Data-Driven Contract Design

Justin Burkett[*] and Maxwell Rosenthal[†]

July 30, 2024

Abstract

We propose a prior-free model of incentive contracting in which the principal's beliefs about the agent's production technology are characterized by revealed preference data. The principal and the agent are each financially risk neutral and the agent's preferences are understood to be quasilinear in effort. Prior to contracting with the agent, the principal observes the output produced by a population of identical agents in best response to finitely many exogenously-specified contracts. She views any technology that rationalizes this data as plausible and evaluates contracts according to their guaranteed expected payoff against the set of all such technologies. This paper does four things. First, we characterize the set of technologies that are consistent with the revealed preference data. Second, we show that robustly optimal contracts are either empirical contracts or *equity bonus contracts* that supplement mixtures of contracts from the data with equity payments. Third, we provide conditions under which these optimal contracts append equity payments to only a single contract from the data. Fourth, and finally, we show that all of our results generalize without complication to a setting in which there might be arbitrary forms of unobserved heterogeneity within the population of agents.

1 Introduction

Moral hazard is the workhorse model of financial incentive provision and therefore enjoys many important practical applications. While a rich theoretical literature has delivered many interesting and important insights about properties that well designed incentive schemes ought to have (Holmström [2017]), the optimal contracts prescribed by these models are contingent on information that principals are unlikely to possess in practice. Unlike in other

[*]Georgia Institute of Technology; justin.burkett@gatech.edu.
[†]Georgia Institute of Technology; rosenthal@gatech.edu.

areas of mechanism design in which a theoretical pricing exercise might be supported by the extensive literature on demand estimation or an auction design problem by the empirical auctions literature, operationalizing the design of incentive contracts faces some special challenges. In the classical model (Holmström [1979]), for instance, optimal contracts are highly sensitive to the likelihood ratios induced by the output distributions associated to actions that are by assumption unobserved. Even in weakly structured models (Chassang [2013], Carroll [2015]) that are less informationally demanding, it is still presumed that the principal has at least some prior knowledge about the relationship between effort and output, even though she again never observes effort. How might an analyst give practical and data-driven advice to a real-world principal on how to design contracts then?

This paper studies a moral hazard problem in which the principal possesses data X describing the performance of one or more exogenously specified contracts. The *observations* (F_i, w_i) in her data consist of a contract w_i and the distribution of output F_i produced by a population of identical agents[1] in best response to w_i. The agent is risk neutral, suffers additive disutility from unobserved effort, and enjoys limited liability. The principal understands the agent's preferences but observes neither the effort cost e_i associated with F_i nor the identity of other actions that might also be available to the agent. She insists on an objective analysis and therefore incorporates only the information that can be inferred from direct observation into her designs. However, because she observes best responses to only finitely many contracts, the principal's model of the agent's technology is only partially identified. In accordance with her desire to act on objective information about the agent's capabilities, the principal evaluates contracts according to their guaranteed payoff against the set of technologies that rationalize the data.

There are two avenues open to the principal in pursuit of robust optimality. First, the principal might pay the agent with an empirical contract. Because the agent accepts recommendations from the principal whenever he is indifferent between two or more actions, the worst-case for an *empirical contract* w_i is simply the best response F_i recorded in the data. Second, the principal might pay the agent with a novel contract. Because of the principal's permissive model of the agent's technology, it is not prima facie obvious that such a contract should provide any non-trivial payoff guarantee whatsoever, much less an

[1] While we interpret our model as one in which the revealed preference data are generated by a population of agents, our formal model is one of interaction between a principal and a single representative member of that population. Accordingly, we use singular language when referring to the agent throughout most of the paper. In Section 7, we model heterogeneity within the agent population as stochasticity in the representative agent's technology. We discuss the interpretation of the data in both that section and in the conclusion.

optimal one. In fact, that intuition turns out to be overly pessimistic, because the data are indeed informative about the agent's best response to novel contracts. As we show, there are interesting cases of our model in which novel contracts not only provide non-trivial guarantees but also robustly outperform all empirical contracts.

Our analysis is structured as follows. First, we characterize the rationalizability of data sets that contain finitely many contract–action pairs. Because effort costs are additive, a system of linear inequalities bound the unobserved effort costs e corresponding to the observed output distributions F chosen in best response to each contract w. By analogy to a classical problem from network theory, we give a simple necessary and sufficient condition for there to exist effort costs that rationalize these finite data sets. We then apply this result to (i) characterize the rationalizability of the principal's data X and (ii) characterize the rationalizability of the appended data set $X \cup \{(F, w)\}$ for observations (F, w) that do not appear in the principal's data. In turn, we obtain not only a tight condition under which the principal's problem is well defined, but also a tight characterization of incentive compatibility for novel contracts.

In the second part of our analysis, we use our characterization of incentive compatibility to reduce our model to a robust contracting problem in the spirit of Carroll [2015]. We develop a multidimensional version of Carroll's support line argument and use that argument to prove that optimal contracts are either empirical contracts or *equity bonus contracts* that pay the agent with a convex combination of one or more empirical contracts plus a share of the principal's profit. These contracts are optimal because they leverage both the revealed preference information in the data and also the desirable preference-alignment properties of linear contracts. We show via a simple two-observation example in Section 5 that there are interesting cases of our model in which optimal contracts are indeed novel.

Third, we study the relationship between the form of the optimal equity bonus contract and the characteristics of the revealed preference data X. We confirm via example that the optimal contract is sometimes a convex combination of multiple empirical contracts plus an equity payment and is therefore at least somewhat complex. We then show that this complexity is related to what the data reveal about the quasiconcavity of the agent's payoffs. Specifically, if the data are supermodular in an appropriate sense, then two related facts follow. First, the agent's preferences over observed actions are provably quasiconcave. Second, and more importantly, optimal equity bonus contracts add an equity payment to a single contract from the data. In addition to simplicity, this result lends tractability to our problem because the optimal *one-contract mixture* equity bonus contract is completely characterized by an adaptation of a straightforward calculation that appears in Chassang [2013] and Carroll [2015]. This lies in contrast to the general case, where the optimal contract can not in general be characterized by first-order conditions.

Fourth, and finally, we relax our assumption that the data are generated by a population of identical agents and allow for arbitrary types of unobserved heterogeneity within that population. We show via elementary arguments that the guarantee provided by any contract in this heterogeneous-agent environment is identical to the guarantee that it provides in the homogeneous-agent environment considered throughout the preceding parts of the paper. In doing so, we verify not only that all of our preceding results hold in the heterogeneous-agent environment but also confirm more broadly that unobserved agent heterogeneity is inconsequential for data-driven robust contract design.

Relation to Literature This work continues our earlier studies of data driven prior-free mechanism design. First and foremost, our model is an application of the revealed preference robust mechanism design framework developed in Rosenthal [2019] to the study of moral hazard (rather than multidimensional screening, as in that paper). Elsewhere, Burkett and Rosenthal [2022] provides a statistical theory of simple contracts for risk-averse agents in a model that emphasizes finite sample issues outside of the scope of the present paper.

Our model might be viewed as an alternative interpretation of the robust contracting exercise developed by Carroll [2015] and elaborated on in Dai and Toikka [2022], Kambhampati [2023], and elsewhere. Our paper is distinguished from Carroll's analysis by the structure of our informational environment. In particular, the principal in Carroll's model knows (i) the output distributions and effort costs associated with a known subset of the agent's actions and (ii) lower bounds on effort costs for potential unknown actions. She otherwise entertains the possibility that the agent's technology is any superset of the known actions, without restriction. In our model, the principal's knowledge of the agent's technology is instead characterized by revealed preference data, and we impose no exogenous bounds on relative effort costs.

Less closely related is other work in the robust mechanism design space, including Carroll [2017]. There, the author studies a robust multidimensional screening problem in which the principal knows the marginal distribution of each dimension of the agent's private information, but not the joint distribution. The optimal mechanism is a "no-bundling" solution that concatenates the solutions to each of the one-dimensional problems. Setting aside any subtleties that might arise from observing lower bounds rather than exact valuations, there is an interpretation of Carroll's model as one in which the principal infers the marginal distributions from a family of non-panelled cross-sectional data sets. Under that interpretation, the author's no-bundling solution is analogous to a counterfactual result for our setting under which optimal contracts must be empirical. This lies in contrast to our actual solution, in which optimal contracts are sometimes novel. At the same time, our principal possesses "more" information than Carroll's, because she draws information from one panelled data set rather than a series of independent cross-sectional data sets.This is true by definition in the homogeneous-agent environment and true "on path" in the heterogeneous environment, wherein worst-case agent populations are indeed homogeneous.

Along separate lines, this paper is conceptually related to a similarly-motivated study of data-driven incentive design by Georgiadis and Powell [2022]. As in our paper, the principal in their model has access to contract-output data and uses this data to make inferences about the performance of novel contracts. The authors identify payoff-improving perturbations of existing contracts under the hypothesis that the agent chooses from an unknown action set that is known by the principal to be differentiable. Aside from differences in our model and our solution concept, there are two main distinctions between our analysis and theirs. First, while their analysis is local in nature, we identify globally optimal contracts. Second, while the authors later show how to extend their analysis to global optimality under particular assumptions about the agent's technology, our model imposes no structure on the agent's technology beyond rationalizability of the principal's data, even when we allow for heterogeneity in the agent population. In contrast, their smoothness assumption is substantive rather than technical, given the nature of the authors' analysis.

Finally, the concurrent paper Antic and Georgiadis [2022] studies a similar model with intersecting results. Present in their paper but absent from ours is (i) a detailed study of the two-observation problem, with a direct proof that one-contract mixtures are optimal for those cases; and (ii) an interesting empirical exercise in which the authors apply their analytical results to an experimental data set from elsewhere in the literature. Present in our paper but lacking in theirs is (i) our first theorem on the complete characterization of the set of rationalizing technologies; (ii) our third theorem on the optimality of one-contract mixtures for supermodular data, which includes not only two-observation data sets as a special case but also a much richer class of problems; and (iii) our fourth theorem on heterogeneous agent populations.

The paper is structured as follows. First, we introduce the model in Section 2. Next, we treat the revealed preference problem in Section 3, which yields a condition for rationalizing the data, as well as any previously unobserved action. This condition is then used to set up and solve the robust contracting problem in Section 4. Section 5 discusses conditions under which the principal can find a novel contract that improves upon the contracts in the data. Section 6 considers special cases of the model optimal contracts take an especially simple form, including most importantly the aforementioned supermodular case. Lastly, we extend the model to allow for heterogeneous agent data in Section 7. Section 8 concludes. Proofs and supporting technical material appear in the Appendix.

2 Model

2.1 Actions and payoffs

The agent chooses *action* (F, e) from production *technology* $\mathcal{A}$ and produces *output* y in set $\mathcal{Y} \subset \mathbb{R}$.[2] Output is distributed according to $F \in \Delta(\mathcal{Y})$ and the agent suffers unobserved effort cost $e \in \mathbb{R}$. The principal pays the agent with a continuous *contract* $w : \mathcal{Y} \to \mathbb{R}_+$ to provide incentives for effort. The agent is a risk-neutral expected utility maximizer with quasilinear effort costs and the principal is herself financially risk-neutral. Accordingly, given action (F, e) and contract w, the agent's payoff and the principal's payoff are respectively

$$\int_{\mathcal{Y}} w(y)\mathrm{d}F - e \qquad\qquad \int_{\mathcal{Y}} (y - w(y))\mathrm{d}F.$$

Ties in the agent's problem are broken in favor of the principal. Given technology $\mathcal{A}$ and contract w, we write

$$c(\mathcal{A}|w) = \underset{(F,e)\in\mathcal{A}}{\arg\max}\left(\int_{\mathcal{Y}} w\mathrm{d}F - e\right) \qquad \Pi(\mathcal{A}|w) = \max_{(F,e)\in c(\mathcal{A}|w)} \int_{\mathcal{Y}} (y - w(y))\mathrm{d}F$$

for the agent's preferred actions and the principal's resulting payoff, respectively.

2.2 The production technology

The principal's beliefs about the agent's technology $\mathcal{A}$ are characterized by revealed preference data. We call a pair (F, w) consisting of output distribution F and contract w an

[2] In the interest of brevity, we occasionally refer to output distributions F as "actions", without explicit reference to the associated effort cost.

observation, and suppose that the principal observes a finite set X of observations prior to contracting with the agent.

Definition 1. Technology $\mathcal{A}$ *rationalizes* X if there exists a map $e : X \to \mathbb{R}$ such that $(F, e(F, w)) \in c(\mathcal{A}|w)$ for each $(F, w) \in X$.

The principal entertains the possibility that the agent's technology is any technology $\mathcal{A}$ which rationalizes X. We make three clarifying comments about the relationship between X and its rationalizing technologies $\mathcal{A}$:

1. If $\mathcal{A}$ rationalizes the data X and $(F, w) \in X$ for some contract w, then $(F, e) \in \mathcal{A}$ for some effort e.

2. If X includes distinct observations that yield the same output distribution F, then there might be fewer actions in $\mathcal{A}$ than observations in X.

3. If $\mathcal{A}$ includes actions (F, e) with output distributions F that are not constituent to any observation $(F, w) \in X$, then there might be more actions in $\mathcal{A}$ than observations in X.

We write $\mathscr{A}$ for the set of technologies $\mathcal{A}$ that rationalize X. It is sometimes convenient to index the data, in which case we write (F_i, w_i) to refer to a typical observation in the data and e_i for the value of a particular rationalizing effort cost for F_i. We reserve the subscripted notation F_i, w_i, e_i for this purpose.

2.3 The principal's objective

The principal's objective is to choose the contract w that maximizes her payoff guarantee

$$\Pi(w) = \inf_{\mathcal{A} \in \mathscr{A}} \Pi(\mathcal{A}|w).$$

2.4 Formalities and notation

The set of real numbers $\mathbb{R}$ has the Euclidean topology. The set of outputs $\mathcal{Y}$ is compact and contracts are continuous. We normalize $\min \mathcal{Y} \equiv 0$ and write $\delta(x)$ for the degenerate distribution at x. We occasionally write $\langle F, w \rangle$ for the expectation of contract w against output distribution F. When stating examples with finitely many outputs, we interpret the set of outputs $\mathcal{Y}$, output distributions in $\Delta(\mathcal{Y})$, and contracts w as ordered tuples.

We normalize each contract w associated with an observation $(F, w) \in X$ by supposing that the minimum payment assigned by w is 0. This is without loss of generality,[3] and we do not require $w(0) = 0$. Furthermore, we restrict attention to production technologies $\mathcal{A}$ with only finitely many actions, noting that this is again without loss of generality because (i) the principal's data X includes only finitely many observations and (ii) the principal's guarantee $\Pi(w)$ for any particular contract w is achieved by some technology $\mathcal{A}$ that includes at most $|X| + 1$ actions. Finally, we streamline some definitions and proofs by suppressing edge cases in which none of the observations in the data are profitable for the principal.

Assumption 1. The revealed preference data X contain at least one observation (F, w) with $\langle F, y - w \rangle > 0$.

Were every observation in the data to yield a non-positive payoff, it might sometimes be optimal for the principal to pay the agent with the linear bonus contract $w = y$ that pays the agent all of the output that he produces. Aside from ruling out those pathological data sets, Assumption 1 does not affect our results.

3 Revealed preference

This section considers the question of which data sets can be rationalized by some assignment of effort costs. The characterization in Theorem 1, which essentially reduces to an assertion that a particular set of linear inequalities has a solution, serves two purposes for the principal. First, it exactly characterizes what is meant by rationalizability of the data in this setting. More importantly, we use Theorem 1 to characterize the set of rationalizable output distributions at *any* contract the principal could offer the agent. This characterization significantly reduces the complexity of principal's optimal contracting problem, which we return to in Section 4.

The problem of identifying which finite sets of observations are rationalizable is equivalent to determining whether there exists an assignment of effort costs under which all of the agent's incentive compatibility conditions can jointly be satisfied. The map e *rationalizes* X if

$$\langle F - G, w \rangle \geq e(F, w) - e(G, v) \quad \forall (F, w), (G, v) \in X. \tag{1}$$

[3] The set of technologies that rationalize observation (F, w) is identical to the set of technologies that rationalize observation $(F, w + \kappa)$ for any constant κ. Accordingly, the normalizations $\min w_i = 0$ are for notational convenience only.

This problem is itself isomorphic to a classical network theory problem. To make the connection, let $\langle F - G, w\rangle$ be the length of an arc from the node (G, v) to the node (F, w). A finite sequence of observations $P = \{(H_1, w_1), \ldots, (H_n, w_n)\}$ is a *path* through the data with length

$$\ell(P) = \sum_{i=1}^{n-1} \langle H_{i+1} - H_i, w_{i+1}\rangle$$

A path $C = \{(H_1, w_1), \ldots, (H_n, w_n), (H_1, w_1)\}$ that terminates at its starting point is a *cycle*.

There is nothing unusual about data sets X in which some paths have negative length, since an alternative action may give the agent a higher expected wage at a much higher effort cost; however, it is critical to the revealed preference problem that there are no cycles with strictly negative length. To understand why, consider some cycle $C = \{(H_1, w_1), \ldots, (H_n, w_n), (H_1, w_1)\}$. To each arc in the cycle, there is an associated inequality from (1) given by $\langle H_{i+1} - H_i, w_{i+1}\rangle \geq e(H_{i+1}, w_{i+1}) - e(H_i, w_i)$. Summing up these inequalities over all arcs in the cycle implies $\ell(C) \geq 0$. Therefore, all cycles must have non-negative length if (1) is to have a solution. The converse is also true: if all cycles having non-negative length, then there exists a solution to (1).

Theorem 1. *The following statements are equivalent:*

1. The data X are rationalizable.

2. All cycles in X have non-negative length.

3. There is a shortest path between any two observations in X.

We give a brief proof of Theorem 1 in the appendix, drawing on the network-theoretic framework developed in Vohra [2011]; there are several closely related results in the literature on revealed preference and rationalizability, including Afriat [1967] and Rochet [1987]. We make use of this framework and its terminology in subsequent results.

As the result shows, the lengths of the paths and cycles in the data X determine its rationalizability. Both of these objects have economic interpretations. For any pair of observations, the length of the shortest path from one to the other bounds the effort cost difference between the two associated actions. We write s_{ij} for the length of the shortest path in X from observation (F_i, w_i) to observation (F_j, w_j). It follows from the argument preceding Theorem 1 that s_{ij} is a tight upper bound for the difference in effort costs $e_j - e_i$ associated to actions F_j and F_i.

The length of a cycle in the data then indicates the amount of slack available for constructing effort costs that are consistent with the constituent inequalities in the cycle. Recall that each arc in the data is associated with an inequality in (1). The length of a cycle then corresponds to value taken by the left-hand side when these inequalities are summed over all constituent arcs. The right-hand side of the sum is zero if it is a cycle. In the extreme case, a zero-length cycle indicates that all inequalities must hold with equality, and hence that the agent *must* be indifferent along the cycle. A cycle with strictly positive length, on the other hand, allows for an assignment of effort costs in which some (or all) of the constituent inequalities are slack.

Aside from offering a pair of interpretations of rationalizability, Theorem 1 plays two major roles in our analysis. First, it yields the tightest-possible condition under which there are technologies that rationalize the principal's data X.

Corollary 1. *If all cycles in the data X have non-negative length, then the set of technologies which rationalize the principal's data X is non-empty.*

Second, Theorem 1 yields a characterization of what output distributions F would be incentive compatible for the agent in response to an arbitrary contract w. In particular, the theorem implies that the appended data set $X \cup \{(F, w)\}$ is rationalizable if and only if every cycle in $X \cup \{(F, w)\}$ itself has non-negative length. Equivalently, we obtain the condition

$$\langle F_i - F, w_i \rangle + s_{ij} + \langle F - F_j, w \rangle \geq 0 \quad \forall i, j,$$

where we remind the reader that we write $(F_i, w_i), (F_j, w_j)$ to indicate observations in X.

Corollary 2. *If all cycles in the data X have non-negative length, then output distribution F is rationalizable at contract w given data X if and only if*

$$\langle F, w - w_i \rangle \geq V_i(w) \quad \forall i, \tag{2}$$

where $V_i(w) \equiv \max_j \langle F_j, w - w_i \rangle - \Delta(i, j)$ and $\Delta(i, j) \equiv s_{ij} + \langle F_i - F_j, w_i \rangle \geq 0$ is the length of the shortest cycle in the data that includes the arc from j to i.

Throughout the paper, we refer to the quantities $w - w_i$ as *supplemental payments*. Each of the constants $\Delta(i, j)$ is fixed by the data and has a natural interpretation. First, note that because $\Delta(i, j)$ is the length of a cycle it is nonnegative if the data can be rationalized.

Next, notice that rationalizability requires that

$$\langle F_i - F_j, w_i \rangle \geq e_i - e_j \geq -s_{ij},$$

where the second inequality is a lower bound on $e_i - e_j$ determined by a shortest path from i to j. The difference $\Delta(i, j) = \langle F_i - F_j, w_i \rangle + s_{ij}$ is therefore the largest rationalizable gain in payoff that the agent could realize from picking (F_i, e_i) over (F_j, e_j) at w_i. Equivalently, it is the largest loss the agent could experience from being forced to select (F_j, e_j) at w_i. This shows up as a cost for the principal if she wants the agent to select an action like F_j at w_i. The $\Delta(i, j)$ terms play a key role in the construction of novel contracts, and we further develop the intuition for their importance below.

As is consistent with Theorem 1, we adopt the following condition and maintain it throughout the rest of the paper, without explicit citation. A cycle is *non-degenerate* if it passes through at least two distinct observations.

Assumption 2. All non-degenerate cycles in the data X have strictly positive length.

Assumption 2 is sufficient for rationalizability and also rules out — via its exclusion of data sets with zero-length cycles — knife-edge cases wherein there are provable indifferences in the data, so that for example F_j is guaranteed incentive compatible at w_i for distinct observations i, j.[4]

Going forward, we say distribution F is *rationalizable* at contract w if the data set $X \cup \{(F, w)\}$ is rationalizable, meaning there exists a rationalizing technology $\mathcal{A} \in \mathscr{A}$ and an effort cost e such that $(F, e) \in c(\mathcal{A}|w)$. These distributions are the candidate worst-case distributions for novel contracts.

Finally, before proceeding, we show through example that one cannot simplify the non-negative-cycles condition to a pairwise condition.

Example 1. Suppose that there are four output states and consider the three-observation data set X with

$$w_1 = (3, 0, 4, 0) \qquad w_2 = (4, 3, 0, 0) \qquad w_3 = (0, 4, 3, 0)$$
$$F_1 = (1/2, 0, 0, 1/2) \qquad F_2 = (0, 1/2, 0, 1/2) \qquad F_3 = (0, 0, 1/2, 1/2).$$

[4] In earlier versions of this paper we allowed for zero-length cycles in the data. While allowing for these knife edge special cases does not change any of our results, it complicates both proofs and exposition to the potential disservice of the reader. Details are available from the authors.

Now, the pairwise conditions $\langle F_i, w_i - w_j \rangle \geq \langle F_j, w_i - w_j \rangle$ are apparently necessary for rationalizability. These conditions, which are satisfied for all i, j in X, suggest the effort bounds

$$\begin{aligned} 1/2 &\leq e_1 - e_2 \leq 3/2 \\ 1/2 &\leq e_2 - e_3 \leq 3/2 \\ -3/2 &\leq e_1 - e_3 \leq -1/2. \end{aligned}$$

However, summing up the pairwise bounds for $1, 2$ and $2, 3$ yield

$$1 \leq e_1 - e_3 \leq 3,$$

which is evidently inconsistent with the pairwise $1, 3$ bound. This failure of rationalizability is captured in our condition by for example the cycle $1 \to 2 \to 3 \to 1$, which has length

$$\langle F_2 - F_1, w_2 \rangle + \langle F_3 - F_2, w_3 \rangle + \langle F_1 - F_3, w_1 \rangle = -3/2 < 0.$$

Thus, while the pairwise conditions are indeed necessary, they are insufficient. Instead, the more complex cycle conditions in our characterization must also be satisfied.

4 Optimal contracts

This section provides our general results on the optimality of empirical contracts and equity bonus contracts. First, in Proposition 1, Lemma 1, and Theorem 2, we generalize the support line argument developed in Carroll [2015] to show that optimal novel contracts are equity bonus contracts. Second, in Proposition 2, we show that empirical contracts and equity bonus contracts are uniquely optimal when the output distributions in the principal's data have full support. Third, and finally, we use that proposition to show that our optimal contracts are not weakly dominated by other types of contracts.

It is convenient to first restrict attention to novel contracts that improve upon the contracts in the data.

Definition 2. Contract w is *eligible* if $\Pi(w) > \Pi(w_i)$ for every empirical contract w_i.

In the first step towards identifying the structure of the eligible contracts, we show that the principal's guarantee $\Pi(w)$ can be written as a minimization problem using the inequalities in (2) as constraints. We also establish that at least one of the inequalities must bind at the worst-case F.

Proposition 1. *If the contract w is eligible then*

$$\Pi(w) = \min_{F \in \Delta(\mathcal{Y})} \langle F, y - w \rangle \text{ s.t. } \langle F, w - w_i \rangle \geq V_i(w) \quad \forall i$$
$$\langle F, w - w_i \rangle = V_i(w) \quad \exists i. \tag{3}$$

Proposition 1 serves two purposes. First, the proposition treats the tension between the weak inequalities in the constraints of program (3) and the strict inequalities suggested by our assumption that ties are broken in favor of the principal. This requires special care in boundary cases; details are left to the Appendix for interested readers. Second, by confirming that the worst-case distribution for any eligible contract lies on the boundary of the constraint set, the proposition justifies our use of the separating hyperplane theorem to assign multipliers to the inner minimization problem in the principal's maxmin program. We then use these multipliers to improve arbitrary eligible contracts into equity bonus contracts.

Now, the structure of (3) implies that the principal may restrict attention to a certain class of contracts when optimizing her guarantee. The form taken by these contracts is revealed in the next lemma.

Lemma 1. *If the contract w is eligible, then there exists a constant κ and a vector of non-negative constants $\lambda = (\lambda_1, \ldots, \lambda_n)$ with $\lambda_i > 0$ for at least one i such that*

$$y - w(y) \geq \sum_i \lambda_i (w(y) - w_i(y)) + \kappa \quad \forall y \in \mathcal{Y} \tag{4}$$

$$\Pi(w) = \sum_i \lambda_i V_i(w) + \kappa. \tag{5}$$

The proof of Lemma 1 is based around an application of the separating hyperplane theorem and is essentially a multi-dimensional adaptation of a similar argument that appears in Carroll [2015], applied to transformed payoff vectors of the form $(w(y) - w_1(y), \ldots, w(y) - w_n(y); y - w(y))$ in our analysis rather than $(w(y); y - w(y))$ as in Carroll's analysis.

The constants λ given in the statement of Lemma 1 can be interpreted as the multipliers on the constraints in (3). Rearranging (4) yields the pointwise inequality

$$w(y) \leq \frac{1}{1+\sum_i \lambda_i}\left(y + \sum_i \lambda_i w_i(y) - \kappa\right), \tag{6}$$

the right hand side of which delineates a new, putatively improving contract

$$w^{\lambda,\kappa}(y) \equiv \frac{1}{1+\sum_i \lambda_i}\left(y + \sum_i \lambda_i w_i(y) - \kappa\right) \tag{7}$$

that pays the agent at least as much for each level of output as the original contract w.

We make two comments about this contract here. First, $w^{\lambda,\kappa}$ inherits limited liability from w (for a suitable κ) and continuity from the empirical contracts w_i. Thus, it is a well-defined contract. Second, because $w^{\lambda,\kappa}(y) \geq w(y)$ for all outputs y, we also have $V_i(w^{\lambda,\kappa}) \geq V_i(w)$ for all i. Therefore, for *any* F satisfying the cycle constraints in (2) we have

$$\langle F, y - w^{\lambda,\kappa}\rangle = \sum_i \lambda_i \langle F, w^{\lambda,\kappa} - w_i\rangle + \kappa \geq$$
$$\sum_i \lambda_i V_i(w^{\lambda,\kappa}) + \kappa \geq \sum_i \lambda_i V_i(w) + \kappa = \Pi(w), \tag{8}$$

where the first equality follows from the definition of $w^{\lambda,\kappa}$, the second inequality from (2), and the last equality from Lemma 1. Altogether, we conclude $\Pi(w^{\lambda,\kappa}) \geq \Pi(w)$. Regardless of the form of the original contract w, the improved contract $w^{\lambda,\kappa}$ belongs to the following class.

Definition 3. Contract w is an *equity bonus contract* if there exists a convex combination of empirical contracts $w_0 = \sum_i \beta_i w_i$, a share parameter $\alpha \in (0, 1]$, and a constant $\gamma \geq 0$ such that

$$w(y) = w_0(y) + \alpha(y - w_0(y)) + \gamma.$$

Equity bonus contracts pay the agent a convex combination of empirical contracts $w_0 = \sum \beta_i w_i$ plus a share α of the principal's profit $y - w_0(y)$ and a constant wage γ. When w_0 is itself an empirical contract, these contracts can be interpreted simply as a payment structure in which the principal (i) pays the agent with contract w_i, (ii) appends to that contract a share α of the surplus he generates $(y - w_i(y))$, and (iii) adds a bonus payment (or penalty) γ to the agent's wage. More generally, if w is a mixture of k contracts, then we call w a *k-mixture equity bonus contract.* We delay our discussion of the optimal k until Section 6.3, and for the time being embrace the possibility that optimal equity bonus contracts might include multiple contracts from the data.

In order to state our main result, we introduce the program

$$\max_{\lambda \geq 0, \kappa} \sum_i \lambda_i V_i(w^{\lambda,\kappa}) + \kappa \quad \text{s.t. } w^{\lambda,\kappa}(y) \geq 0 \; \forall y \in \mathcal{Y}. \tag{9}$$

It will be helpful to keep in mind that the quantities $V_i(w)$ have a simple, closed form representation when w is an equity bonus contract $w^{\lambda,\kappa}$. In particular,

$$V_i(w^{\lambda,\kappa}) = \max_j \left(\frac{\langle F_j, y - w_i \rangle + \sum_k \lambda^k \langle F_j, w_k - w_i \rangle}{1 + \sum_k \lambda_k} - \Delta(i,j) \right) + \kappa.$$

Theorem 2. *If* (9) *has a solution* (λ^*, κ^*) *with value* Π^**, then equity bonus contract* w^{λ^*,κ^*} *is optimal and* $\Pi(w^{\lambda^*,\kappa^*}) = \Pi^*$*. If* (9) *does not have a solution and* i *solves* $\max_i \langle F_i, y - w_i \rangle$*, then empirical contract* w_i *is optimal and* $\Pi(w_i) = \langle F_i, y - w_i \rangle$*.*

Theorem 2 formalizes our discussion of the optimality of equity bonus contracts and provides a characterization of optimality within that class of contracts. Mathematically, these contracts are optimal because they impose a tight relationship on the principal's objective $y - w$ and the supplemental payments $w - w_i$ that characterize incentive compatibility. Economically, they are optimal because (i) their empirical component w_i makes direct use of the revealed preference information contained in the principal's data and (ii) their novel component $y - w_i$ provides the agent with incentives to maximize the principal's profit, as linear contracts do in Carroll [2015].

The proof of the theorem verifies that solving (9) is sufficient for solving the principal's problem. The main difficulty in the argument stems from the fact that it need not be true in general that the "expected" combination of incentive compatibility constraints binds for an arbitrary equity bonus contract w, meaning it may not necessarily be the case that the $V_i(w)$ constraint binds if and only if contract w_i is included in w. Consequently, the payoff guarantee given by the equity bonus contract $w^{\lambda,\kappa}$ might in principle exceed (rather than meet) its lower payoff bound $\kappa + \sum \lambda_i V_i(w^{\lambda,\kappa})$. However, Lemma 1 is strong enough to imply not only that the bound is tight at the solution, but also that maximizing the lower bound stated in the objective of (9) is equivalent to maximizing the principal's payoff guarantee, as desired. Details are in the appendix.

Now, Theorem 2 falls short of providing a complete characterization of optimality in part because it does not tell us whether or not there are contracts other than equity bonus contracts (or empirical contracts, as appropriate) that might be optimal. When output distributions have full support, no other optimal contracts exist.

Proposition 2. *Suppose output distribution F_i has full support for every i and let w^* be an optimal contract. If (9) has a solution (λ^*, κ^*) with $\Pi^* > \max_i \langle F_i, y - w_i \rangle$, then w^* is an equity bonus contract. If (9) has a solution (λ^*, κ^*) with $\Pi^* = \max_i \langle F_i, y - w_i \rangle$, then w^* is either an equity bonus contract or an empirical contract. Finally, if (9) does not have a solution, then w^* is an empirical contract.*

By way of interpretation, when the known output distributions F_i have full support, the principal can always improve the agent's payoff for choosing those actions by increasing his wage for any output that he produces. At the same time, the unknown actions F that drive the principal's payoff guarantee assign positive probability only to output states where the principal's payoff $y - w$ is small relative to the supplemental payments $w - w_i$ that incentivize the agent. Thus, by increasing wages outside of the support of those "bad" actions F, the principal further incentivizes the productive actions F_i without incentivizing undesirable actions F. Because equity bonus contracts align the principal's payoff with the agent's incentives, this is strictly desirable from her perspective. Thus, equity bonus contracts are strictly better than their counterparts. Mathematically, Proposition 2 can quickly be understood via (8). There, when empirical actions have full support, the final weak inequality $\geq$ is instead strict.

Finally, as a last result for this section, our uniqueness result immediately yields a result on weak dominance (mostly) as a corollary.

Corollary 3. *Suppose output distribution F_i has full support for every i. If (9) has a solution (λ^*, κ^*) with $\Pi^* > \max_i \langle F_i, y - w_i \rangle$, then there exists an optimal equity bonus contract that is not weakly dominated by any other contract. If (9) has a solution (λ^*, κ^*) with $\Pi^* = \max_i \langle F_i, y - w_i \rangle$, then there exists either an optimal equity bonus contract or an optimal empirical contract that is not weakly dominated by any other contract. Finally, if (9) does not have a solution, then there exists an empirical contract that is not weakly dominated by any other contract.*

The substance of our results on undominated contracts is contained in the uniqueness results of Proposition 2, because robustly optimal contracts are not weakly dominated by contracts that are not themselves maxmin optimal; we address some remaining technicalities around the existence of an undominated contract in the appendix. To the extent that even robustly optimal contracts might be undesirable if they are weakly dominated, the corollary provides further justification for the use of equity bonus contracts and empirical contracts. On the other hand, our result comes with its own caveats. In particular, we have not ruled out that our optimal contracts might be weakly dominated by screening menus of contracts, nor that the principal might be able to achieve a better guarantee by randomizing over contracts.[5] We leave the development of these observations to future work.

5 When are novel contracts optimal?

Theorem 2 shows that optimal novel contracts must be equity-bonus contracts. In this section, we study the conditions under which these contracts exist, discuss the intuition behind these conditions, and show through example that optimal novel contracts exist for some data sets.

It is helpful to think of optimal novel contracts as being modified versions of contracts in the data, where the modifications are designed to robustly induce the agent to select an alternative, more profitable action from the data. To emphasize intuition, in this section we consider the situation where there are two data points, $X = \{(F_1, w_1), (F_2, w_2)\}$ and the principal would earn the most profit from contract w_2 and action F_1, were she able to select any combination of actions and contracts from the data.

Before proceeding, note that rationalizability with two data points implies that there exist e_1 and e_2 satisfying

$$s_{21} = \langle F_1 - F_2, w_1 \rangle \geq e_1 - e_2 \geq -\langle F_2 - F_1, w_2 \rangle = -s_{12}.$$

This condition implies a bound on the agent's loss in surplus from choosing F_1 at w_2,

$$\langle F_2, w_2 \rangle - e_2 - \langle F_1, w_2 \rangle + e_1 \leq \langle F_2 - F_1, w_2 \rangle + s_{21} = \Delta.$$

The Δ term therefore measures the largest rationalizable loss the agent could experience

[5] Screening does not improve the principal's guarantee in the standard robust contracts problem, as is demonstrated in Carroll [2015]; it seems an open question as to whether or not screening might weakly dominate paying the agent with a single contract. Elsewhere, Kambhampati [2023] confirms that there are indeed randomized contracts that guarantee the principal a higher payoff than the best deterministic contract.

from choosing F_1 at w_2 (note that $\Delta = \Delta(2, 1) = \Delta(1, 2)$). Intuitively, the larger is Δ, the more the principal would have to compensate the agent to choose F_1 over F_2 if these were the only available actions.

Theorem 3 implies that with two data points the optimal equity bonus contract, when it exists, is a 1-mixture contract (see Section 6.3). In this case, the relevant 1-mixture contract is w_2, $w = w_2 + \alpha(y - w_2)$. From Theorem 2, the optimal contract maximizes the quantity

$$(1-\alpha)\langle F_1, y - w_2\rangle - \frac{1-\alpha}{\alpha}\Delta$$

as a function of α and, as we show in the proof, the maximized quantity is its guarantee. Maximization yields

$$\alpha^* = \sqrt{\frac{\Delta}{\langle F_1, y - w_2\rangle}}$$

by the same calculation as appears in Chassang [2013], Carroll [2015]. Notice that the size of the optimal equity share α is positively related to Δ, so that larger rationalizable losses Δ require more equity compensation. The profit guarantee from this contract is

$$\Pi(w) = \left(\sqrt{\langle F_1, y - w_2\rangle} - \sqrt{\Delta}\right)^2.$$

Therefore, a sufficient condition for there to exist an optimal novel contract in this setting is

$$\left(\sqrt{\langle F_1, y - w_2\rangle} - \sqrt{\Delta}\right)^2 \geq \max\{\langle F_1, y - w_1\rangle, \langle F_2, y - w_2\rangle\}, \tag{10}$$

which shows that optimal novel contracts exist when Δ is sufficiently small or when the agent is sufficiently close to being indifferent between the two observed actions at w_2. Relatedly, the condition suggests changes to the data that would be particularly valuable to the principal. For example, consider modifying w_1 according to $w_1' = w_1 + \lambda z$, with $\lambda > 0$, $\langle F_1, z\rangle = 0$ and $\langle F_2, z\rangle > 0$.[6] An increase in λ causes Δ to fall (without affecting other terms in (10)), because this change gives the agent more utility from the action F_2 at w_1 without affecting their payoff from F_1. For large enough λ, Δ approaches zero, at which point the agent is provably indifferent between the two observed actions at both observed contracts.

To explore additional implications of (10), the following proposition summarizes a set

[6] Such a z exists if $\Delta > 0$.

of necessary conditions for the optimality of w.

Proposition 3. *Suppose there are two data points and $w = w_2 + \alpha(y - w_2)$ is an optimal contract. Then*

$$\langle F_1, y - w_2 \rangle \geq \Delta \tag{11}$$

$$\langle F_1, w_1 - w_2 \rangle > \langle F_2, w_1 - w_2 \rangle \geq 0 \tag{12}$$

$$\langle F_1 - F_2, y \rangle \geq \langle F_1 - F_2, w_1 \rangle > \langle F_1 - F_2, w_2 \rangle. \tag{13}$$

The conditions in (12) imply that w_1 has a higher expected payment across both actions in the data. As a consequence $e_1 \geq e_2$ for all rationalizing effort costs and we can label F_1 the "higher" action. The conditions in (13) imply that

$$\langle F_1, y - w_2 \rangle - \langle F_2, y - w_2 \rangle \geq \langle F_1 - F_2, w_1 \rangle - \langle F_1 - F_2, w_2 \rangle = \Delta,$$

so the principal gains more from F_1 at w_2 than the agent stands to lose from selecting F_1 at w_2. In other words, the data reveal that there is an increase in surplus if F_1 is selected at w_2 across all rationalizing technologies.

We illustrate the key ideas above through an example.

Example 2. Let $\mathcal{Y} = (0, 30, 60)$ and suppose the data X consists of two observations

$$w_1 = (0, 4, 56) \qquad F_1 = (0, 1/2, 1/2)$$

$$w_2 = (0, 20, 20) \qquad F_2 = (3/4, 0, 1/4).$$

from which we calculate

$$\langle F_1, y - w_1 \rangle = 15 \qquad \langle F_2, y - w_2 \rangle = 10$$

$$\langle F_1, y - w_2 \rangle = 25 \qquad \Delta = 1.$$

Letting $w = w_2 + \alpha^*(y - w_2)$ with $\alpha^* = 1/5$, we get a guaranteed profit of $\Pi(w) = 16$.

To modify w_1 to improve the profit guarantee as suggested by the paragraph following (10), let $z = (0, -1, 1)$ and $w_1' = w_1 + 4z = (0, 0, 60)$, making $\Delta' = 0$ and the optimal profit guarantee $\langle F_1, y - w_2 \rangle = 25$.

6 When are optimal contracts simple?

The form of any particular equity bonus contract depends on the form of the empirical contracts from which it is constructed. In this section, we consider three special cases in which the data satisfy stricter conditions than rationalizability and analyze what implications those conditions have for the features of the optimal contract. First, we show quickly that if empirical contracts are nondecreasing then so too is the optimal contract. Next, we show that if each empirical contract is linear, then the optimal contract is not only linear but in fact empirical. Finally, and most importantly, we show that whenever the data are supermodular in an appropriate sense, then optimal contracts take the especially simple form of equity bonus contracts constructed from a single empirical contract.

6.1 Monotone data

The potential failure of monotonicity in the classical moral hazard problem is occasionally regarded as a peculiarity (Holmström [2017], Innes [1990]). It follows quickly in our environment that whenever all of the contracts in the principal's data are non-decreasing, so too are our optimal contracts.

Proposition 4. *If all empirical contracts w_i are nondecreasing in output, then all equity bonus contracts are nondecreasing in output.*

6.2 Linear data

Linear contracts are a special case of the monotone contracts considered in the last section, so Proposition 4 applies to them as well. However, while we can find examples such as Example 2 in which the principal can guarantee herself a non-trivial improvement over the profits available from the empirical contracts when her data is monotone, there is no such hope when all contracts in the data are linear.

Proposition 5. *If all empirical contracts take the form $w_i(y) = \alpha_i y$ for $\alpha_i > 0$, the principal's optimal contract is empirical.*

While it is apparent from our definition that equity bonus contracts will be linear when all of the empirical contracts in X are linear, our result is stronger. We show that not only is the optimal contract linear, but it is exactly one of the contracts in the principal's data, meaning it uses the same share parameter. The proof uses the fact that when all of the empirical contracts are linear, the cycle constraints in (2) take a simple form. In particular, the data can be arranged in a way to make α_i and $\langle F_i, y\rangle$ increasing in i, and for a novel contract $w = \alpha_0 y$ one can find an i such that $\alpha_i \leq \alpha_0 \leq \alpha_{i+1}$, which then can be shown to imply that

$$\langle F_i, y\rangle \leq \langle F, y\rangle \leq \langle F_{i+1}, y\rangle,$$

for the unknown distribution F. Finally, we show that $\langle F_i, y\rangle = \langle F, y\rangle$ cannot be ruled out, and from here it follows that w cannot improve on w_i.

6.3 Multi-contract mixtures

Example 2 shows that equity bonus contracts can be quite simple in that they supplement a single empirical contract with an equity payment. Furthermore, the resulting 1-mixture contracts are straightforward to optimize because their payoff guarantee derives from a single binding constraint. In particular, as suggested in Example 2, one can find the best 1-mixture contract by selecting a w_i, forming the equity bonus contract $\alpha(y - w_i) + w_i$ and choosing α to maximize the right-hand side of the i^{th} cycle constraint. On the other hand, identifying the solution to the general form of program (9) might in principle be difficult, both because there are many incentive compatibility constraints that might simultaneously bind and also because the program's objective is sometimes non-convex. The next example shows that such possibilities cannot be ignored, as there exist data sets where 1-mixture contracts are strictly suboptimal.

Example 3. Suppose the set of outputs $\mathcal{Y}$ is the vector $(0, \varepsilon_1, \varepsilon_2, \varepsilon_3, \varepsilon_4, 100)$, where the ε_i are distinct and approximately zero. There are four observations (F_i, w_i), each of which is represented by row i in the corresponding matrices below.

$$F = \begin{pmatrix} 0.65 & 0.05 & 0 & 0 & 0.2 & 0.1 \\ 0 & 0 & 0 & 0.2 & 0.8 & 0 \\ 0.65 & 0 & 0.2 & 0.05 & 0 & 0.1 \\ 0 & 0.2 & 0.8 & 0 & 0 & 0 \end{pmatrix} \qquad w = \begin{pmatrix} 0 & 0 & 0 & 0 & 10 & 60 \\ 0 & 0 & 0 & 10 & 0 & 10 \\ 0 & 0 & 10 & 0 & 0 & 60 \\ 0 & 10 & 0 & 0 & 0 & 10 \end{pmatrix}$$

There are two optimal 1-mixture contracts, one constructed from w_2 and the other from w_4. Each yields payoff guarantee 4. However, neither contract is optimal in general. For instance, the 2-mixture contract $w = w_0 + 0.3(y - w_0)$ with $w_0 = (w_2 + w_4)/2$ has a worst-case profit of 4.375. It is straightforward to check that w also improves on the best empirical contract, which yields a profit of 2.

The difficulty inherent in identifying optimal contracts for arbitrary rationalizable data sets stems from the potentially ill-behaved nature of the principal's problem. While Theorem 2 suggests a computational method for finding the optimal equity bonus contract, it does not yield much in terms of analytical insight, at least in part because the objective will typically neither be concave nor convex. On the other hand, if the data satisfy more stringent conditions than rationalizability, then the problem of identifying an optimal contract is much simpler.

To elaborate, rationalizability itself puts only weak requirements on the structure of the agent's technology. In particular, there are rationalizable data sets X under which there is no sense in which the agent's actions can be ordered to make "local" incentive constraints imply the "global" ones, in line with the often invoked first-order approach to the standard model. Example 3 shows that there exist data sets for which the optimal equity bonus contract involves convex combinations of multiple observations from the data. We show next that these two facts are closely related.

First, we address the question of when we can assign an order to the actions in the data. We say that the data are ordered if there is an index on the data such that for each contract w_i actions that are closer to F_i with respect to the indexing scheme are preferred to those further away for *all* effort costs that rationalize the data. Formally,

Definition 4. The data can be *ordered* if there is an index on the data such that for all observations i and all rationalizing choices of effort costs, $\{e_1, \ldots, e_n\}$

$$\langle F_j - F_k, w_i \rangle \geq e_j - e_k,$$

whenever $i < j < k$ or $i > j > k$.

This is analogous to being able to prove from the data that the agent has quasiconcave preferences over the actions in the data. Importantly, note that the ordering of the actions is not independent of the observed contracts in the definition. To conclude that the data are ordered, it is sufficient for the data to satisfy a supermodularity condition.

Definition 5. The data are *supermodular* if[7]

$$\langle F_j - F_{j'}, w_i - w_{i'} \rangle \geq 0 \quad \forall j < j', i < i'. \tag{14}$$

[7]This condition is equivalent to the superficially weaker alternative with restrictions $j' = j + 1$ and $i' = i + 1$.

We emphasize that data sets X that contain only two or fewer observations automatically satisfy our supermodularity condition. Proceeding,

Proposition 6. *If the data are supermodular, then the data can be ordered.*

Proof. Fix an index on the data. If $i < j < k$, we have $\langle F_j - F_k, w_i \rangle \geq \langle F_j - F_k, w_j \rangle \geq e_j - e_k$, where the first inequality is due to supermodularity and the second is due to rationalizability. The $i > j > k$ case is analogous. □

Next, we show that under the supermodularity condition optimal equity bonus contracts cannot incorporate more than one contract from the data. To prove this result, we exploit the additional structure on the cycle constraints imposed by the supermodularity condition. For the remainder of this section, we assume that the supermodularity condition holds and that the index on the data is consistent with this condition holding. In the next lemma, we show that under the supermodularity condition, shortest paths through the data must traverse adjacent nodes.

Lemma 2. *Between data points i and j, there is a shortest path that traverses the sequence of adjacent nodes indexed by $\{i, i+\iota, i+2\iota, \ldots, j-\iota, j\}$ where $\iota = \operatorname{sgn}(j-i)$. Consequently, for any $i < j < k$, $s_{ik} = s_{ij} + s_{jk}$ and $s_{ki} = s_{kj} + s_{ji}$.*

Intuitively, the shortest paths in the data pick up on the agent's most preferred alternative action choices, so shortest paths existing between adjacent nodes coincides with the quasiconcavity of the agent's preferences over actions in the data.

Towards the purpose of proving that the optimal contracts for the principal are 1-mixture contracts, the importance of the prior lemma is that it can be used to show that the problem in Theorem 2 is convex in directions corresponding to the removal of constituent empirical contracts from its mixture. We show this formally in the Appendix (Lemma 9).

In Theorem 3, we also require that the contracts in the data never pay more than output, a condition referred to as two-sided limited liability.

Definition 6. The empirical contracts satisfy *two-sided limited liability* if $w_i(y) \geq 0$ and $y \geq w_i(y)\ \forall i, \forall y \in \mathcal{Y}$.

As a consequence of two-sided limited liability, there is no opportunity for the principal to gain from subtracting a fixed payment from the contracts in the data.[8]

[8] The contracts in the data were already assumed normalized in the sense that each contract assigns a payment of 0 to *some* state. The possibility that we rule out for Theorem 3 is that a convex combination of these contracts admits a contract with a minimum payment strictly greater than 0.

Lemma 3. *Under two-sided limited liability, $\kappa = 0$ in any solution to* (9).

Proof. For any λ, $w^{\lambda,\kappa}(0) = -\kappa/(1 + \sum_i \lambda_i)$, and hence $\kappa \leq 0$ with optimal value $\kappa = 0$. □

With these intermediate results in place, our main result for this subsection is that the supermodularity condition implies the optimality of 1-mixture contracts.

Theorem 3. *If the data are supermodular and satisfy two-sided limited liability, then 1-mixture contracts are optimal.*

The situation in which the data can be ordered via the supermodularity condition is analogous to the sufficiency of local incentive compatibility constraints in the standard moral hazard problem (c.f., Grossman and Hart [1983]), in the sense that in both cases it is impossible for multiple (non-local) incentive constraints to simultaneously bind at an optimum. In turn, because the set of potentially binding constraints is much smaller, the potentially complex task of identifying an optimal contract is much easier.

7 Heterogeneous agents

Thus far, we have assumed that the revealed preference data are generated by a continuum of identical agents equipped with the same production technology. In this section, we show how all of our results generalize without complication to a general setting that allows for broad heterogeneity within the population of agents. In doing so, we make no structural assumptions about the nature of that heterogeneity and we introduce no additional assumptions about the agents or the revealed preference data. Our principal views both homogeneity and all forms of heterogeneity as plausible.

Formally, we model agent heterogeneity as stochasticity in the agent's production technology. For the purposes of illustration, suppose as we have done throughout the paper that the revealed preference data X satisfy the no-negative-cycles criterion. While X might have been generated by a continuum of agents each equipped with the same technology $\mathcal{A}$, it might also have been generated by a continuum of agents equipped with either the technology $\mathcal{B}$ with probability p or the technology $\mathcal{C}$ with probability $1-p$. In the former case, the fixed technology rationalizes the data if the best response $(G, e) = (A_i, a_i)$ to contract w_i given technology $\mathcal{A}$ satisfies $A_i = F_i$ for each i; in the latter case, the heterogeneous technologies rationalize the data if the respective best responses $(B_i, b_i), (C_i, c_i)$ to w_i given technologies $\mathcal{B}, \mathcal{C}$ satisfy $pB_i + (1-p)C_i = F_i$ for each i. Thus, F_i is the "average" best response to w_i across the population of agents. As we assumed earlier in the paper that $\mathcal{A}$ is fixed across observations, we assume here that the technologies $\mathcal{B}, \mathcal{C}$ and the population weight p are similarly constant.

TODO: DOUBLE CHECK INDEXING THROUGHOUT PAPER ON BOTH THE SET X (are we using K or N or n or none of the above? We should do this consistently and make sure there's nothing in this section that clashes with it

Definition 7. A finite family of technologies $\mathcal{A}_1, \ldots, \mathcal{A}_J$ and a probability $P \in \Delta^J$ *rationalize* the revealed preference data X if for each i, j there exists an action $(G_{ij}, e_{ij}) \in c(\mathcal{A}_j|w_i)$ such that

$$\sum_{j=1}^{J} G_{ij} P_j = F_i.$$

We write $\mathcal{P}$ for the set of all such rationalizing stochastic technologies and

$$\Pi(P|w) = \sum_{j=1}^{J} \Pi(A_j|w) P_j \qquad\qquad \Pi^{\mathcal{P}}(w) = \inf_{P \in \mathcal{P}} \Pi(P|w)$$

for the principal's payoff given distribution P and her payoff guarantee against the set of all such distributions $\mathcal{P}$, respectively, with dependency on the identities of $\mathcal{A}_1, ..., \mathcal{A}_J$ suppressed in our notation for the purposes of readability.

As above, a finitely supported distribution of technologies P rationalizes the data if the average best response to contract w_i under distribution P is F_i. We emphasize for clarity's sake that the both the support $\mathcal{A}_1, \ldots, \mathcal{A}_J$ of P and the size of that support J are variables and unknown to the principal. Thus, the rationalizing set $\mathcal{P}$ contains both degenerate technology distributions (in the homogeneous agent case) and non-degenerate distributions with support of arbitrary (finite) size. We require finite support both to avoid unnecessary technicalities and, more importantly, to facilitate an interpretation of our heterogeneous revealed preference exercise as one in which there are a large number of agents within each "type".

Remark 1. The set of rationalizing technologies $\mathscr{A}$ is embedded in the set of rationalizing technology distributions $\mathcal{P}$ by the map $P = \delta(\mathcal{A})$.

Theorem 4. *For all contracts w, $\Pi^{\mathcal{P}}(w) = \Pi(w)$.*

Theorem 4 shows that our identical-agents assumption is without loss of generality from the perspective of worst-case payoff guarantees. The proof is straightforward and based around the observation that if action (F_j, e_j) is a best response to contract w under technology $\mathcal{A}_j$ for each j, then the average action $(\sum F_j P_j, \sum e_j P_j)$ is a best response to w under the average technology $\sum \mathcal{A}_j P_j$ for any distribution P. To conclude this section, we clarify the role of the no-negative-cycles assumption in characterizing the rationalizability of the revealed preference data by distributions of technologies.

Remark 2. There exists a distribution P of technologies that rationalize the revealed preference data X if and only if X has no cycles of negative length.

In the proof of Theorem 4 we show that the average of any distribution of rationalizing technologies is itself a rationalizing technology. Thus, there are no data sets that can be rationalized by heterogeneous agent technologies that can not also be rationalized by a single fixed technology. In turn, we obtain Remark 2.

8 Conclusion

We study a principal-agent problem in which the principal's information about the agent's technology is characterized by revealed preference data. We show that robustly optimal contracts append equity payments to the empirical contracts in the data. Both of these features contrast sharply with the classic moral hazard model, which assumes the principal is endowed with complete information about the agent's capabilities and yields optimal contracts that are both complex and may have features that are not typically seen in practice.

Throughout our analysis, we take the data X as being exogenously generated. We imagine the data as being generated by a population of agents that are each paid with a standard contract that periodically adjusts, perhaps in response to external wage pressures. Alternatively, we imagine that different populations of agents within the same organization might be employed in the same capacity but be subject to different incentive schemes, either because of heterogeneous hiring circumstances or because they belong to different units within the same institution. When interpreting our results along these lines, it will be important to keep in mind that our generalization to heterogeneous agents requires only that the *population* that generates the data is the same as the population from which the agent is drawn, just as researchers in the social and natural sciences use data to make out-of-sample predictions.

Setting aside the stationarity of the contracting environment, the most important remaining avenue for future work seems to be to relax our assumption that the principal observes the true distribution of output produced in best response to each of the contracts in her data, rather than a finite sample. We leave this problem, which is of significant practical importance, to interested researchers.

Altogether, we view our analysis here as only one perspective on a much broader empirical contract design exercise in which one might consider different types of data sets and contracting environments.

References

S.N. Afriat. The construction of a utility function from expenditure data. *International Economic Review*, 8:67–77, 1967. 3

Nemanja Antic and George Georgiadis. Robust contracts: A revealed preference approach. Working Paper, 2022. 1

Justin Burkett and Maxwell Rosenthal. Statistical uncertainty and coarse contracts. Working Paper, 2022. 1

Gabriel Carroll. Robustness and linear contracts. *American Economic Review*, 105(2): 536–563, 2015. 1, 1, 4, 4, 4, 5, 5

Gabriel Carroll. Robustness and separation in multidimensional screening. *Econometrica*, 85(2):453–488, 2017. 1

Sylvain Chassang. Calibrated incentive contracts. *Econometrica*, 81(5):1935–1971, 2013. 1, 5

Tianjiao Dai and Juuso Toikka. Robust Incentives for Teams. *Econometrica*, Forthcoming, 2022. 1

George Georgiadis and Michael Powell. A/b contracts. *American Economic Review*, 112 (1):267–303, 2022. 1

Sandford J. Grossman and Oliver D. Hart. An analysis of the principal-agent problem. *Econometrica*, 51(1):7–45, 1983. 6.3

Bengt Holmström. Moral hazard and observability. *The Bell Journal of Economics*, 10(1): 74–91, 1979. 1

Bengt Holmström. Pay for performance and beyond. *American Economic Review*, 107(7): 1753–1777, 2017. 1, 6.1

Robert D. Innes. Limited liability and incentive contracting with ex-ante action choices. *Journal of Economic Theory*, 52(1):45–67, 1990. 6.1

Ashwin Kambhampati. Randomization is optimal in the robust principal-agent problem. *Journal of Economic Theory*, 207, 2023. 1, 5

Jean-Charles Rochet. A necesary and sufficient condition for rationalizability in a quasi-linear context. *Journal of Mathematical Economics*, 16:191–200, 1987. 3

Maxwell Rosenthal. A revealed preference approach to multidimensional screening. Working Paper, 2019. 1

Rakesh V. Vohra. *Mechanism Design: A Linear Programming Approach*. Cambridge University Press, New York, NY, 2011. 3, A.1

A Appendix

Throughout the appendix, we make use of the notation

$$C_i(w) \equiv \arg\max_j \left(\langle F_j, w - w_i \rangle - \Delta(i,j)\right)$$

for the set of empirical actions that achieve the maximum $V_i(w)$.

A.1 Theorem 1

Proof of Theorem 1. The equivalence of the last two statements is implied by Corollary 3.4.2 of Vohra [2011] (page 34). The first is equivalent to the third because, as discussed in that text, the system in (1) is isomorphic to the dual of a shortest path problem in the data. See also Section 7.1 of that text for a direct connection to the problem of rationalizing a set of choice data with a quasilinear utility function. □

A.2 Supporting results

Lemma 4. *Suppose that X is a set of observations with the no-negative-cycles property and let $(F, w) \in X$. There exists a rationalization $e : X \to \mathbb{R}$ of X in which F is strictly incentive compatible at w if and only if every non-degenerate cycle through (F, w) in X has strictly positive length.*

Proof. For each $(G, v) \in X$, write $P(G, v)$ for the shortest path from (F, w) to (G, v) in X and $C(G, v)$ for the cycle $P(G, v) \to (F, w)$. We prove the if statement first and then the only if statement.

$\Longleftarrow$ Define effort map $e : X \to \mathbb{R}$ by $e(F, w) = 0$ and $e(G, z) = \ell(P(G, z))$. First, we claim e yields the desired strict preference for action F at contract w. Accordingly, choose $(G, v) \in X$ distinct from (F, w). Because $C(G, v)$ is a cycle through (F, w), we have $\ell(C(G, v)) > 0$ by hypothesis. In turn, we obtain

$$\langle F - G, w\rangle + \ell(P(G, v)) > 0.$$

It follows immediately that $\langle F - G, w\rangle > e(F, w) - e(G, v)$, as claimed. Next, we claim e rationalizes X. Choose $(H_1, z_1), (H_2, z_2) \in X$ arbitrarily. Because $P(H_1, z_1) \to (H_2, z_2)$ is a path from (F, w) to (H_2, z_2) but $P(H_2, z_2)$ is the shortest such path, we have

$$\ell(P(H_1, z_1)) + \langle H_2 - H_1, z_2\rangle \geq \ell(P(H_2, z_2)).$$

In turn, we obtain $\langle H_2 - H_1, z_2\rangle \geq e(H_2, z_2) - e(H_1, z_1)$ from the definition of e. Thus, e rationalizes X.

$\Longrightarrow$ Per Theorem 1, X is weakly rationalizable if and only if X has the no-negative-cycles property. Accordingly, if X includes any cycles through (F, w) with negative length, then there is no rationalization of X under which F is even weakly incentive compatible at w. Suppose alternatively that X has the no-negative-cycles property but C has length zero for some non-degenerate cycle through (F, w). Let $e : X \to \mathbb{R}$ be any rationalization of X. Enumerate $C(g, v)$ by $(H_1, z_1), \ldots, (H_n, z_n)$ and write $e_i \equiv e(H_n, z_n)$. First, because

$(H_1, z_1) = (H_n, z_n)$ by definition, we have

$$\sum_i (e_i - e_{i-1}) = 0.$$

In turn, because C has length zero, we have

$$\sum_i \langle H_i - H_{i-1}, z_i \rangle = \sum_i (e_i - e_{i-1}). \tag{15}$$

Finally, because e rationalizes X, we have for all i

$$\langle H_i - H_{i-1}, z_i \rangle \geq e_i - e_{i-1}. \tag{16}$$

It follows jointly from (15) and (16) that $\langle H_i - H_{i-1}, z_i \rangle = e_i - e_{i-1}$ for all i. Because $(H_i, z_i) = (F, w)$ for some i and C is non-degenerate by hypothesis, $\langle F - G, w \rangle = e(F, w) - e(G, v)$ for some $(G, v) \in X$.

□

Lemma 5. *If $j \in C_i(w)$ then $\langle F_j, w - w_k \rangle \geq V_k(w)$ for all k.*

Proof. Let $j \in C_i(w)$ and $l \in C_k(w)$. First, by definition of C_i we have

$$\langle F_j, w - w_i \rangle - \Delta(i, j) \geq \langle F_l, w - w_i \rangle - \Delta(i, l). \tag{17}$$

Substituting in the definition of $\Delta(i, j), \Delta(i, l)$ into (17) yields

$$\langle F_j - F_l, w \rangle \geq s_{ij} - s_{il}. \tag{18}$$

Next, because s_{jk} is the length of the shortest path from j to k and s_{il} is the shortest path from i to l, we have

$$\langle F_k - F_j, w_k \rangle \geq s_{jk} \tag{19}$$

$$s_{ij} + s_{jk} + s_{kl} \geq s_{il}. \tag{20}$$

Together, (18), (19), (20) imply

$$\langle F_j - F_l, w \rangle + \langle F_k - F_j, w_k \rangle + s_{kl} \geq 0. \tag{21}$$

Finally, substituting the definition of $\Delta(k, l)$ into (21) yields $\langle F_j, w - w_k \rangle \geq \langle F_l, w - w_k \rangle - \Delta(k, l) = V_k(w)$, as claimed. □

Corollary 4. *For each observation $(F_i, w_i) \in X$ the principal's payoff guarantee is $\Pi(w_i) = \langle F_i, y - w_i \rangle$.*

Proof. Choose $(F_i, w_i) \in X$. First, *for every* rationalization $\mathcal{A}$ of X, output distribution F_i is weakly incentive compatible given contract w_i and technology $\mathcal{A}$. Second, from Lemma 4, *there exists* a rationalization $\mathcal{B}$ of X under which output distribution F_i is strictly incentive compatible given contract w_i and technology $\mathcal{B}$. Because ties are broken in favor of the principal, these two facts imply $\Pi(w_i) = \langle F_i, y - w_i \rangle$, as claimed. □

Lemma 6. *If w is eligible, then there exists an output distribution F such that $\langle F, w - w_i \rangle > V_i(w)$ for all i.*

Proof. Consider any contract w. Write $S \subset \mathbb{R}^n$ for the free disposal convex hull of the set of supplemental payment vectors $(w(y) - w_1(y), \ldots, w(y) - w_n(y))$ across the set of outputs $\mathcal{Y}$ and write $\nu \equiv (V_1(w), \ldots, V_n(w))$. We argue by contraposition. In particular, suppose that there does not exist an element $s \in S$ with $s_i > \nu_i$ for all i, and let I be the set of boundary indices i with the property

$$s \in S, s \geq \nu \implies s_i = \nu_i. \tag{22}$$

Part 1 First, we claim $C_i(w) = \{i\}$ for each $i \in I$. To see why, choose $i \in I$ and suppose to the contrary

$$\langle F_j, w - w_i \rangle - \Delta(i, j) = V_i(w) \tag{23}$$

for some j distinct from i. Lemma 5 implies $\forall k \ \langle F_j, w - w_k \rangle \geq V_k(w)$. In turn, because $i \in I$ by hypothesis, (22) implies

$$\langle F_j, w - w_i \rangle = V_i(w). \tag{24}$$

From (23) and (24) we obtain $\Delta(i, j) = 0$. This contradicts Assumption 2, and hence we conclude $C_i(w) = \{i\}$ for each $i \in I$.

Part 2 Second, we claim $I = \{i\}$ for some $i \in I$. To see why, suppose to the contrary that I contains distinct indices i, j. From the first step, we have

$$\langle F_i, w - w_i \rangle = V_i(w) \tag{25}$$
$$\langle F_j, w - w_j \rangle = V_j(w). \tag{26}$$

From Lemma 5 and from (22), we also have

$$\langle F_j, w - w_i \rangle = V_i(w) \tag{27}$$
$$\langle F_i, w - w_j \rangle = V_j(w). \tag{28}$$

Altogether, (25)–(28) imply $\langle F_i - F_j, w_i - w_j \rangle = 0$. This again contradicts Assumption 2 and hence we conclude $I = \{i\}$ for some i. We write i^* for the lone element of the set I.

Part 3 Third, note that S is convex and that v belongs to the boundary of S. Accordingly, the supporting hyperplane theorem provides a normal vector $\eta \in \mathbb{R}^n \setminus \{0\}$ such that $\sum \eta_i v_i \geq \sum \eta_i s_i$ for all $s \in S$. We make two claims about η. First, because S contains elements s with s_i arbitrarily negative for all i, it must be that $\eta_i \geq 0$ for all i. Second, by definition of the set I, for each $i \neq i^*$ there exists a supplemental payoff vector $\sigma \in S$ with $\sigma \geq v$ and $\sigma_i > v_i$. In turn, because s is convex, there exists a $\sigma^* \in S$ with $\sigma^* \geq v$ and $\sigma_i^* > v_i$ for all $i \neq i^*$. Consequently, $\eta_i = 0$ for all $i \neq i^*$ and $\eta_{i^*} > 0$. Altogether, we have

$$\forall F \in \Delta(\mathcal{Y}) \ \langle F_{i^*}, w - w_{i^*} \rangle \geq \langle F, w - w_i^* \rangle. \tag{29}$$

Thus, $\operatorname{supp}(F_{i^*})$ is contained in the set of maximizers for $w(y) - w_{i^*}(y)$ and hence there is a constant $\gamma \in \mathbb{R}$ such that $w = w_{i^*} + \gamma$ for all outputs y in $\operatorname{supp}(F_{i^*})$. We claim $\gamma \geq 0$. For the purposes of contradiction, suppose $\gamma < 0$. There are two subcases to consider. First, if $U \equiv \arg\max_y \ w(y) - w_{i^*}(y)$ intersects $V \equiv \arg\min_y w_{i^*}(y)$, then $w(z) = w_{i^*}(z) + \gamma < w_{i^*}(z) < 0$ for some $z \in U \cap V$. Thus, w does not satisfy limited liability and therefore is not a contract. Alternatively, if U does not intersect V, then $w(z) < w_{i^*}(z) + \gamma$ for some $z \in V$. Because $w_{i^*}(z) = 0$ and $\gamma < 0$, we have $w(z) < 0$. We again conclude that w is not a contract. In either case, we conclude $\gamma \geq 0$, as claimed.

Part 4 Fourth, and finally, Lemma 4 implies that there exists a rationalization $\mathcal{A}$ of X with the property that output distribution F_{i^*} is strictly incentive compatible given technology $\mathcal{A}$ and contract w_{i^*}. From (29), it follows that F_{i^*} is also strictly incentive compatible given technology $\mathcal{A}$ and contract w. Altogether, we have $\Pi(w) \leq \langle F_{i^*}, y - w_{i^*}\rangle - \gamma \leq \langle F_{i^*}, y - w_{i^*}\rangle$. Corollary 4 implies $\langle F_{i^*}, y - w_{i^*}\rangle = \Pi(w_{i^*})$, and thus we have $\Pi(w) \leq \Pi(w_{i^*})$. This contradicts our hypothesis that w is eligible and therefore completes the proof.

□

A.3 Proposition 1

Proof of Proposition 1. Consider contract w and suppose w is eligible. There are three parts to the argument. □

Part 1 To see that the claimed lower bound is valid, note that if $\langle F, w - w_i\rangle < V_i(w)$ for any i then there is no technology $\mathcal{A} \in \mathscr{A}$ in which F maximizes the agent's payoff, per Corollary 2. Consequently,

$$\Pi(w) \geq \min_{F\in\Delta(\mathcal{Y})} \langle F, y - w\rangle \text{ such that } \forall i\ \langle F, w - w_i\rangle \geq V_i(w). \tag{30}$$

Part 2 To see that the lower bound in (30) holds with equality, let F achieve the minimum. There are two subcases to consider. First, if $\langle F, w - w_i\rangle > V_i(w)$ for every i, then there exists a technology $\mathcal{A} \in \mathscr{A}$ in which F uniquely maximizes the agent's payoff per Lemma 4. Accordingly, $\langle F, y - w\rangle$ is indeed an upper bound for $\Pi(w)$. Second, if $\langle F, w - w_i\rangle = V_i(w)$ for at least one i, then per our hypothesis that w is eligible Lemma 6 provides for the existence of a distribution G such that $\langle G, w - w_i\rangle > V_i(w)$ for every i. In turn, for every $\lambda < 1$ the distribution

$$F^\lambda \equiv \lambda F + (1 - \lambda)G$$

satisfies $\langle F^\lambda, w - w_i\rangle > V_i(w)$ for every i. Because $\langle F^\lambda, y - w\rangle \to \langle F, y - w\rangle$, the quantity $\langle F, y - w\rangle$ is again an upper bound for the principal's guarantee.

Part 3 Finally, to see that $\langle F, w - w_i\rangle = V_i(w)$ for at least one i, note that if $\langle F, w - w_i\rangle > V_i(w)$ for every i then the perturbation

$$F^\lambda \equiv \lambda F + (1 - \lambda)\delta(0)$$

itself satisfies $\langle F^\lambda, w - w_i\rangle > V_i(w)$ for every i. Moreover, it follows from the eligibility of w that $\langle F^\lambda, y - w\rangle < \langle F, y - w\rangle$. This contradicts our hypothesis that F achieves the minimum in (30) and so we conclude that $\langle F, y - w\rangle = V_i(w)$ for at least one i, as claimed.

A.4 Lemma 1

Proof of Lemma 1. Write $S \subset \mathbb{R}^{n+1}$ for the convex hull of the set of vectors

$$(w(y) - w_1(y), \dots, w(y) - w_n(y); y - w(y))$$

across the set of outputs $\mathcal{Y}$ and $T \subset \mathbb{R}^{n+1}$ for the set of points $(v_1, \dots, v_n; \pi)$ with

$$v_i > V_i(w) \qquad\qquad \pi < \Pi(w).$$

We make two remarks about the sets S, T. First, per Proposition 1, they are disjoint. Second, each is evidently convex. In turn, the separating hyperplane theorem provides for the existence of constants $\lambda_1, \dots, \lambda_n, \mu, \kappa$ such that

$$\lambda_1 v_1 + \cdots + \lambda_n v_n - \mu\pi \leq -\kappa \qquad (v_1, \dots, v_n, \pi) \in S \tag{31}$$

$$\lambda_1 v_1 + \cdots + \lambda_n v_n - \mu\pi \geq -\kappa \qquad (v_1, \dots, v_n, \pi) \in T, \tag{32}$$

with $\lambda_1, \dots, \lambda_n, \mu$ not all zero. We establish four facts about these multipliers.

First, because includes points $(v_1, \dots, v_n, \pi)$ with v_i arbitrarily large and π arbitrarily small, $\lambda_1, \dots, \lambda_n, \mu$ are each non-negative.

Second, we claim $\mu > 0$. To see why, suppose to the contrary that $\mu = 0$. Because $\lambda_1, \dots, \lambda_n, \mu$ are not all 0, we have $\lambda_i > 0$ for at least one index i and $\lambda_1 V_1(w) + \cdots + \lambda_n V_n(w) \geq \lambda_1 v_1 + \cdots + \lambda_n v_n$ for all supplemental payment vectors $(v_1, \dots, v_n)$ in the projection of S onto its first n coordinates. Per Lemma 6, this contradicts eligibility, and so we conclude that $\mu > 0$.

Third, we claim $\lambda_i > 0$ for at least one i. Were this not the case, then $\mu\pi \geq \mu\Pi(w)$ for all values of π in the projection of S onto coordinate $n + 1$. Because $\mu > 0$ and $\min \pi \leq 0 - w(0) \leq 0$, we obtain $\Pi(w) \leq 0$. This again contradicts eligibility.

Fourth, and finally, let F^* be the distribution that achieves the minimum in (3) and write $v_i^* = \langle F^*, w - w_i\rangle$. We claim $\lambda_i = 0$ for all i with $v_i^* > V_i(w)$. To see why, note that $v_j^* = V_j(w)$ for at least one j by hypothesis. Accordingly, the payoff vector $(v_1^*, \dots, v_n^*, \Pi(w))$ lies in the closure of both S and T and thus

$$\mu\Pi(w) = \lambda_1 v_1^* + \cdots + \lambda_n v_n^* + \kappa \leq \lambda_1 V_1(w) + \cdots + \lambda_n V_n(w) + \kappa.$$

Because (i) $v_i^* \geq V_i(w)$ for all i by hypothesis and (ii) $\lambda_i \geq 0$ for all i per the first fact above, it follows immediately that $\lambda_i = 0$ for all i with $v_i^* > V_i(w)$. Altogether, normalizing each multiplier $\mu, \lambda_1, \dots, \lambda_n$ by factor $1/\mu$ yields the claimed equality

$$\Pi(w) = \lambda_1 V_1(w) + \cdots + \lambda_n V_n(w) + \kappa.$$

□

A.5 Theorem 2

Lemma 7. *For every contract w there exists an i such that*

$$\begin{aligned} \langle F_i, w - w_j \rangle &> V_j(w) \quad j \neq i \\ \langle F_i, w - w_j \rangle &= V_i(w) \quad j = i. \end{aligned}$$

Proof. Fix w and take F_i to be a strictly rationalizable empirical action at w. Let $X' = X \cup \{(F_i, w)\}$. Lemma 4 then implies that all non-degenerate cycles through (F_i, w) have strictly positive length. In turn, for any $j \neq i$ and any k

$$\begin{aligned} 0 &< \langle F_i - F_k, w \rangle + s_{jk} + \langle F_j - F_i, w_j \rangle \\ &= \langle F_i, w - w_j \rangle - \langle F_k, w - w_j \rangle + \langle F_j - F_k, w_j \rangle + s_{jk} \\ &= \langle F_i, w - w_j \rangle - \langle F_k, w - w_j \rangle + \Delta(j,k). \end{aligned}$$

The right-hand side of the first inequality is the length of a non-degenerate cycle in X', while remaining inequalities use the definitions of s_{ik}, s_{jk} and $\Delta(j,k)$. It follows that for all $j \neq i$

$$\langle F_i, w - w_j \rangle > \max_k \langle F_k, w - w_j \rangle - \Delta(j,k) = V_j(w)$$

Note that if $j = i$

$$V_i(w) = \max_k \langle F_k, w - w_i \rangle - \Delta(j,k) \geq \langle F_i, w - w_i \rangle,$$

since $\Delta(i,i) = 0$. The reverse inequality holds because F_i is rationalizable at w. $\square$

Corollary 5. *If $w_0 = \sum_i \beta_i w_i$ is a convex combination of empirical contracts with $\beta_i < 1$ for all i, then $\sum_i \beta_i V_i(w_0) < 0$.*

Proof. By Lemma 7, if $\beta_i < 1$ for all i, then there is a F_j such that

$$\sum_i \beta_i V_i(w_0) < \sum_i \beta_i \langle F_j, w_0 - w_i \rangle = 0.$$

$\square$

Lemma 8. *If w^{λ^t,κ^t} is a sequence of equity bonus contracts with $\sum_i \lambda_i^t \to \infty$, then* $\limsup \Pi(w^{\lambda^t,\kappa^t}) \leq \max_{(F_i,w_i)\in X} \Pi(w_i)$.

Proof. Consider a sequence (λ^t, κ^t) of equity bonus contract parameters with $\sum_i \lambda_i^t \to \infty$ and κ^t chosen so that $w^t \equiv w^{\lambda^t,\kappa^t}$ satisfies limited liability for each t. Define

$$\Lambda^t = \sum_i \lambda_i^t \qquad \alpha^t = \frac{1}{1+\Lambda^t} \qquad \beta_i^t = \frac{\lambda_i^t}{\Lambda^t}$$

and note that (i) $\sum \beta_i^t = 1$ for all t; (ii) $\beta_i^t \in [0,1]$ for all i; and (iii) κ^t is bounded. Passing to a subsequence if necessary, we have $(\beta^t, \kappa^t) \to (\beta^*, \kappa^*)$ for some vector of weights β^* and some constant κ^*. Because $\alpha^t \to 0$, w^t converges uniformly on $\mathcal{Y}$ to $w^* \equiv \sum_i \beta_i^* w_i + \kappa^*$. There are two cases to consider. $\square$

Case 1. Suppose $\beta_i^* = 1$ for some i. Because $w^t \rightrightarrows w^*$ and $\Delta(i,j) > 0$ for all $j \neq i$, we have $C_i(w^t) = \{i\}$ for all t sufficiently large and hence $V_i(w^t) = \langle F_i, w^t - w_i \rangle$. Rearranging terms yields the equality

$$\Pi(w^t) = \sum_j \lambda_j^t V_j(w^t) + \kappa^t = \frac{\lambda_i^t}{1+\lambda_i^t} \langle F_i, y - w_i \rangle + \kappa^t$$

and hence $\limsup \Pi(w^t) = \langle F_i, y - w_i \rangle + \kappa^* = \Pi(w_i) + \kappa^*$, where the second equality follows from Corollary 4. Lastly, because $w^t \rightrightarrows w_i$ and $\min w_i = 0$, $\kappa^* \leq 0$ and thus $\limsup \Pi(w^t) \leq \Pi(w_i)$.

Case 2. Suppose $\beta_i^* < 1$ for all i. First, note that

$$\sum_i \lambda_i^t V_i(w^t) + \kappa^t = \Lambda^t \sum_i \beta_i^t V_i(w^t) + \kappa^t.$$

Second, because $w^t \rightrightarrows w^*$ and $(\beta^t, \kappa^t) \to (\beta^*, \kappa^*)$, we have

$$\sum_i \beta_i^t V_i(w^t) \to \sum_i \beta_i^* V_i(w^*) \qquad\qquad \kappa^t \to \kappa^*.$$

Third, from Corollary 5,

$$\sum_i \beta_i^* V_i(w^*) < 0.$$

Altogether, we have $\sum_i \lambda_i^t V_i(w^t) + \kappa^t \to -\infty$.

Proof of Theorem 2. First, suppose that (9) has a solution (λ^*, κ^*) with value $\Pi^* \equiv \sum_i \lambda_i^* V_i(w^*) + \kappa^*$, where we write $w^* \equiv w^{\lambda^*,\kappa^*}$. We have

$$\Pi(w^*) \geq \sum \lambda^* V_i(w^*) + \kappa^* \geq \max_{(F_i, w_i) \in X} \Pi(w_i), \tag{33}$$

where the first inequality follows from (8) and the second follows from the argument given in the first case of the proof of Lemma 8. We claim $\Pi^* \geq \Pi(w)$ for all contracts w, including w^* itself. To see why, suppose to the contrary that $\Pi(\tilde{w}) > \Pi^*$ for some contract $\tilde{w}$. Because $\Pi^* \geq \max_i \Pi(w_i)$ per (33), $\tilde{w}$ is eligible. In turn, Lemma 1 yields multipliers $\tilde{\lambda}$ and a constant $\tilde{\kappa}$ such that $\Pi(\tilde{w}) = \sum_i \tilde{\lambda}_i V_i(\tilde{w}) + \tilde{\kappa}$. Thus, $\sum_i \tilde{\lambda}_i V_i(\tilde{w}) + \tilde{\kappa} > \sum_i \lambda_i^* V_i(w^*) + \kappa^*$. However, this can not be the case because (i) $(\tilde{\lambda}, \tilde{\kappa})$ is feasible for program (9) by construction of $w^{\tilde{\lambda},\kappa}$ and (ii) (λ^*, κ^*) is a solution to that program. We conclude both that w^* is optimal and that $\Pi(w^*) = \Pi^*$.

Second, suppose instead that (9) does not have a solution. Write M for its supremum and let (λ^t, κ^t) be a feasible sequence with

$$\sum_i \lambda_i^t V_i(w^t) + \kappa^t \nearrow M,$$

where we write $w^t \equiv w^{\lambda^t,\kappa^t}$. Because the sum $\sum_i \lambda_i V_i(w^{\lambda,\kappa}) + \kappa$ is continuous in (λ, κ) and the set of feasible solutions to (9) is closed, it must be that at least one of λ^t, κ^t is unbounded. More specifically, because each pair (λ^t, κ^t) is feasible and therefore satisfies limited liability, κ^t is bounded and λ^t must therefore be unbounded. The rest of the argument follows immediately from Lemma 8. □

A.6 Proposition 2

Proof of Proposition 2. We argue in cases.

Case 1. Suppose that (9) has a solution and that its value Π^* strictly exceeds the payoff from the best empirical contract $\max_i \langle F_i, y - w_i \rangle$. By definition, optimal contracts in these cases are eligible. Accordingly, suppose w is an eligible contract that is not an equity bonus contract. Apply Lemma 1 to obtain constants λ, κ such that $\Pi(w) = \sum_i \lambda_i V_i(w) + \kappa$, and consider equity bonus contract $w^{\lambda,\kappa}$. We have already argued in the body that $\Pi(w^{\lambda,\kappa}) \geq \Pi(w)$; we claim now that this improvement is strict under the full support assumption. To see why, recall from (8) that

$$\Pi(w^{\lambda,\kappa}) \geq \sum_i \lambda_i V_i(w^{\lambda,\kappa}) + \kappa \geq \sum_i \lambda_i V_i(w) + \kappa = \Pi(w).$$

Because (i) contracts are continuous by definition; (ii) $w^{\lambda,\kappa}(y) \geq w(y)$ for all y by construction; and (iii) $w^{\lambda,\kappa}$ and w are distinct because w is not an equity bonus contract, there exists an open set $\mathcal{O} \subset \mathcal{Y}$ such that $w^{\lambda,\kappa}(y) > w(y)$ for all y in $\mathcal{O}$. Consequently, $V_i(w^{\lambda,\kappa}) > V_i(w)$ for all i and hence

$$\Pi(w^{\lambda,\kappa}) \geq \sum_i \lambda_i V_i(w^{\lambda,\kappa}) + \kappa > \sum_i \lambda_i V_i(w) + \kappa = \Pi(w).$$

Thus, w is not an optimal contract.

Case 2. Suppose that (9) has a solution and that its value Π^* is exactly the payoff from the best empirical contract $\max_i \langle F_i, y - w_i \rangle$. Now, because $\max_i \langle F_i, y - w_i \rangle > 0$ by assumption, w is optimal only if w yields a positive payoff guarantee. Accordingly, let $\Pi(w) > 0$ and suppose that w is neither an empirical contract nor an equity bonus contract.

Strictly speaking, we can not immediately apply Lemma 1 and replicate the above argument because eligibility is a hypothesis of the Lemma. However, the proof of that Lemma can quickly be strengthened to show that if each F_i has full support, then the desired constants λ, κ exist unless $w = w_i + \gamma$

for some non-negative constant γ. To see why, note that the only place the hypothesis $\Pi(w) > \max_i \Pi(w_i)$ is used in the proof of Lemma 1 (other than via its implication that $\Pi(w) > 0$, which holds here by hypothesis) or any of its dependencies is in Part 4 of the proof of Lemma 6. However, in Part 3 of that same proof, the conclusion that there exists an index i^* and a constant $\gamma \geq 0$ such that $w(y) = w_{i^*}(y) + \gamma$ for all y in $\text{supp}(F_{i^*})$ simplifies under our full support assumption to $w(y) = w_{i^*}(y) + \gamma$ for all y. Thus, either (i) Lemma 1 applies and w can be strictly improved to an equity bonus contract, by the same argument as in the first case; or (ii) $\gamma = 0$ and $w = w_{i^*}$, so that w is itself an empirical contract; or (iii) $\gamma > 0$ and $\Pi(w) = \Pi(w_{i^*}) - \gamma < \Pi(w_{i^*})$, so that w is strictly suboptimal. Because we have excluded the second case by hypothesis, we conclude that w is not an optimal contract.

Case 3. Suppose that (9) does not have a solution and further that w is not an empirical contract. As before, if w is optimal, then $\Pi(w) > 0$; assume as such. Now, if w is an equity bonus contract, then the argument provided in the second step of this proof provides for the existence of constants $\lambda_i + \kappa$ such that $\Pi(w) = \sum \lambda_i V_i(w) + \kappa$. However, because (9) does not have a solution, there is a strictly better set of constants and the corresponding equity bonus contract generated by those contracts yields a strictly better guarantee than w. Accordingly, suppose instead that w is neither an empirical contract nor an equity bonus contract. The same argument as the one used in the second case implies that w is strictly suboptimal.

□

A.7 Corollary 3

Proof of Corollary 3. Corollary 3 almost follows from the uniqueness results in Proposition 2. Because optimal contracts are not weakly dominated by suboptimal contracts, we are left to verify the existence of an undominated contract within the class of optimal contracts. In all three cases of the problem, the set of optimal contracts is compact in the standard topology on $\mathbb{R}^{|X|+1}$ (when equity bonus contracts are viewed as standard convex combinations of the empirical contracts and the output vector). Furthermore, the principal's payoff is a continuous function of output distributions F in the topology of weak convergence, and the set of distributions F satisfying the constraints of Proposition 1 is itself compact in the same topology. Consequently, the problem of identifying an undominated contract can be interpreted as a normal form game between the principal (choosing contracts from the set of optimal contracts) and an adversary (choosing output distributions from the set of distributions that satisfy the constraints of Proposition 1). Standard results based on Zorn's lemma provide for the existence of undominated strategies in normal form games with compact action sets and continuous payoff functions. **TODO**: CITE □

A.8 Proposition 3

Proof of Proposition 3. The requirement in (11) is obvious. For (12) note

$$
\begin{aligned}
\left(\sqrt{\langle F_1, y - w_2\rangle} - \sqrt{\Delta}\right)^2 &\geq \langle F_1, y - w_1\rangle \\
\langle F_1, y - w_2\rangle - 2\sqrt{\langle F_1, y - w_2\rangle \Delta} + \Delta &\geq \langle F_1, y - w_1\rangle \\
\langle F_1, y - w_2\rangle - \Delta &\geq \langle F_1, y - w_1\rangle \\
\langle F_2, w_1 - w_2\rangle &\geq 0
\end{aligned}
$$

where the third line follows from $\sqrt{\langle F_1, y - w_2\rangle \Delta} > \sqrt{\Delta^2} = \Delta$. Similarly,

$$
\begin{aligned}
\left(\sqrt{\langle F_1, y - w_2\rangle} - \sqrt{\Delta}\right)^2 &\geq \langle F_2, y - w_2\rangle \\
\langle F_1, y - w_2\rangle - \Delta &\geq \langle F_2, y - w_2\rangle \\
\langle F_1 - F_2, y\rangle &\geq \langle F_1 - F_2, w_1\rangle.
\end{aligned}
$$

The remaining inequalities follow from $\Delta > 0$. □

A.9 Proposition 5

Proof of Proposition 5. Define $y_i \equiv \langle F_i, y\rangle$ so that $w_i = \alpha_i y_i$. Note that any novel equity bonus contract is a linear contract, as in $w_0 = \alpha_0 \langle F, y\rangle = \alpha_0 y_0$ where y_0 is unknown. Note also that since $\min(y) = 0$, adding a constant can only reduce the principal's guarantee. First, we show that the data can be ordered. For any i and j, there is a cycle length in the data corresponding to the inequality

$$
\langle F_j - F_i, \alpha_j y\rangle + \langle F_i - F_j, \alpha_i y\rangle = (\alpha_i - \alpha_j)(y_i - y_j) \geq 0,
$$

so $\alpha_i > \alpha_j$ if and only if $y_i > y_j$. Index the data so that $i < j$ implies $y_i < y_j$ and hence $\alpha_i < \alpha_j$. Note also that the supermodularity condition holds trivially here.

The cycle constraints in (2) then imply through an analogous inequality that there is an i such that $\alpha_i \leq \alpha_0 \leq \alpha_{i+1}$ and $y_i \leq y_0 \leq y_{i+1}$. Fix that i for the remainder of the proof.

We next show that $y_0 = y_i$ satisfies all of the cycle constraints. To show this, note that for the cycle constraints in (2) to be satisfied it is sufficient for the following inequality to hold for all k and j:

$$\alpha_k(y_k - y_0) + s_{kj} + \alpha_0(y_0 - y_j) \geq 0,$$

and these hold for $y_0 = y_i$ if

$$\alpha_k(y_k - y_i) + s_{kj} + \alpha_0(y_i - y_j) \geq 0.$$

The key observation is that $\alpha_0(y_i - y_j) > s_{ji}$. Using $\delta = \text{sgn}(i - j)$, this follows from

$$\begin{aligned}\alpha_0(y_i - y_j) - s_{ji} &= \alpha_0(y_i - y_j) - \sum_{k=j+\delta}^{i} \alpha_k(y_k - y_{k-\delta}) \\ &= \sum_{k=j+\delta}^{i} (\alpha_0 - \alpha_k)(y_k - y_{k-\delta}) > 0,\end{aligned}$$

where the first equality follows from Lemma 2. Therefore

$$\alpha_k(y_k - y_i) + s_{kj} + \alpha_0(y_i - y_j) \geq \alpha_k(y_k - y_i) + s_{kj} + s_{ji} \geq 0,$$

because the $\alpha_k(y_k - y_i) + s_{kj} + s_{ji}$ is the length of a cycle in the data and nonnegative by assumption. Therefore, by Proposition 1

$$\Pi(w_0) \leq (1 - \alpha_0)y_i \leq (1 - \alpha_i)y_i,$$

and the principal has no opportunity to improve on the contract w_i. □

A.10 Theorem 3

Proof of Lemma 2. First, we claim that there is a shortest path that traverses adjacent nodes. Suppose to the contrary that for any shortest path through nodes indexed by $\{n_1, n_2, \ldots\}$ there is a k such that $|n_{k+1} - n_k| > 1$ (i.e., the path skips adjacent nodes between n_k and

n_{k+1}). Consider any such path and any such node n_k. Notice that supermodularity implies the inequality in

$$\begin{aligned}\langle F_{n_{k+1}} - F_{n_k}, w_{n_{k+1}}\rangle &= \sum_{\ell=0}^{n_{k+1}-n_k-1} \langle F_{n_k+\ell+1} - F_{n_k+\ell}, w_{n_{k+1}}\rangle \\ &\geq \sum_{\ell=0}^{n_{k+1}-n_k-1} \langle F_{n_k+\ell+1} - F_{n_k+\ell}, w_{n_k+\ell+1}\rangle.\end{aligned}$$

For this path to be a shortest path it must be that the inequality holds with equality, but this implies the existence of another shortest path that does not jump from n_k to n_{k+1} but instead traverses the sequence $\{n_k, n_k+1, \ldots, n_{k+1}-1, n_{k+1}\}$, a contradiction.

Next, suppose that there is no shortest path following the particular sequence of adjacent nodes $\{i, i+\iota, i+2\iota, \ldots, j-\iota, j\}$ where $\iota = \text{sgn}(j-i)$. In such a case, any shortest path through adjacent nodes must visit at least one node more than once. Thus there is a cycle in the path, and in order for this path to be shortest this cycle must have zero length. Therefore, there is another shortest path with this cycle removed. By successively removing such cycles we get the claimed shortest path, contradicting the supposition. The final statement of the lemma is an obvious consequence. □

We introduce two more lemmas before proving the result. The first lemma analyzes the following modified (but equivalent) version of (9), where we use the fact that $\kappa = 0$ under two-sided limited liability.

$$\begin{aligned} \max_{\gamma} \quad & \sum_{i,j} \gamma_{ij} \left(\langle F_j, w^\gamma - w_i\rangle - \Delta(i,j)\right) \\ \text{s.t.,} \quad & w^\gamma(y) = \frac{y + \sum_{i,j}\gamma_{ij} w_i(y)}{1 + \sum_{i,j}\gamma_{ij}} \\ & w^\gamma(y) \geq 0, \quad \forall y \in \mathcal{Y}. \end{aligned} \tag{34}$$

Any optimal γ for this problem corresponds to an optimal (λ, κ) for (9) where $\lambda_i = \sum_j \gamma_{ij}$ and $\kappa = 0$. Note that if γ is optimal for (34) and $\gamma_{ij} > 0$ then it must be that $j \in C_i(w^\gamma)$.

Lemma 9. *Let $\gamma \geq 0$ satisfy $\gamma_{ii'} > 0$ and $\gamma_{jj'} > 0$. If*

$$\langle F_{i'} - F_{j'}, w_i - w_j\rangle \geq 0, \tag{35}$$

then there is a $\gamma' \neq \gamma$ with either $\gamma'_{ii'} = 0$ or $\gamma'_{jj'} = 0$ that yields a weakly higher value of (34).

Proof of Lemma 9. Suppose that the optimal γ in (34) is such that $\gamma_{ii'} > 0$ and $\gamma_{jj'} > 0$. Consider the adjustments to γ given by $\gamma_{ii'} + \delta$ and $\gamma_{jj'} - \delta$ for some $\delta \in [-\gamma_{ii'}, \gamma_{jj'}]$. If γ is optimal, it must be optimal to set $\delta = 0$. However, the second-order effect on the objective of increasing δ is

$$\frac{2}{1 + \sum_k \lambda_k} \langle F_{i'} - F_{j'}, w_i - w_j \rangle \geq 0.$$

Therefore, under the supposition of the Lemma the objective is convex in δ and hence is optimized on the boundary, $\delta \in \{-\gamma_{ii'}, \gamma_{jj'}\}$. □

Lemma 10. *Shortest path lengths are submodular in the sense that*

$$s_{i'j'} - s_{i'j} \leq s_{ij'} - s_{ij} \quad \forall i' > i, j' > j.$$

Proof of Lemma 10. Fix j and j' and consider the difference $d_i \equiv s_{ij'} - s_{ij}$. We want to show that d_i is non-increasing in i for any such $j < j'$. Using Lemma 2,

$$d_i = \begin{cases} s_{jj'} & i \leq j < j' \\ s_{ij'} - s_{ij} & j < i < j' \\ -s_{j'j} & j < j' \leq i. \end{cases}$$

It is immediate that $s_{jj'} \geq s_{ij'} - s_{ij} \geq -s_{j'j}$ for any $j \leq i \leq j'$, and consequently, $d_i \geq d_{i'}$ if either $i \leq j$ or $i' \geq j'$. In the remaining case, $j < i < i' < j'$,

$$s_{i'j'} - s_{i'j} - s_{ij'} + s_{ij} = -s_{ii'} - s_{i'i} \leq 0,$$

due to the non-negative cycle length condition. □

Proof of Theorem 3. As in Lemma 9, this proof refers to the problem in (34). Suppose there are no optimal one-contract mixtures, but there is an optimal multi-contract mixture, parameterized by some γ. For such a mixture, it must be that $\gamma_{ii'} > 0$ and $\gamma_{jj'} > 0$ for some $i \neq j$, i', and j'. Without loss of generality, take $i < j$.

It is also without loss to consider cases where $\gamma_{ik} = 0$ for all $k \neq i'$ and $\gamma_{jk} = 0$ for all $k \neq j'$. If, for example, $\gamma_{ik} > 0$ for some $k \neq i'$ and γ is optimal, the problem in (34) would also be optimized by shifting the ik weight to ii' as in

$$\gamma'_{\ell m} = \begin{cases} 0 & \text{if } \ell m = ik \\ \gamma_{ii'} + \gamma_{ik} & \text{if } \ell m = ii' \\ \gamma_{\ell m} & \text{otherwise .} \end{cases}$$

We show that whenever w_i and w_j are involved in the proposed optimal contract, as indicated by $\gamma_{ii'} > 0$ and $\gamma_{jj'} > 0$, there exists another γ' that sets either $\gamma'_{ii'} = 0$ or $\gamma'_{jj'} = 0$ that yields a weakly higher value for the objective. Hence, there is a weakly better mixture contract that removes either w_i or w_j. By successively removing such contracts we arrive at a weakly better 1-mixture contract.

There are two cases. First, if $i' \leq j'$, supermodularity immediately implies the sufficient condition in Lemma 9 for a mixture involving w_i and w_j to be weakly worse than a contract only involving one of w_i or w_j (i.e., $\langle F_{i'} - F_{j'}, w_i - w_j \rangle \geq 0$).

Alternatively, if $i' > j'$, the optimality of γ implies[9]

$$\begin{aligned} \langle F_{i'}, w^{\gamma,\kappa} \rangle - \langle F_i, w_i \rangle - s_{ii'} &\geq \langle F_{j'}, w^{\gamma,\kappa} \rangle - \langle F_i, w_i \rangle - s_{ij'} \\ \langle F_{j'}, w^{\gamma,\kappa} \rangle - \langle F_j, w_j \rangle - s_{jj'} &\geq \langle F_{i'}, w^{\gamma,\kappa} \rangle - \langle F_j, w_j \rangle - s_{ji'}, \end{aligned}$$

which implies

$$s_{jj'} - s_{ji'} \leq \langle F_{j'} - F_{i'}, w^{\gamma,\kappa} \rangle \leq s_{ij'} - s_{ii'}. \tag{36}$$

Together, Lemma 10 and (36) imply

$$\begin{aligned} \langle F_{i'}, w^{\gamma,\kappa} \rangle - \langle F_i, w_i \rangle - s_{ii'} &= \langle F_{j'}, w^{\gamma,\kappa} \rangle - \langle F_i, w_i \rangle - s_{ij'} \\ \langle F_{j'}, w^{\gamma,\kappa} \rangle - \langle F_j, w_j \rangle - s_{jj'} &= \langle F_{i'}, w^{\gamma,\kappa} \rangle - \langle F_j, w_j \rangle - s_{ji'}. \end{aligned}$$

Using the second equation, we find if setting $\gamma_{jj'} > 0$ is optimal, so is taking the weight off

[9]Note for any i and j

$$\langle F_j, w^{\gamma,\kappa} - w_i \rangle - \Delta(i,j) = \langle F_j, w^{\gamma,\kappa} \rangle - \langle F_i, w_i \rangle - s_{ij}$$

of jj' and placing it on ji', meaning

$$\gamma'_{k\ell} = \begin{cases} 0 & \text{if } k\ell = jj' \\ \gamma_{jj'} + \gamma_{ji'} & \text{if } k\ell = ji' \\ \gamma_{k\ell} & \text{otherwise .} \end{cases}$$

Since, γ' satisfies the conditions of Lemma 9 with $\langle F_{i'} - F_{i'}, w_i - w_j \rangle = 0$, we find that there is a weakly better contract that does not involve both of w_i and w_j. □

A.11 Theorem 4

Proof of Theorem 4. Choose $P \in \mathcal{P}$ arbitrarily and enumerate its support $\mathcal{A}_1, \ldots, \mathcal{A}_T$. Because P rationalizes the data, there exists selections (F_{it}, e_{it}) from the choice sets $c(\mathcal{A}_1|w_i), \ldots, c(\mathcal{A}_T|w_i)$ such that $F_i = \sum_t F_{it}P_t$ for each i. Extend the index i on the revealed preference data X by writing $w_0 = w$ for the principal's contract and $(F_{01}, e_{01}), \ldots, (F_{0T}, e_{0T})$ for the elements of the choice sets $c(\mathcal{A}_1|w_0), \ldots, c(\mathcal{A}_T|w_0)$ that minimize the principal's payoff with ties broken in favor of the principal. Observe that

$$\Pi(P|w_0) = \langle F_{01}, y - w_0 \rangle P_1 + \cdots + \langle F_{0T}, y - w_0 \rangle P_T = \langle F_0, y - w_0 \rangle. \tag{37}$$

We proceed by constructing a fixed technology $\mathcal{A} \in \mathscr{A}$ with the property $\Pi(\mathcal{A}|w_0) = \Pi(P|w_0)$. Toward that end, define output distribution $F_0 = \sum_t F_{0t}P_t$ and effort cost vector $(e_0, \ldots, e_n)$ by $e_i = \sum_t e_{it}P_t$. Let $\mathcal{A} = \{(F_0, e_0), \ldots, (F_n, e_n)\}$ and choose i, j arbitrarily. Since $\langle F_{it}, w_i \rangle - e_{it} \geq \langle F_{jt}, w_i \rangle - e_{jt}$ for all t,

$$\begin{aligned} \langle F_i, w_i \rangle - e_i &= \sum_{t=1}^{T} (\langle F_{it}, w_i \rangle - e_{it}) P_t \\ &\geq \sum_{t=1}^{T} (\langle F_{jt}, w_i \rangle - e_{jt}) P_t = \langle F_j, w_i \rangle - e_j. \end{aligned}$$

This establishes two facts. First, it follows immediately that $\mathcal{A}$ rationalizes the data. Second, it follows almost immediately that

$$\langle F_0, y - w_0 \rangle = \Pi(\mathcal{A}|w_0). \tag{38}$$

The only difficulty in the latter lies in ruling out that the existence of an index i with

$$\langle F_i, w_0\rangle - e_i = \langle F_0, w_0\rangle - e_0 \tag{39}$$

$$\langle F_i, y - w_0\rangle > \langle F_0, y - w_0\rangle. \tag{40}$$

To see why this can not be the case, note that

$$(39) \implies \forall t : \langle F_{it}, w_0\rangle - e_{it} = \langle F_{0t}, w_0\rangle - e_{0t}$$

$$(40) \implies \exists t : \langle F_{it}, y - w_0\rangle > \langle F_{0t}, y - w_0\rangle.$$

To see why, observe that if the agent strictly preferred F_{0t} over F_{it} at any component technology t, then he would necessarily strictly prefer F_0 to F_i. Similarly, if the principal strictly preferred F_i to F_0, then she must strictly prefer F_{it} to F_{0t} for at least one component technology t. However, we began by choosing (F_{0t}, e_{0t}) as the principal's preferred element of $c(\mathcal{A}_t|w_0)$. Accordingly, no such index exist, and hence (38) holds. Together, (37) and (38) imply $\Pi(P|w_0) = \Pi(\mathcal{A}|w_0)$, as claimed. In summary, we have shown that for every stochastic technology P that rationalizes the data, there exists a fixed technology $\mathcal{A}$ that (i) also rationalizes the data and (ii) yields the same payoff for the principal as P. Together, these two facts imply $\Pi^{\mathcal{P}}(w) \geq \Pi(w)$. Jointly with Remark 1, this establishes $\Pi^{\mathcal{P}}(w) = \Pi(w)$ for all contracts, as claimed. □