Let $\alpha \in (0,1)$ and take a sequence (t_k) in $\mathbb{N}$ such that $1 \leq t_k < k$ and $\alpha_k := t_k/k \to \alpha$. By definition of t_k^* as a maximizer in (5.14), we have for all (sufficiently large) k the following inequality:

$$\frac{U^0(\frac{t_k}{k})}{\sum_{s=1}^{t_k} \frac{1}{k-s}} \leq \frac{U^0(\frac{t_k^*}{k})}{\sum_{s=1}^{t_k^*} \frac{1}{k-s}}.$$

Taking the limit $k \to \infty$ on both sides, using continuity of U^0 and Lemma 7 on the left and (5.29) on the right, gives

$$\frac{U^0(\alpha)}{-\ln(1-\alpha)} = \lim_{k\to\infty} \frac{U^0(\frac{t_k}{k})}{\sum_{s=1}^{t_k} \frac{1}{k-s}} \leq \lim_{k\to\infty} \frac{U^0(\frac{t_k^*}{k})}{\sum_{s=1}^{t_k^*} \frac{1}{k-s}} = \bar{v}. \tag{5.30}$$

The inequality extends from $\alpha \in (0,1)$ to all $\alpha \in [0,1]$ by taking left and right limits.

It remains to show that $\bar{v} \leq \max_{\alpha\in[0,1]} \frac{U^0(\alpha)}{-\ln(1-\alpha)}$. If $\alpha^* \in (0,1)$, then this is immediate from (5.30) because in this case the limit on the right equals $\frac{U^0(\alpha^*)}{-\ln(1-\alpha^*)}$ by continuity of U^0 and Lemma 7.

To complete the proof, let $\alpha^* = 0$. Then $U^0(\frac{t_k^*}{k}) \to U^0(0) \leq 0$ and $(\sum_{s=1}^{t_k^*} \frac{1}{k-s})^{-1} \to 0$ by Lemma 7. Because the limit of their ratio is $\bar{v} > 0$ by (5.29), we must have $U^0(0) = 0$. This implies that A^0 contains a zero cost action. Furthermore, for any $\alpha > 0$ small enough, the agent must choose a zero cost action as A^0 is finite. Thus, for any $\alpha > 0$ small enough,

$$U^0(\alpha) = \alpha \max_{(\pi,c)\in A^0:\, c=0} \pi y, \quad \text{and} \quad V(\alpha|A^0) = (1-\alpha) \max_{(\pi,c)\in A^0:\, c=0} \pi y.$$

Let e^0 denote the expected output of an action attaining the above maxima. Recall from (5.16) that $\alpha_k^* = t_k^*/k$ is the largest slope to receive positive weight under the optimal random linear contract p^k for the grid Γ_k. Because $\alpha_k^* \to 0$, for all k large enough, the agent will choose an action with expected output e^0 if the true technology is the known technology A^0, regardless of which contract in the support of p^k is realized. Hence, for all k large enough,

$$V(\Gamma_k) = V(p^k) \leq V(p^k|A^0) = \sum_{t=1}^{t^*} p_t^k V(\tfrac{t}{k}|A^0) = \sum_{t=1}^{t^*} p_t^k (1 - \tfrac{t}{k}) e^0 \leq e^0.$$

By (5.29) this implies that $\bar{v} \leq e^0$. On the other hand, for any $\alpha > 0$ small enough,

$$\frac{U^0(\alpha)}{-\ln(1-\alpha)} = \frac{\alpha e^0}{-\ln(1-\alpha)} \to e^0 \quad \text{as } \alpha \to 0.$$

We conclude that if $\alpha^* = 0$, then $\bar{v} \leq e^0 = \frac{U^0(0)}{-\ln(1-0)} \leq \max_{\alpha \in [0,1]} \frac{U^0(\alpha)}{-\ln(1-\alpha)} \leq \bar{v}$. □

The *gain from randomization* is the ratio between the optimal guarantee and the optimal deterministic guarantee, or $\sup_{p \in \Delta(\mathbb{R}^n_+)} V(p) / \sup_{w \in \mathbb{R}^n_+} V(w)$. We say that the gain is positive if the ratio is greater than 1. By Proposition 6 and (3.11), the two guarantees are

$$\sup_{p \in \Delta(\mathbb{R}^n_+)} V(p) = \max_{\alpha \in [0,1]} \frac{U^0(\alpha)}{-\ln(1-\alpha)} \quad \text{and} \quad \sup_{w \in \mathbb{R}^n_+} V(w) = \max_{\alpha \in [0,1]} \frac{1-\alpha}{\alpha} U^0(\alpha).$$

It is then straightforward to compute the gain from randomization for any given known technology A^0 (which has to be specified in order to compute $U^0(\alpha)$).

Kambhampati (2023) showed that the gain from randomization is positive under principal-optimal tie-breaking, provided there is an optimal deterministic contract different from the zero contract. We can replicate his result for adversarial tie-breaking as follows.

Corollary 3. *If there exists an optimal deterministic contract (which is then necessarily different from the zero contract), then the gain from randomization is positive.*

Proof. Suppose there exists an optimal deterministic contract. Then some linear contract $\alpha^*_D > 0$ is optimal by Proposition 2. Note that $\alpha^*_D < 1$ and $U^0(\alpha^*_D) > 0$ because the non-triviality of A^0 implies that the optimal deterministic guarantee is positive. Hence,

$$\max_{\alpha \in [0,1]} \frac{U^0(\alpha)}{-\ln(1-\alpha)} - \max_{\alpha \in [0,1]} \frac{1-\alpha}{\alpha} U^0(\alpha) \geq \left(\frac{1}{-\ln(1-\alpha^*_D)} - \frac{1-\alpha^*_D}{\alpha^*_D} \right) U^0(\alpha^*_D) > 0,$$

where the last inequality is because the expression in the parentheses is positive on $(0, 1)$. □

Corollary 3 gives a sufficient condition for the gain from randomization to be positive in terms of (3.11): The gain is positive if the maximum on the right is attained by some $\alpha^*_D > 0$. As noted in Section 3, a simple sufficient condition for this is that every action in A^0 with positive expected output has a positive cost.

We can give a tight characterization in terms of our formula for the optimal guarantee.

Corollary 4. *The gain from randomization is positive if and only if*

$$\frac{U^0(0)}{-\ln(1-0)} < \max_{\alpha \in [0,1]} \frac{U^0(\alpha)}{-\ln(1-\alpha)}.$$

Proof. Let $\bar{v}$ be the maximum on the right-hand side of the above inequality (i.e., $\bar{v}$ is the optimal guarantee). Suppose first that the inequality is not satisfied. Then

$$\bar{v} = \frac{U^0(0)}{-\ln(1-0)} = \lim_{\alpha \to 0^+} \frac{U^0(\alpha)}{-\ln(1-\alpha)} \left(= \frac{U^0(0)}{0} \right).$$

As in the proof of Proposition 6, we deduce that we must have $U^0(0) = 0$ and the known technology A^0 must thus contain a zero cost action. Moreover, for any $\alpha > 0$ small enough, the agent's choice from A^0 must be a zero cost action that has the maximum expected output among zero cost actions in A^0. Letting e^0 denote this maximum expected output, we have

$$\bar{v} = \lim_{\alpha \to 0^+} \frac{U^0(\alpha)}{-\ln(1-\alpha)} = \lim_{\alpha \to 0^+} \frac{\alpha e^0}{-\ln(1-\alpha)} = e^0.$$

On the other hand, the optimal deterministic guarantee is then also at least e^0, because

$$\lim_{\alpha \to 0^+} \frac{1-\alpha}{\alpha} U^0(\alpha) = \lim_{\alpha \to 0^+} \frac{1-\alpha}{\alpha} \alpha e^0 = e^0.$$

We conclude that there is no gain from randomization.

Suppose then that the inequality is satisfied. If there exists an optimal deterministic contract, then the gain from randomization is positive by Corollary 3. So suppose this is not the case. Letting $\bar{v}_D$ denote the optimal deterministic guarantee, we then have

$$0 < \bar{v}_D = \lim_{\alpha \to 0^+} \frac{1-\alpha}{\alpha} U^0(\alpha) \left(= \frac{U^0(0)}{0} \right),$$

and thus $U^0(0) = 0$, which in turn implies that $U^0(\alpha) \geq 0$ for all α. Therefore,

$$\bar{v}_D = \lim_{\alpha \to 0^+} \frac{1-\alpha}{\alpha} U^0(\alpha) \leq \lim_{\alpha \to 0^+} \frac{U^0(\alpha)}{-\ln(1-\alpha)} = \frac{U^0(0)}{-\ln(1-0)} < \bar{v},$$

where the first inequality follows because $(1-\alpha)/\alpha < -1/\ln(1-\alpha)$ for all $\alpha \in (0,1)$. $\square$

The following example shows that the gain from randomization can be arbitrarily large.

Example 1. Suppose the known technology A^0 consists of a single action (π, c), whose expected output πy is normalized to 1 and $c < 1$. Then $U^0(\alpha) = \alpha - c$. A straightforward calculation using (3.11) gives the optimal deterministic guarantee $\max_{\alpha \in [0,1]} \frac{1-\alpha}{\alpha}(\alpha - c) = (1-\sqrt{c})^2$. On the other hand, Proposition 6 implies that the optimal guarantee is

$$\max_{\alpha \in [0,1]} \frac{\alpha - c}{-\ln(1-\alpha)}.$$

The first-order condition to this problem is $-\ln(1-\alpha) - (1-\alpha)^{-1}(\alpha - c) = 0$, which can readily be verified to have a unique solution, $\alpha^*(c)$, which is a continuous increasing function of c, with $\lim_{c\to 0}\alpha^*(c) = 0$ and $\lim_{c\to 1}\alpha^*(c) = 1$. The Implicit Function Theorem implies that its derivative on $(0,1)$ is $(\alpha^*)'(c) = -1/\ln(1-\alpha^*(c))$. Substituting back into the objective then gives the optimal guarantee $1-\alpha^*(c)$. The gain from randomization is thus

$$g(c) := \frac{\sup_{p\in\Delta(\mathbb{R}^n_+)} V(p)}{\sup_{w\in\mathbb{R}^n_+} V(w)} = \frac{1-\alpha^*(c)}{(1-\sqrt{c})^2}.$$

From the properties of $\alpha^*(c)$ it follows that the gain $g(c)$ is a continuous function of c with $g(0) = 1$. Using l'Hôpital's rule twice shows that $\lim_{c\to 1} g(c) = \infty$. This example thus exhibits the entire range of possibilities as c varies.

5.4 The Limit "Contract"

Propositions 4 and 5 imply that the random linear contract defined by (5.16) approximates the optimal guarantee arbitrarily well as the grid size k grows without a bound. The following result confirms that such a maximizing sequence of contracts is in a sense the strongest result we can hope for as the optimal guarantee is not attained by any finite grid.

Corollary 5. *The optimal guarantee is only attained asymptotically:* $V(\Gamma) < \sup_{p\in\Delta(\mathbb{R}^n_+)} V(p)$ *for every grid* Γ.

Proof. For each $k \in \mathbb{N}_{>0}$, let Γ_k be a grid of size k. Then $V(\Gamma_k)$ converges to the optimal guarantee as $k\to\infty$ by Proposition 4. We prove the corollary by showing that this sequence has a strictly increasing subsequence.

Fix $\hat{k}$ large enough for the grid $\Gamma_{2^m\hat{k}}$ to be eligible for all $m \in \mathbb{N}$. We claim that $V(\Gamma_{2^m\hat{k}}) < V(\Gamma_{2^{m+1}\hat{k}})$ for all m. To see this, fix m. Let $k := 2^m\hat{k}$ so that $2^{m+1}\hat{k} = 2k$. Define t^* for the grid Γ_k as in (5.14). Then (5.15) implies that

$$V(G_{2k}) - V(G_k) \geq \frac{U^0(\frac{2t^*}{2k})}{\sum_{s=1}^{2t^*}\frac{1}{2k-s}} - \frac{U^0(\frac{t^*}{k})}{\sum_{s=1}^{t^*}\frac{1}{k-s}} = \left(\frac{1}{\sum_{s=1}^{2t^*}\frac{1}{2k-s}} - \frac{1}{\sum_{s=1}^{t^*}\frac{1}{k-s}}\right)U^0(\tfrac{t^*}{k}).$$

By Proposition 5, $U^0(\frac{t^*}{k}) > 0$. It thus suffices to show that $\sum_{s=1}^{t^*}\frac{1}{k-s} > \sum_{s=1}^{2t^*}\frac{1}{2k-s}$. We have

$$\begin{aligned}\sum_{s=1}^{t^*}\frac{1}{k-s} = \sum_{s=1}^{t^*}\frac{2}{2k-2s} &> \sum_{s=1}^{t^*}\left(\frac{1}{2k-(2s-1)} + \frac{1}{2k-2s}\right)\\ &= \frac{1}{2k-1} + \frac{1}{2k-2} + \cdots + \frac{1}{2k-(2t^*-1)} + \frac{1}{2k-2t^*} = \sum_{s=1}^{2t^*}\frac{1}{2k-s}.\end{aligned}$$ $\square$

We note that for any eligible grid Γ, the amount by which $V(\Gamma)$ falls short of the optimal guarantee can be calculated directly as the difference (or ratio) between (5.15) and the formula for the optimal guarantee in Proposition 6.

Given the above result, it becomes of interest to understand the limiting behavior of our maximizing sequence. We show next that whenever the gain from randomization is positive, it has a weakly convergent subsequence (which is, of course, itself also a maximizing sequence) whose limit is a particular continuous distribution on the unit interval.

Proposition 7. *Fix $\bar{k}$ such that every grid of size $k \geq \bar{k}$ is eligible. For all $k \geq \bar{k}$, let F_k be the distribution function for the random linear contract (5.16) and let t_k^* be the corresponding choice in (5.14). If the gain from randomization is positive, then the sequence $(F_k)_{k \geq \bar{k}}$ has a subsequence that converges weakly to the distribution function F defined by*

$$F(\alpha) := \frac{\ln(1-\alpha)}{\ln(1-\alpha^*)} \quad \forall \alpha \in [0, \alpha^*],$$

where $\alpha^* \in (0,1)$ is the limit of $\frac{t_k^*}{k}$ along the subsequence, and it satisfies

$$\alpha^* \in \underset{\alpha \in [0,1]}{\arg\max} \frac{U^0(\alpha)}{-\ln(1-\alpha)}.$$

Proof. Suppose the gain from randomization is positive. By extracting a subsequence if necessary, we may assume that $\frac{t_k^*}{k} \in (0,1)$ converges to some $\alpha^* \in [0,1]$. It was shown in the proof of Proposition 6 that $\alpha^* < 1$, and that it attains the maximum of $\frac{U^0(\alpha)}{-\ln(1-\alpha)}$ on $[0,1]$. In particular, if $\alpha^* = 0$, then $\frac{U^0(0)}{-\ln(1-0)} = \max_{\alpha \in [0,1]} \frac{U^0(\alpha)}{-\ln(1-\alpha)}$, which contradicts the gain from randomization being positive by Corollary 4. Thus, $\alpha^* \in (0,1)$.

By (5.16), each distribution function $F_k : [0,1] \to \mathbb{R}_+$ satisfies

$$F_k(\alpha) = \begin{cases} 0 & \text{if } \alpha \in [0, \frac{1}{k}), \\ \sum_{u=1}^{t} \frac{1}{k-u} / \sum_{s=1}^{t_k^*} \frac{1}{k-s} & \text{if } \alpha \in [\frac{t}{k}, \frac{t+1}{k}) \text{ and } 1 \leq t \leq t_k^*, \\ 1 & \text{if } \alpha \in [\frac{t_k^*+1}{k}, 1]. \end{cases}$$

The three different cases can equivalently be expressed as (i) $\alpha k \in [0,1)$, (ii) $\alpha k \in [t, t+1)$ and $1 \leq t \leq t_k^*$, and (iii) $\alpha k \in [t_k^* + 1, k]$. Using the floor function, we have

$$F_k(\alpha) = \begin{cases} 0 & \text{if } \lfloor \alpha k \rfloor = 0, \\ \sum_{u=1}^{\lfloor \alpha k \rfloor} \frac{1}{k-u} / \sum_{s=1}^{t_k^*} \frac{1}{k-s} & \text{if } 1 \leq \lfloor \alpha k \rfloor \leq t_k^*, \\ 1 & \text{if } \lfloor \alpha k \rfloor > t_k^*. \end{cases}$$

We want to show that $F_k(\alpha) \to F(\alpha)$ for all $\alpha \in [0, 1]$. Note first that if $\alpha = 0$, then $F_k(0) = 0 = F(0)$ for all k. Suppose then that $\alpha > \alpha^* = \lim \frac{t_k^*}{k} = \lim \frac{t_k^*+1}{k}$. Then, for all k large enough, we have $\alpha \geq \frac{t_k^*+1}{k}$, or $\alpha k \geq t_k^* + 1$. Because $t_k^* \in \mathbb{N}$, we then also have $\lfloor \alpha k \rfloor \geq t_k^* + 1$ for all k large enough. This implies that, eventually, $F_k(\alpha) = 1 = F(\alpha)$.

Consider now $\alpha \in (0, \alpha^*)$. Define the sequence $(t_k)_{k \geq \bar{k}}$ in $\mathbb{N}$ by $t_k := \lfloor \alpha k \rfloor$. We observe that for k large enough, $1 \leq t_k < k$. Moreover, we have $\frac{t_k}{k} \to \alpha$, because

$$\alpha \geq \frac{t_k}{k} \geq \frac{\alpha k - 1}{k} \to \alpha \quad \text{as } k \to \infty.$$

From $\lim \frac{t_k}{k} = \alpha < \alpha^* = \lim \frac{t_k^*}{k}$ it follows that, eventually, $\lfloor \alpha k \rfloor = t_k \leq t_k^*$, i.e., we are in the second case of F_k. Thus, applying Lemma 7 to the numerator and the denominator gives

$$\lim_{k \to \infty} F_k(\alpha) = \lim_{k \to \infty} \frac{\sum_{u=1}^{\lfloor \alpha k \rfloor} \frac{1}{k-u}}{\sum_{s=1}^{t_k^*} \frac{1}{k-s}} = \frac{-\ln(1-\alpha)}{-\ln(1-\alpha^*)} = F(\alpha). \tag{5.31}$$

It remains to consider the case $\alpha = \alpha^*$. By extracting a subsequence, we may assume that $\frac{t_k^*}{k}$ converges monotonically to α^*. If it converges from below, then $\alpha^* \geq \frac{t_k^*}{k}$, and thus $\lfloor \alpha^* k \rfloor \geq t_k^*$, implying that $F_k(\alpha^*) = 1 = F(\alpha^*)$. If instead convergence is from above, then $\lfloor \alpha^* k \rfloor \leq \alpha^* k \leq t_k^*$, and the limit follows by setting $\alpha = \alpha^*$ in (5.31). □

The limiting distribution function F in Proposition 7 puts positive density to the non-degenerate interval $[0, \alpha^*]$. It can thus be interpreted as a "continuous random linear contract." Even if there is no way to exactly implement such a continuous distribution in practice, it provides an approximation (in the topology of convergence in distribution) to the approximately optimal random linear contracts from Proposition 5 for k large.

6 Concluding Remarks

We have shown that linear contracts remain robustly optimal when randomization over contracts is allowed, in that random linear contracts can (weakly) outperform all other contracts. The gain from randomization is positive, except for a corner case, and can be arbitrarily large. Whenever the gain is positive, the principal's optimal payoff guarantee can be approximated and asymptotically achieved by the simple grid-based random linear contracts from Proposition 5, which are in turn well-approximated by the continuous distribution in Proposition 7 when the grid is large. As the principal is indifferent among contracts in the support, these contracts do not require commitment to randomization.

We conclude by commenting briefly on possible extensions and on the role of tie-breaking.

6.1 Extensions

We have deliberately kept the model streamlined to allow for a simple and transparent analysis of the merits of linear contracts under randomization. A detailed account of possible extensions is beyond the scope of this paper, but we note here some that are immediate.

The simplest way to impose a participation constraint is to assume that the known technology A^0 contains an action that results in zero output at no cost to the agent—this is subsumed by our general analysis. If instead we assume that any contract has to deliver the agent an expected payoff (net of cost) of at least $\bar{u} > 0$, then this can be handled by replacing $U^0(w)$ with $\max\{U^0(w), \bar{u}\}$ on the right-hand side of (4.3). Because each affine contract x^t constructed in the proof of Proposition 3 dominates the corresponding original contract w^t pointwise, it satisfies the so-modified participation constraint (4.3). It then follows that for every random contract, there exists a random affine contract that (weakly) outperforms it, and thus random affine contracts are optimal as a class.

Similarly, it is straightforward to allow for multi-dimensional output y from an arbitrary finite set Y, with the payoff to the principal being $v(y)$ for some function $v : Y \to \mathbb{R}_+$. Modifying the primal problem (4.1–4.5) in the obvious way, the same argument gives an improvement from moving to a random contract that is linear in $v(y)$.[13]

It also follows from our results that if we restrict attention to linear contracts in any eligible grid Γ, then randomizing over menus of linear contracts cannot increase the principal's guarantee. This is because the proof of Proposition 5 constructs a saddle-point (p, A) for the max-min problem (5.17). The guarantee of any randomization over menus of linear contracts is no better than its performance against this particular technology A. Moreover, against A, the randomization over menus reduces to a randomization over the contracts the agent chooses from each menu given A. The resulting payoff is thus no higher than $V(p)$, because p is an optimal randomization over linear contracts against A. Combining this observation with Lemma 6 and Proposition 3 suggests that randomizing over menus cannot help in general.

6.2 Role of Tie-breaking

As noted in Section 3, the optimal deterministic guarantee is the same under adversarial and principal-optimal tie-breaking. To sketch the argument for why the same is true of the optimal guarantee, notice that we can ensure strict incentives for the agent by adding a small $\varepsilon > 0$ to the right-hand sides of constraints (4.2) and (4.3). As the value of program

[13]Specifically, replace each y_i with $v(y_i)$ in the primal objective (4.1) and on the right-hand side of the dual constraint (4.7). The rest of the argument then goes through with obvious modifications.

(4.1–4.5) is continuous in the right-hand sides, it is then immediate that the guarantee of any random contract p for which the program is feasible for some $\varepsilon > 0$ is unaffected by the tie-breaking assumption. Contracts that violate feasibility for all $\varepsilon > 0$—most notably, contracts whose support contains the zero contract—have to be handled by approximation, but the conclusion is the same; we omit the details in the interest of space.[14]

Appendix

A.1 Proof of Lemma 1

We first prove a stronger result for an ε-perturbation of the problem and then establish Lemma 1 via taking the limit $\varepsilon \to 0$.

Given $\varepsilon \geq 0$, we introduce a relaxation of the incentive compatibility constraint (4.2):

$$\pi^t w^t - c_t \geq \pi^s w^t - c_s - \varepsilon \quad \forall t, s \in T : t \neq s. \tag{A.1}$$

The *ε-perturbed primal problem* consists of solving (4.1) subject to (A.1), (4.3), (4.4), and (4.5). It is feasible for all $\varepsilon \geq 0$, because it is a relaxation of the feasible primal problem (4.1–4.5). Let $V(p; \varepsilon)$ denote the value of the ε-perturbed primal problem. We clearly have $V(p; \varepsilon) \geq -\max_{i,t} w_i^t$, and hence an optimal solution exists.

Because the perturbation ε only enters the right-hand sides of some primal constraints, in the dual it only enters the objective. The *ε-perturbed dual problem* consists of solving

$$V(p; \varepsilon) = \max_{\kappa, \lambda, \mu} \sum_{t \in T} \lambda_t U(w^t) + \sum_{t \in T} \mu_t - \varepsilon \sum_{t \in T} \sum_{s \neq t} \kappa_{st} \tag{A.2}$$

subject to (4.7–4.9).

Lemma A.1. *Let $\varepsilon > 0$. If $(\kappa^*, \lambda^*, \mu^*)$ is an optimal solution to the ε-perturbed dual problem, then $G(\kappa^*)$ is acyclic.*

Proof. Let $\varepsilon > 0$ and let $(\kappa^*, \lambda^*, \mu^*)$ be an optimal solution to the ε-perturbed dual. Suppose toward contradiction that $G(\kappa^*)$ contains a cycle. By relabeling if necessary, we can assume without loss of generality that the cycle involves vertices $C := \{1, \ldots, m\}$ for $2 \leq m \leq k$,

[14] This perturbation argument also explains why an existence problem arises in the deterministic case under adversarial tie-breaking only when the zero contract would be uniquely optimal under principal-optimal tie-breaking. Given any other (optimal) contract, constraint (3.2) in the program for the deterministic guarantee can be perturbed to ensure strict incentives. However, if w is the zero contract and $U^0(0) = 0$, then it is not feasible to satisfy (3.2) as a strict inequality and thus adversarial tie-breaking can lead to a strictly lower value to the program, necessitating the approximation of the zero contract with a sequence of positive slopes.

with $\kappa^*_{t,t+1} > 0$ for all $t \in C$, where $m + 1 = 1$ by convention. Similarly, we can assume without loss of generality that $p_1 = \min_{t \in C} p_t$.

By complementary slackness, the incentive compatibility constraints in (A.1) corresponding to $\kappa^*_{t,t+1}$ for $t \in C$ hold with equality. Letting $\{(\pi^t, c_t)\}_{t \in T}$ denote an optimal solution to the ε-perturbed primal, we thus have

$$\pi^t w^t - c_t - \pi^{t+1} w^t + c_{t+1} = -\varepsilon \quad \forall t \in C. \tag{A.3}$$

Summing the equalities over $t \in C$ and recalling that $m + 1 = 1$ gives

$$\sum_{t \in C} (\pi^t - \pi^{t+1}) w^t = -m\varepsilon. \tag{A.4}$$

We will show that it is then possible to strictly lower the value of the ε-perturbed primal, contradicting optimality.

Construct a new solution for the ε-perturbed primal as follows. Let $(\hat{\pi}^t, \hat{c}_t) := (\pi^t, c_t)$ for all $t \notin C$. For each $t \in C$, let

$$(\hat{\pi}^t, \hat{c}_t) := (1 - \delta_t)(\pi^t, c_t) + \delta_t(\pi^{t+1}, c_{t+1}),$$

where $\delta_t := p_1/p_t \in (0, 1]$. As the action assigned to each contract $t \notin C$ is the same as before, their contribution to the objective is unchanged. But the contracts in C now yield

$$\begin{aligned}
\sum_{t \in C} p_t \hat{\pi}^t (y - w^t) &= \sum_{t \in C} p_t [(1 - \delta_t)\pi^t + \delta_t \pi^{t+1}](y - w^t) \\
&= \sum_{t \in C} [(p_t - p_1)\pi^t + p_1 \pi^{t+1}](y - w^t) \\
&= \sum_{t \in C} p_t \pi^t (y - w^t) + p_1 \sum_{t \in C} (\pi^t - \pi^{t+1})(w^t - y) \\
&= \sum_{t \in C} p_t \pi^t (y - w^t) - p_1 m\varepsilon,
\end{aligned}$$

where the last equality follows by (A.4) and the fact that $\sum_{t \in C} (\pi^t - \pi^{t+1}) y = 0$ because the sum is cyclical. Therefore, the new solution $\{(\hat{\pi}^t, \hat{c}_t)\}_{t \in T}$ is a strict improvement over the supposed optimal solution $\{(\pi^t, c_t)\}_{t \in T}$, provided we show that it is feasible.

To show feasibility, we note first that each $(\hat{\pi}^t, \hat{c}_t)$ satisfies the probability and non-negativity constraints (4.4) and (4.5). For $t \notin C$ this is trivial as $(\hat{\pi}^t, \hat{c}_t) = (\pi^t, c_t)$. For $t \in C$ this follows because $(\hat{\pi}^t, \hat{c}_t)$ is by definition a convex combination of two actions, each of which satisfies (4.4) and (4.5).

Note then that, given any contract w^t, the agent's payoff from the new action $(\hat{\pi}^t, \hat{c}_t)$ is weakly higher than from the old action (π^t, c_t) by construction:

$$\hat{\pi}^t w^t - \hat{c}_t = \begin{cases} \pi^t w^t - c_t + \delta_t(\pi^{t+1} w^t - c_{t+1} - \pi^t w^t + c_t) = \pi^t w^t - c_t + \delta_t \varepsilon & \text{if } t \in C, \\ \pi^t w^t - c_t & \text{if } t \notin C, \end{cases}$$

where the first case follows by the definition of $(\hat{\pi}^t, \hat{c}_t)$ and equation (A.3). This implies that the new solution satisfies the participation constraint (4.3) for all t, since

$$\hat{\pi}^t w^t - \hat{c}_t \geq \pi^t w^t - c_t \geq U^0(w^t).$$

It remains to verify the incentive compatibility constraint (A.1). Fix contracts t and $s \neq t$. Observe that by construction of $(\hat{\pi}^s, \hat{c}_s)$, we have $(\hat{\pi}^s, \hat{c}_s) = (1-\beta)(\pi^s, c_s) + \beta(\pi^{s+1}, c_{s+1})$ for $\beta \in (0,1]$ (if $s \in C$) or for $\beta = 0$ (if $s \notin C$). Therefore,

$$\hat{\pi}^t w^t - \hat{c}_t \geq \pi^t w^t - c_t \geq (1-\beta)(\pi^s w^t - c_s - \varepsilon) + \beta(\pi^{s+1} w^t - c_{s+1} - \varepsilon) = \hat{\pi}^s w^t - \hat{c}_s - \varepsilon,$$

where the second inequality follows because $\{(\pi^t, c_t)\}_{t \in T}$ satisfies (A.1) by assumption.

We conclude that the new solution $\{(\hat{\pi}^t, \hat{c}_t)\}_{t \in T}$ is feasible, a contradiction. □

To prove Lemma 1, for each $n \in \mathbb{N}$, let $(\kappa_n^*, \lambda_n^*, \mu_n^*)$ be an optimal extreme point solution to the $1/n$-perturbed dual.[15] The feasible set is independent of the perturbation parameter and it contains only finitely many extreme points. Hence, one of them appears infinitely often as $n \to \infty$. Thus, by extracting a subsequence if necessary, we may assume that the sequence is constant. That is, there is an extreme point $(\kappa^*, \lambda^*, \mu^*)$ such that $(\kappa_n^*, \lambda_n^*, \mu_n^*) = (\kappa^*, \lambda^*, \mu^*)$ for all n. By Lemma A.1, the associated graph $G(\kappa^*)$ is acyclic.

We claim that $(\kappa^*, \lambda^*, \mu^*)$ is optimal in the 0-perturbed dual (4.6–4.9). To see this, note that the perturbed dual objective $f(\kappa, \lambda, \mu, \varepsilon) := \sum_{t \in T} \lambda_t U(w^t) + \sum_{t \in T} \mu_t - \varepsilon \sum_{t \in T} \sum_{s \neq t} \kappa_{st}$ is jointly continuous in $(\kappa, \lambda, \mu, \varepsilon)$. Therefore, given any feasible solution (κ, λ, μ), we have

$$f(\kappa, \lambda, \mu, 0) = \lim_n f(\kappa, \lambda, \mu, 1/n) \leq \lim_n f(\kappa^*, \lambda^*, \mu^*, 1/n) = f(\kappa^*, \lambda^*, \mu^*, 0),$$

where the inequality is because $(\kappa^*, \lambda^*, \mu^*)$ is optimal for all n by construction. We conclude that $(\kappa^*, \lambda^*, \mu^*)$ is an optimal solution to (4.6–4.9) such that $G(\kappa^*)$ is acyclic.

[15] To see that the feasible set of the ε-perturbed dual has an extreme point (and thus an extreme point that is optimal, because the optimal value is finite) even though there are free variables, let $\mu'_t := \min_{i \in I} p_t(y_i - w_i^t)$ for all $t \in T$. Then $(\kappa, \lambda, \mu) = (0, 0, \mu')$ is an extreme point.

A.2 Proof of Lemma 4

Let v be the optimal value of problem (5.6–5.9).

Step 1: $v \geq V(p)$. Let $\{e_t\}$ be an optimal solution to (5.6–5.9). We show first that we may assume without loss of optimality that $e_t \leq y_n$ for all $t \in T$. Suppose to the contrary that $e_t > y_n$ if and only if $t > \bar{t}$ for some cutoff type $\bar{t} \in T$. Consider reducing e_t to y_n for all $t > \bar{t}$. This lowers the value of the objective (5.6) at least weakly (and strictly if $p_t(1-\alpha_t) > 0$ for any $t > \bar{t}$). Because $e_{\bar{t}} \leq y_n$, the new solution satisfies the monotonicity constraint (5.8). It also satisfies the participation constraint (5.7): For any $t > \bar{t}$, the left-hand side is now

$$\begin{aligned}\sum_{s=1}^{\bar{t}}(\alpha_s - \alpha_{s-1})e_s + \sum_{s=\bar{t}+1}^{t}(\alpha_s - \alpha_{s-1})y_n &= \sum_{s=1}^{\bar{t}}(\alpha_s - \alpha_{s-1})e_s + (\alpha_t - \alpha_{\bar{t}})y_n \\ &\geq U^0(\alpha_{\bar{t}}) + (\alpha_t - \alpha_{\bar{t}})y_n \\ &\geq U^0(\alpha_{\bar{t}}) + U^0(\alpha_t) - U^0(\alpha_{\bar{t}}) = U^0(\alpha_t),\end{aligned}$$

where the first inequality follows by the $\bar{t}$-instance of (5.7), and the second inequality follows because—letting (e_t^0, c_t^0) denote the optimal choice of type t from A^0—we have

$$U^0(\alpha_t) - U^0(\alpha_{\bar{t}}) \leq \alpha_t e_t^0 - c_t^0 - (\alpha_{\bar{t}} e_t^0 - c_t^0) = (\alpha_t - \alpha_{\bar{t}})e_t^0 \leq (\alpha_t - \alpha_{\bar{t}})y_n.$$

We conclude that the new solution is optimal and it has $e_t \leq y_n$ for all t, as desired.

Define the costs $\{c_t\}$ recursively by setting $c_1 = 0$ and $c_{t+1} = \alpha_t(e_{t+1} - e_t) + c_t$ for $t = 1, \dots, k-1$. We will show that the solution $\{(e_t, c_t)\}$ so constructed is feasible in (5.1–5.5), which implies that $v = \sum_t p_t(1-\alpha_t)e_t \geq V(p)$.

Because $e_t \leq y_n$ for all t, constraints (5.8–5.9) and the construction of c_t imply constraints (5.4–5.5). Similarly, a straightforward calculation shows that (5.7) and the construction of c_t imply (5.3). Finally, $\{(e_t, c_t)\}$ satisfies all upward adjacent incentive compatibility constraints in (5.2) with equality by construction, i.e., $\alpha_t e_t - c_t = \alpha_t e_{t+1} - c_{t+1}$ for all $t < k$. This and the monotonicity constraint (5.8) imply the downward adjacent constraints $\alpha_{t+1}e_{t+1} - c_{t+1} \geq \alpha_{t+1}e_t - c_t$ for all $t < k$.[16] Together the upward and downward adjacent constraints then imply all the remaining constrains in (5.2).[17] We conclude that $\{(e_t, c_t)\}$ is

[16]This standard result follows from $\alpha_{t+1}(e_{t+1} - e_t) \geq \alpha_t(e_{t+1} - e_t) = c_{t+1} - c_t$, where the inequality uses monotonicity (and $\alpha_{t+1} > \alpha_t$) and the equality uses the equality in the upward adjacent constraint.

[17]This standard result follows from Milgrom and Shannon's (1994) Monotone Selection Theorem. To sketch a direct proof for the current setting, note that the constraints between any two adjacent types α_t and α_{t-1} imply that the allocation is monotone: $e_t \geq e_{t-1}$. Therefore, for any $s \geq t$, we have

$$0 \leq \alpha_t e_t - c_t - (\alpha_t e_{t-1} - c_{t-1}) \leq \alpha_s e_t - c_t - (\alpha_s e_{t-1} - c_{t-1}),$$

where the first equality follows by the downward adjacent constraint for type t, and the second is because feasible in (5.1–5.5) as we wanted to show.

Step 2: $v \leq V(p)$. The argument has two parts. We show first that problem (5.1–5.5) has an optimal solution $\{(e_t, c_t)\}$ that satisfies all upward adjacent incentive constraints in (5.2) with equality. Then we show that the allocations $\{e_t\}$ are feasible in (5.6–5.9), which implies that $v \leq \sum_t p_t(1-\alpha_t)e_t = V(p)$.

For the first part, let $\{(e_t, c_t)\}$ be the optimal solution to (5.1–5.5) that has the smallest number of upward adjacent incentive constraints hold with strict inequality. Suppose toward contradiction that this number is greater than zero. Then there is a type $s-1 \in T$ with

$$\alpha_{s-1}e_{s-1} - c_{s-1} > \alpha_{s-1}e_s - c_s. \tag{A.5}$$

We construct a new solution $\{(\hat{e}_t, \hat{c}_t)\}$ where $(\hat{e}_t, \hat{c}_t) = (e_t, c_t)$ for every type $t \neq s$, but where type s gets a new bundle $(\hat{e}_s, \hat{c}_s)$ defined as the unique solution to

$$\alpha_s e_s - c_s = \alpha_s \hat{e}_s - \hat{c}_s \quad \text{and} \quad \alpha_{s-1}e_{s-1} - c_{s-1} = \alpha_{s-1}\hat{e}_s - \hat{c}_s. \tag{A.6}$$

(That is, $(\hat{e}_s, \hat{c}_s)$ is the unique intersection of type $s-1$'s and type s's indifference lines through their old bundles.) Note that $\{(\hat{e}_t, \hat{c}_t)\}$ satisfies type $s-1$'s upward adjacent incentive constraint as an equality, and thus it has one less slack upward adjacent constraint than $\{(e_t, c_t)\}$. Moreover, we have $\hat{e}_s < e_s$, because subtracting the second equality in (A.6) from the first one gives

$$(\alpha_s - \alpha_{s-1})\hat{e}_s = \alpha_s e_s - c_s - (\alpha_{s-1}e_{s-1} - c_{s-1}) < \alpha_s e_s - c_s - (\alpha_{s-1}e_s - c_s) = (\alpha_s - \alpha_{s-1})e_s,$$

where the inequality follows by (A.5). Therefore, $\{(\hat{e}_t, \hat{c}_t)\}$ attains at least weakly lower value of the objective (5.1) than $\{(e_t, c_t)\}$ (and a strictly lower value if $p_s(1-\alpha_s) > 0$). This means that we have a contradiction to $\{(e_t, c_t)\}$ being the optimal solution with the smallest number of slack upward adjacent incentive constraints, if we show that $\{(\hat{e}_t, \hat{c}_t)\}$ is feasible.

To show feasibility, note first that the participation constraint (5.3) is satisfied because only type s's bundle was changed, but type s is indifferent between (e_s, c_s) and $(\hat{e}_s, \hat{c}_s)$ by construction. As for the incentive constraints in (5.2), it suffices to verify all upward and downward adjacent constraints (see footnote 17). The only affected types are then $s-1$, s, and $s+1$. Type $s-1$'s upward constraint holds with equality by construction. Type s's upward and downward constraints continue to hold because type s is indifferent between its old and new bundles. To see that type $s+1$'s downward adjacent constraint is satisfied, note

$e_t \geq e_{t-1}$ and $\alpha_s \geq \alpha_t$. Stringing together inequalities as t varies gives $\alpha_s e_s - c_s \geq \alpha_s e_{s-1} - c_{s-1} \geq \cdots \geq \alpha_s e_1 - c_1$, thus establishing all downward constraints for type s. Upward constraints are shown analogously.

that type s's indifference gives $0 = \alpha_s(e_s - \hat{e}_s) - (c_s - \hat{c}_s) < \alpha_{s+1}(e_s - \hat{e}_s) - (c_s - \hat{c}_s)$, where the inequality follows because $\alpha_{s+1} > \alpha_s$ and $\hat{e}_s < e_s$. Thus, type $s+1$ views type s's new bundle to be worse than type s's old bundle, which in turn is weakly worse than type $s+1$'s own bundle (which hasn't changed). Finally, note that the adjacent incentive constraints between types $s-1$ and s imply that $\hat{e}_s \geq e_{s-1} \geq 0$ and $\hat{c}_s \geq c_{s-1} \geq 0$, and above we argued that $\hat{e}_s < e_s \leq y_n$. Thus $\{(\hat{e}_t, \hat{c}_t)\}$ also satisfies (5.4) and (5.5), and it is therefore feasible, a contradiction. We conclude that there is an optimal solution to (5.1–5.5) that satisfies all upward adjacent constraints in (5.2) with equality.

To complete step 2, we verify that if $\{(e_t, c_t)\}$ is an optimal solution to (5.1–5.5) that satisfies all upward adjacent incentive constraints with equality, then $\{e_t\}$ is feasible in (5.6–5.9). The incentive constraints (5.2) imply the monotonicity constraint (5.8), as usual. The non-negativity constraint (5.9) is immediate from (5.5). Finally, using the binding upward adjacent incentive constraints to solve for c_t for $t > 1$ gives $\alpha_t e_t - c_t = \sum_{s=1}^{t}(\alpha_s - \alpha_{s-1})e_s - c_1$. Then (5.7) follows from (5.3), because $c_1 \geq 0$ by (5.5).

A.3 Proof of Lemma 6

Fix a random linear contract p with $\mathrm{supp}(p) = \{\alpha_1, \ldots, \alpha_k\}$, and label the slopes such that $0 \equiv \alpha_0 \leq \alpha_1 < \cdots < \alpha_k$. Suppose the guarantee $V(p)$ is positive (otherwise the claim is trivial). Note that $V(p)$ is weakly less than the best guarantee obtained by optimizing the probabilities of the contracts in the support of p, i.e.,

$$V(p) \leq v := \max\{V(p) : p \in \Delta(\mathbb{R}^n_+) \text{ and } \mathrm{supp}(p) \subset \{\alpha_1, \ldots, \alpha_k\}\}.$$

Following the beginning of the proof of Proposition 5 without the substitutions $\alpha_t = \frac{t}{k}$ and $\alpha_t - \alpha_{t-1} = k^{-1}$ (which do not apply to our "irregular grid") shows that v is the optimal value to the following dual pair of linear programs, which differ from their analogs (5.18–5.21) and (5.22–5.26) only in that we have imposed an upper bound on e_t; we will see momentarily that this bound is without loss of optimality. Writing $\theta_0 = \theta_k \equiv 0$, the programs are:

Primal (P)

$$\max_{p,\lambda,\theta,\varphi} \sum_{t\in T} \lambda_t U^0(\alpha_t) + y_n \sum_{t\in T} \varphi_t \tag{A.7}$$

$$\text{s.t.} \quad \sum_{s=t}^{k} \lambda_s(\alpha_t - \alpha_{t-1}) - \theta_t + \theta_{t+1} - \varphi_t - p_t(1-\alpha_t) \le 0 \qquad \forall t \in T, \qquad (e_t) \tag{A.8}$$

$$\sum_{t\in T} p_t = 1, \qquad (\delta) \tag{A.9}$$

$$p_t, \lambda_t, \varphi_t \ge 0 \qquad \forall t \in T. \tag{A.10}$$

Dual (D)

$$\min_{\{e_t\},\delta} \delta \tag{A.11}$$

$$\text{s.t.} \quad \delta - (1-\alpha_t)e_t \ge 0 \qquad \forall t \in T, \qquad (p_t) \tag{A.12}$$

$$\sum_{s=1}^{t} (\alpha_s - \alpha_{s-1})e_s \ge U^0(\alpha_t) \qquad \forall t \in T, \qquad (\lambda_t) \tag{A.13}$$

$$e_{t+1} - e_t \ge 0 \qquad \forall t \in T \setminus \{k\}, \qquad (\theta_t) \tag{A.14}$$

$$-e_t \ge -y_n \qquad \forall t \in T, \qquad (\varphi_t) \tag{A.15}$$

$$e_t \ge 0 \qquad \forall t \in T. \tag{A.16}$$

We observe first that if $\alpha_1 = 0$, then there is an optimal solution with $p_1 = 0$. To see this, note that setting $e_1 = 0$ is without loss of optimality in (D): This relaxes (A.12) and (A.14), and it does not affect (A.13) because $(\alpha_1 - \alpha_0)e_1 = 0e_1 = 0$ for any $e_1 \ge 0$. Because $\delta = V(p) > 0$, constraint (A.12) and complementary slackness then implies that $p_1 = 0$. We may thus assume in what follows that $\alpha_1 > 0$.

We argue then that the shadow prices on the monotonicity constraints (A.14) and the upper bounds (A.15) are zero. Ignore (A.15) momentarily. Fixing δ at its optimal value, we can without loss of optimality set e_t to satisfy (A.12) at equality:[18] This relaxes (A.13) and satisfies (A.14) with strict inequality. A straightforward but tedious calculation then verifies that this solution can be further modified by setting $e_t = \max\{\delta(1-\alpha_t)^{-1}, y_n\}$ for all t without affecting the objective.[19] We conclude that $\theta_t = 0$ and $\varphi_t = 0$ for all t.

The upshot is that the (nonempty) sets of optimal solutions to the above dual programs are bounded. For (D) this is obvious because e_t is bounded by feasibility and δ is unique by optimality. For (P), note that p_t is bounded by feasibility and above we argued that

[18] If $\alpha_k = 1$, then this is impossible for $t = k$, but instead we may set $e_k = \max\{e_{k-1} + 1, y_n + 1\}$.

[19] See the beginning of the proof of Lemma 4 for an analogous argument.

$\theta_t = \varphi_t = 0$ for all t. Then (A.8) bounds $\lambda_t \geq 0$ from above for all t (because $\alpha_1 > 0$).

Because the sets of optimal solutions to (P) and (D) are nonempty and bounded, it follows from Theorem 1 in Robinson (1977) that the optimal value v is continuous in the slopes α_t in some open neighborhood of $(\alpha_1, \dots, \alpha_k)$. It is then clear that given any $\varepsilon > 0$, any sufficiently fine grid Γ satisfies $v(\Gamma) \geq v - \varepsilon \geq V(p) - \varepsilon$ as desired.

A.4 Proof of Lemma 7

Fix a sequence (t_k) as in the statement of the lemma. Let $\alpha_k := t_k/k$. Then

$$\sum_{s=1}^{t_k} \frac{1}{k-s} = \sum_{s=1}^{\alpha_k k} \frac{1}{k-s} = \sum_{n=(1-\alpha_k)k}^{k-1} \frac{1}{n} = \sum_{n=1}^{k-1} \frac{1}{n} - \sum_{n=1}^{(1-\alpha_k)k-1} \frac{1}{n} = H_{k-1} - H_{(1-\alpha_k)k-1},$$

where H_ℓ is a harmonic number, i.e., the sum of the first ℓ terms of the harmonic series. It can be approximated using the Euler-Maclaurin formula $H_\ell = \ln(\ell) + \gamma + \varepsilon_\ell$, where $\gamma \approx 0.57772$ is the Euler-Mascheroni constant and ε_ℓ is an error term that converges to zero as $\ell \to \infty$. This implies that, for some error term $\tilde{\varepsilon}_k$ convergent to zero as $k \to \infty$, we have

$$\sum_{s=1}^{t_k} \frac{1}{k-s} = \ln\left(\frac{k-1}{(1-\alpha_k)k-1}\right) + \tilde{\varepsilon}_k. \tag{A.17}$$

Suppose first that $\lim \alpha_k = \lim \frac{t_k}{k} = \alpha < 1$. Fix $\eta > 0$ small enough so that $\alpha + \eta < 1$. Note that for any k large enough so that $\alpha_k \in (\alpha - \eta, \alpha + \eta)$, we have

$$\frac{k-1}{(1-\alpha+\eta)k-1} \leq \frac{k-1}{(1-\alpha_k)k-1} \leq \frac{k-1}{(1-\alpha-\eta)k-1}.$$

Applying Stolz-Cesàro theorem to the bounds (which are ratios of monotone divergent sequences) squeezes the expression in the middle to the interval $[(1-\alpha+\eta)^{-1}, (1-\alpha-\eta)^{-1}]$ as $k \to \infty$. As $\eta > 0$ was arbitrary, this shows that the argument of $\ln(\cdot)$ in (A.17) converges to $(1-\alpha)^{-1}$ as $k \to \infty$. Sending $k \to \infty$ on both sides of (A.17) then shows that the sum on the left converges to $-\ln(1-\alpha)$ by continuity of $\ln(\cdot)$.

Suppose then that $\lim \alpha_k = \lim \frac{t_k}{k} = \alpha = 1$. By moving to a subsequence if necessary, we may assume that (α_k) is monotone. Then, for any fixed $\bar{k}$ and any $k \geq \bar{k}$, we have

$$\frac{k-1}{(1-\alpha_k)k-1} \geq \frac{k-1}{(1-\alpha_{\bar{k}})k-1}.$$

Therefore,

$$\liminf_{k\to\infty} \frac{k-1}{(1-\alpha_k)k-1} \geq \liminf_{k\to\infty} \frac{k-1}{(1-\alpha_{\bar{k}})k-1} = \frac{1}{1-\alpha_{\bar{k}}} \to \infty \quad \text{as } \bar{k} \to \infty,$$

where the equality is by the Stolz-Cesàro theorem. We conclude that the right-hand side of (A.17) diverges to ∞, and so does the sum on the left.

References

Antic, N. (2021): "Contracting with Unknown Technologies," Discussion paper, Northwestern University.

Antic, N., and G. Georgiadis (2023): "Robust Contracts: A Revealed Preference Approach," Discussion paper, Northwestern University.

Brooks, B., and S. Du (2024): "On the Structure of Informationally Robust Optimal Mechanisms," *Econometrica*, forthcoming.

Burkett, J., and M. Rosenthal (2023): "Data-Driven Contract Design," Discussion paper, Georgia Institute of Technology.

Carroll, G. (2015): "Robustness and Linear Contracts," *American Economic Review*, 105(2), 536–563.

——— (2019): "Robustness in Mechanism Design and Contracting," *Annual Review of Economics*, 11, 139–166.

Carroll, G., and L. Bolte (2023): "Robust contracting under double moral hazard," *Theoretical Economics*, 18(4), 1623–1663.

Dai, T., and J. Toikka (2022): "Robust Incentives for Teams," *Econometrica*, 90(4), 1583–1613.

Ellsberg, D. (1961): "Risk, Ambiguity, and the Savage Axioms," *The Quarterly Journal of Economics*, pp. 643–669.

Holmström, B., and P. Milgrom (1987): "Aggregation and Linearity in the Provision of Intertemporal Incentives," *Econometrica*, pp. 303–328.

Kambhampati, A. (2023): "Randomization Is Optimal in the Robust Principal-Agent Problem," *Journal of Economic Theory*, 207, 105585.

——— (2024): "Robust Performance Evaluation of Independent Agents," *Theoretical Economics*, forthcoming.

Ke, S., and Q. Zhang (2020): "Randomization and Ambiguity Aversion," *Econometrica*, 88(3), 1159–1195.

Lai, H.-C., and S.-Y. Wu (1992): "On Linear Semi-Infinite Programming Problems: An Algorithm," *Numerical Functional Analysis and Optimization*, 13(3–4), 287–304.

Marku, K., S. Ocampo Diaz, and J.-B. Tondji (2024): "Robust Contracts in Common Agency," *The Rand Journal of Economics*, forthcoming.

Milgrom, P., and C. Shannon (1994): "Monotone Comparative Statics," *Econometrica*, 62(1), 157–180.

Peng, B., and Z. G. Tang (2024): "Optimal Robust Contract Design," Discussion paper, Shanghai University of Finance and Economics.

Raiffa, H. (1961): "Risk, Ambiguity, and the Savage Axioms: Comment," *The Quarterly Journal of Economics*, 75(4), 690–694.

Robinson, S. (1977): "A Characterization of Stability in Linear Programming," *Operations Research*, 25(3), 435–447.

Rosenthal, M. (2023): "Robust Incentives for Risk," *Journal of Mathematical Economics*, 109, 102893.

Saito, K. (2015): "Preferences for Flexibility and Randomization Under Uncertainty," *American Economic Review*, 105(3), 1246–71.

Walton, D., and G. Carroll (2022): "A General Framework for Robust Contracting Models," *Econometrica*, 90(5), 2129–2159.