

Combinatorial Contracts Beyond Gross Substitutes*

Paul Dütting[†]

Michal Feldman[‡]

Yoav Gal Tzur[§]

Abstract

We study the combinatorial contracting problem of Dütting et al. [13], in which a principal seeks to incentivize an agent to take a set of costly actions. In their model, there is a binary outcome (the agent can succeed or fail), and the success probability and the costs depend on the set of actions taken. The optimal contract is linear, paying the agent an α fraction of the reward. For gross substitutes (GS) rewards and additive costs, they give a poly-time algorithm for finding the optimal contract. They use the properties of GS functions to argue that there are poly-many “critical values” of α , and that one can iterate through all of them efficiently in order to find the optimal contract.

In this work we study to which extent GS rewards and additive costs constitute a tractability frontier for combinatorial contracts. We present an algorithm that for *any* rewards and costs, enumerates all critical values, with poly-many demand queries (in the number of critical values). This implies the tractability of the optimal contract for any setting with poly-many critical values and efficient demand oracle. A direct corollary is a poly-time algorithm for the optimal contract in settings with supermodular rewards and submodular costs. We also study a natural class of matching-based instances with XOS rewards and additive costs. While the demand problem for this setting is tractable, we show that it admits a super-polynomial number of critical values. On the positive side, we present (pseudo-) polynomial-time algorithms for two natural special cases of this setting. Our work unveils a profound connection to sensitivity analysis, and designates matching-based instances as a crucial focal point for gaining a deeper understanding of combinatorial contract settings.

1 Introduction

Contract theory is one of the pillars of microeconomic theory. In its most prominent model of Holmström [24] and Grossman and Hart [19], a principal seeks to delegate the execution of a task to an agent. The agent may perform one of n costly actions $A = \{1, \dots, n\}$, unobservable by the principal. Each action leads to a stochastic reward (outcome) for the principal. The basic incentive problem is that the agent, who bears the cost of executing the task, does not directly benefit from the stochastic reward that is generated. To incentivize the agent to exert effort, the principal designs a contract that specifies the payments for each outcome. This induces a game between the principal and the agent. Given a contract, the agent replies with an action that maximizes payment minus cost. The principal in turn tries to find the contract that maximizes reward minus payment, given the agent’s best response.

In the vanilla contract problem, where the agent may only choose a single action, the optimal contract can be computed in polynomial time: This is done by computing for each action individually, using a linear program, the minimum payment required to incentivize this action (if at all possible), then pick the best action among those [19].

In many real-life settings, the contract problem is combinatorial in one or more dimensions of the problem. In such settings, the approach above might not yield an efficient algorithm. A recent line of work embarks on the design of efficient algorithms for combinatorial settings [14, 13, 3, 17].

The Model. We study the model introduced by [13] where the principal seeks to hire an agent to run a project, and the agent can perform *any subset* of actions from a fixed set $A = \{1, 2, \dots, n\}$. The project can either succeed or fail. If the project succeeds, the principal gets a reward, which we normalize to be 1, otherwise, the reward is 0. Each subset of actions S induces a success probability of the project, specified by a set function

*Accepted as a “soft merge” with Deo-Campo Vuong, Dughmi, Patel, and Prasad [12]. This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement No. 866132), by the Israel Science Foundation (grant number 317/17), by an Amazon Research Award, and by the NSF-BSF (grant number 2020788). The full version of the paper can be accessed at <https://arxiv.org/pdf/2309.10766.pdf>

[†]Google Research, Zurich, Switzerland. Email: duetting@google.com

[‡]Tel Aviv University, Israel. Email: mfeldman@tauex.tau.ac.il

[§]Tel Aviv University, Israel. Email: yoavgaltzur@mail.tau.ac.il

$f : 2^A \rightarrow [0, 1]$. Since $f(S)$ is also the expected reward of a set of actions, we also refer to f as the reward function. The cost of a set is given by a set function $c : 2^A \rightarrow \mathbb{R}_{\geq 0}$. We assume to have value oracle access to both f and c .

For this binary setting, the optimal contract takes a linear form, where the principal pays a fraction $\alpha \in [0, 1]$ of her reward if the project succeeds and nothing otherwise [7]. In this case, the agent picks a set of actions $S \subseteq A$ which maximizes $\alpha f(S) - c(S)$, and the principal chooses a contract α such that $(1 - \alpha)f(S)$ is maximized — with S being the agent’s best response to contract α . Motivated by a connection to combinatorial auctions, interpreting $f(S)$ as the value for a set of items and $\frac{1}{\alpha}c(S)$ as (bundle) prices, we will also refer to the agent’s best response for a given α , as the agent’s *demand*.

Values of α at which the agent’s demand changes are called *critical values*, and it is easy to observe that the optimal contract α occurs at a critical value.

State of the Art. The main result of [13] is a polynomial time algorithm when the success probability function f is gross substitutes (GS), and the cost function c is additive.

To prove this result the authors of [13] show that for GS f and additive c there can be at most $O(n^2)$ critical values of α . In addition, they exploit that for GS valuations and item prices the agent’s demand problem can be solved greedily [36], to obtain an efficient algorithm that given a critical value of α finds the next critical value $\alpha' > \alpha$. Thus, the optimal contract can be solved by iterating efficiently through critical values of α , and choosing the value of α that maximizes the principal’s utility. Crucially, the properties of GS are used both to provide an upper bound on the number of critical values, and to iterate efficiently through these values.

For the case in which f is submodular and c is additive, the authors of [13] show that there may be exponentially many critical values of α . Furthermore, they show that computing the optimal contract for this case is NP-hard.

1.1 Our Contribution In this work, we explore to which extent gross substitutes success probability f and additive costs c constitutes a tractability frontier for combinatorial contracts.

Enumerating All Critical Values with a Demand Oracle. As a first result, we present an algorithm that for *any* f and c finds all critical values of α , using polynomially (in the number of critical values) many calls to an algorithm for the agent’s demand problem.

Proposition (Proposition 3.1). Given oracle access to the agent’s demand, there exists an algorithm which returns all the critical values. The running time of this algorithm is linear in the number of critical values.

An immediate algorithmic implication of this result is an efficient algorithm for the optimal contract problem in scenarios with a small (i.e., polynomial) number of critical values, given access to a demand oracle. In scenarios with a poly-time algorithm for a demand query, this gives a poly-time algorithm for the optimal contract.

Our algorithm is recursive. It finds the critical values by querying the agent’s demand oracle at two end points of a given segment $[\alpha, \alpha'] \subseteq [0, 1]$. If the agent’s best response for the two contracts is the same, i.e., $S_\alpha = S_{\alpha'}$, there are no critical values in (α, α') . Otherwise, the oracle answers differ, i.e., $S_\alpha \neq S_{\alpha'}$, and the procedure is applied recursively for the sub-segments $[\alpha, \gamma]$ and $[\gamma, \alpha']$, where γ is the contract for which the agent is indifferent between S_α and $S_{\alpha'}$. That is, γ is such that $\gamma f(S_\alpha) - c(S_\alpha) = \gamma f(S_{\alpha'}) - c(S_{\alpha'})$.

This algorithm has been discovered before in different contexts, in particular in the field of sensitivity analysis of combinatorial optimization problems [22], where it is referred to as the Eisner-Severance technique [18]. We use this algorithm, and the connection to sensitivity analysis more generally, to make progress on the tractability frontier for optimal contracts beyond gross substitutes rewards and additive costs.

Complementary Actions. A first setting for which we show the tractability of the optimal contract is the case of supermodular f and submodular c . This captures complementary actions with non-decreasing marginal rewards and non-increasing marginal costs.

Complementary actions appear in many real-world applications, specifically when the agent can perform actions that are infrastructural in nature (e.g., buying a computer, learn a new technology) and the infrastructure can be used for various purposes.

In this setting, the agent’s demand problem can be solved efficiently using value queries (via reduction to submodular minimization). We show that the number of critical values in this setting is at most n . In combination with Proposition 3.1 we thus obtain:

Theorem (Theorem 4.1). If the success probability function is monotone supermodular, and the cost function is monotone submodular, the optimal contract can be computed in polynomial time.

Matching-based Rewards and Costs. In the second setting that we consider, actions are edges in a bipartite graph $G = (V \cup U, A)$. We think of the right-hand side vertices $u \in U$ as resources, and the left-hand side vertices $v \in V$ as tasks. Every matching in G is associated with a success probability, and the success probability for a set of edges S is the maximal success probability of a matching contained in S . Similarly, each set of edges S (independent of whether it forms a matching or not) is associated with a cost.

A natural starting point, and the focus of this work, is where each edge is associated with a cost $c_{v,u}$ and a success probability $f_{v,u}$. The cost for a set of edges S is given by the sum of costs $c_{v,u}$ of $(v, u) \in S$. The success probability of a set of edges S is the max-weight matching in the subgraph induced by S , with respect to weights $\{f_{v,u}\}_{(v,u) \in A}$. In this case, the success probability $f(S)$ is XOS, but not necessarily submodular.¹

The agent's best response in this setting corresponds to a max-weight matching computation, so a demand oracle can be computed efficiently. However, as our main technical result, we show a super-polynomial lower bound on the number of critical values. We prove this lower bound through a non-trivial reduction from the parametric shortest $s - t$ path problem in sensitivity analysis, and a family of hard instances due to Carstensen [8] and Mulmuley and Shah [34], to an instance of the agent's best response problem.

Theorem (Theorem 5.2). For matching based rewards and additive costs, the number of critical values may be super-polynomial.

We also show tractability of the optimal contract for two important special cases. First, we utilize the work of Carstensen [8] to derive a quasi-polynomial upper bound on the number of critical values where costs and rewards are integers (up to scaling required for normalization).

Proposition (Corollary 5.2). For matching based rewards and additive costs, when costs and rewards are integers, there is a quasi-polynomial time algorithm for finding the optimal contract.

We also give a polynomial-time algorithm for the case where the cost of each edge depends only on one side of the graph, i.e., all edges incident to the same vertex in, say U , have the same cost. We prove this result by a non-immediate reduction to the case of gross substitutes f [13]. The crux of this argument is that while f when viewed as a function of a subset of edges is not gross substitutes (and not even submodular), in this special case it's possible to reformulate the agent's utility problem in terms of the vertices of U with additive costs and an OXS (and hence gross substitutes) reward function (see Section 2 for definitions).

Theorem (Theorem 5.4). For matching based rewards and additive costs when the costs depend only on one side, the optimal contract can be computed in polynomial time.

Our work unveils a profound connection between contracts and sensitivity analysis. It also designates matching-based instances as a pivotal focal point beyond gross substitutes (and even beyond submodular) for further research in this field. This is because they admit an efficient demand oracle, but at the same time they may have a super-polynomial number of critical values (as we show in Theorem 5.2). Our work thus opens up two exciting directions for future work: either (1) finding a poly-time algorithm for computing the optimal contract (which would inevitably have to sidestep enumerating all critical values), or (2) showing a hardness result (which would necessarily not rely on the hardness of demand queries). Either of these solutions would contribute to a deeper comprehension of the overall landscape.

1.2 Related Work The classic hidden-action principal agent problem was formulated by Holmström [24] and Grossman and Hart [19], who also gave a poly-time algorithm for finding the optimal contract based on linear programming. For the types of structured contracting settings as we study them in this work, this approach however does not yield a poly-time algorithm. Holmström and Milgrom [25] consider a dynamic contracting setting, and show that linear contracts are optimal in that setting. In more recent work, Carroll [7] shows the (max-min) robust optimality of linear contracts. In his model, the principal knows a fixed subset of the agent's

¹For each matching $M \subseteq A$ in the graph G , let $x_M \in \mathbb{R}_{\geq 0}^{|A|}$ be its corresponding vector, with $x_M(a) = f_a$ if $a \in M$ and 0 otherwise. Then f is XOS, as $f(S) = \max_M x_M \cdot \chi_S$, where χ_S is the characteristic vector of the set S .

actions and show that the max-min optimal contract is linear. Dütting et al. [16] show the max-min optimality of linear contracts also hold when the principal is aware of the agent's actions, but only the expected value of the outcome distributions are known. They also discuss how well linear contracts approximate the optimal one. In our model with binary outcomes linear contract are optimal.

An exponential blow-up in complexity also naturally arises when the principal seeks to hire a *team of agents*. Babaioff et al. [3] introduced the combinatorial agency model, in which a team of agents, each chooses whether to exert effort or not, is incentivized by the principal. The outcome depends on the different combination of effort-making agents, whose structure is also expressed by a set function. In later work Babaioff et al. [5], [4] study the effects of mixed strategy and free-riding in this setting. These earlier works focused on the case where the success probability is encoded as a read-once network. The recent work of Dütting et al. [14] studies complement free reward functions – XOS, submodular and subadditive. The key result of that work are poly-time algorithms for XOS and submodular rewards that yield $O(1)$ -approximations to the optimal contract with poly-many demand and value oracle calls. A related but different multi-agent contracting problem was recently studied by Castiglioni et al. [11], who consider the case where the principal can observe each agent's individual outcome (and base her contracts on this information).

The exponential blow-up can stem also from a *structured outcome space*, which is studied by Dütting et al. [17]. They consider scenarios with a larger outcome space (of size m), whose description size is $\text{poly}(n, \log(m))$. The goal is to find an optimal or near-optimal contract in time polynomial in the succinct description of the instance.

Ho et al. [23], Zhu et al. [38], and Dütting et al. [15] consider a contract design question from an online learning perspective. Kleinberg and Kleinberg [30] and Kleinberg and Raghavan [31] consider problems that can be thought of as contract design without money. Guruganesh et al. [20, 21], Castiglioni et al. [9, 10] and Alon et al. [2, 1] consider settings in which screening (that is, hidden types) is combined with moral hazard (hidden actions).

2 Model and Notation

We consider a principal-agent setting, where the agent can choose any subset of a given action set $A = \{1, 2, \dots, n\}$, and the outcome is binary (success or failure). A success outcome gives the principal a reward $r \in \mathbb{R}_{\geq 0}$ (and a failure gives the principal zero reward). Every set of actions $S \subseteq A$ is associated with a cost (incurred by the agent) and a success probability, which are given by set functions $c : 2^A \rightarrow \mathbb{R}_{\geq 0}$ and $f : 2^A \rightarrow [0, 1]$, respectively.

The principal cannot observe the agent's actions, only the final outcome. Thus, the principal incentivizes the agent to take costly actions by offering a contract $t : \{0, 1\} \rightarrow \mathbb{R}_{\geq 0}$, which specifies a payment for every outcome. As standard in the literature, we impose a limited liability constraint, meaning that payments go only from the principal to the agent. The agent observes the contract t , and responds by taking a set of actions S .

The principal's goal — our goal in this paper — is to find the contract that maximizes her utility, given that the agent responds with the action set that maximizes the agent's utility. This contract is referred to as the *optimal contract*.

In this binary outcome setting, it is without loss of generality to assume that $t(0) = 0$ [13]. Furthermore, we normalize the reward in case of a success to 1, i.e., $r = 1$, so $f(S)$ is the expected reward for a set of actions S . In this normalized setting, we often refer to f as the *reward function*. A contract can now be described using a single parameter $\alpha \in [0, 1]$, which is the fraction of the reward ($r = 1$) that is paid to the agent, that is, $t(1) = \alpha$. The agent and principal's utilities are then given by:

$$u_a(\alpha, S) = \alpha f(S) - c(S) \quad \text{and} \quad u_p(\alpha, S) = (1 - \alpha)f(S).$$

Given a contract α , the agent's *best response*, is a set of actions $S^* \subseteq A$ that maximizes $u_a(\alpha, S)$ among all $S \subseteq A$. As standard in the literature, we assume that the agent breaks ties in favor of the principal. Namely, for any two sets S, S' such that $u_a(\alpha, S) = u_a(\alpha, S')$, the agent prefers the set with the higher value of f .

Drawing from the combinatorial auctions literature, we refer to the set of sets that maximize the agent's utility under a given contract α as the agent's *demand* under this contract. By employing a consistent tie-breaking rule, we can narrow this collection down to a single set, S_α . We can now write the agent and principal's utilities under a contract α as $u_a(\alpha) = u_a(\alpha, S_\alpha)$ and $u_p(\alpha) = u_p(\alpha, S_\alpha)$.

Classes of Set Functions. For any set function $f : 2^A \rightarrow \mathbb{R}$, we denote by $f(a | S)$ the *marginal value* of $a \in A$ given the set $S \subseteq A$, that is $f(a | S) = f(S \cup \{a\}) - f(S)$. We consider the following set functions.

- f is *supermodular* if for every two sets $S, T \subseteq A$ such that $S \subseteq T$ and every $a \in A$, $f(a | S) \leq f(a | T)$.
- f is *XOS* if it is a maximum over additive function, i.e. there exists $f_1, \dots, f_k \in \mathbb{R}^{|A|}$ such that for any $S \subseteq A$, $f(S) = \max_i f_i \cdot \chi_S$, where $\chi_S \in \{0, 1\}^{|A|}$ is the characteristic vector of S .
- f is *submodular* if for every two sets $S, T \subseteq A$ such that $S \subseteq T$ and every $a \in A$, $f(a | S) \geq f(a | T)$. Note that f is submodular if and only if $-f$ is supermodular.
- f is *gross-substitutes (GS)* if for any two price vector $p, q \in \mathbb{R}_{\geq 0}^{|A|}$, with $p \leq q$ (point-wise), For every set $S \in \mathcal{D}(p)$ there exists a set $T \in \mathcal{D}(q)$ such that T contains every action a such that $p_a = q_a$, where $\mathcal{D}(p) = \max_{S \subseteq A} f(S) - p \cdot \chi_S$.²
- $f : 2^U \rightarrow \mathbb{R}_{\geq 0}$ is *OXS* if there exists a bipartite graph $G(V \cup U, E)$ with non-negative weights on the edges, such that for every $T \subseteq U$, $f(T)$ is the total weight of the max weight matching on the subgraph induced by T .

The last four classes are part of the complement-free hierarchy [32], where containment is strict (i.e., $OXS \subset GS \subset Submodular \subset XOS$).

2.1 Structure of Optimal Contracts We begin by making a few observations on the structure of the utilities of both the agent and the principal. For a fixed set of actions $S \subseteq A$, the agent’s utility, $u_a(\alpha, S) = \alpha f(S) - c(S)$, is an affine function of α . Thus, the utility of the agent $u_a(\alpha) = \alpha f(S_\alpha) - c(S_\alpha)$, is the upper envelope of these 2^n affine functions. As such, it is continuous, monotonically non-decreasing, convex and piece-wise linear, as illustrated in Figure 1.




Figure 1: The utilities of the principal and the agent as a function of α (critical values marked in green). $A = \{1, 2\}$, the rewards are $f(\emptyset) = 0$, $f(\{1\}) = 0.2$, $f(\{2\}) = 0.4$, $f(\{1, 2\}) = 0.6$ and the costs $c(\emptyset) = 0$, $c(\{1\}) = 0.05$, $c(\{2\}) = 0.15$, $c(\{1, 2\}) = 0.3$.

The values of α at which the agent’s demand changes are called *critical values*; the set of critical values is denoted by $\mathcal{C}_{f,c}$.

For every interval of α between two critical values, the agent’s demand is fixed at some set S , and the principal’s utility is $u_p(\alpha) = (1 - \alpha)f(S)$. Thus, the principal’s utility function is linear-in-parts, and the best contract within each such interval is its left-most point (which is a critical value). Indeed, within the interval, the agent’s demand does not change, and α increases, thus the principal’s utility $u_p(\alpha) = (1 - \alpha)f(S)$ decreases. This immediately implies that the optimal contract must be a critical value. This is cast in the following observation.

OBSERVATION 2.1. (DÜTTING ET AL. [13]) *For every cost and reward functions $c : 2^A \rightarrow \mathbb{R}_{\geq 0}$ and $f : 2^A \rightarrow [0, 1]$, there exist $k \leq 2^n + 1$ critical values, $0 = \alpha_0 < \alpha_1 < \dots < \alpha_k \leq 1$, such that*

²In auction literature this is known as the demand for f under price p .

1. For every contract $x \in [0, 1]$, $S_x = S_{\alpha_{i(x)}}$, where $i(x) = \arg \max_j \{\alpha_j \mid \alpha_j \leq x\}$.
2. For every $0 \leq i < j \leq k$, $f(S_{\alpha_i}) < f(S_{\alpha_j})$. Furthermore, $c(S_{\alpha_i}) < c(S_{\alpha_j})$ under consistent tie-breaking.
3. The optimal contract, $\alpha^* = \arg \max_{\alpha \in [0, 1]} u_p(\alpha^*)$, satisfies $\alpha^* \in \{\alpha_0, \alpha_1, \dots, \alpha_k\} = \mathcal{C}_{f,c}$.

Observation 2.1 simplifies the optimal contract problem as it allows us to restrict attention to the set of critical values. And yet, even with this restriction, the problem is far from being trivial, due to two main challenges. First, an instance may admit exponentially many critical values. Second, unlike the agent's utility, the principal's utility is not well-behaved. For example, it may be neither monotone nor continuous in α .

Note also that the cost and reward functions, c and f , are set functions of exponential size (in the number of actions n). We consider two oracle models to these functions, named value and demand oracles. A value oracle for f (respectively, for c) receives a set $S \subseteq A$ and returns $f(S)$ (resp., $c(S)$). A demand oracle for f, c receives a contract $\alpha \in [0, 1]$ and returns the agent's best response, S_α . Note that this corresponds to a demand query with bundle prices, where the value function is $f(S)$ and the price of a bundle S is $c(S)/\alpha$. In the case where c is additive we obtain the canonical definition of a demand query for f with item prices.

2.2 A General Scheme for the Optimal Contract Following the above observations, the authors in [13] suggest a scheme for computing the optimal contract (Algorithm 4.1 in their paper): Going over all critical values of α and pick the one that maximizes the principal's utility. This algorithm runs in polynomial time if the following three conditions hold:

1. The number of critical values, $|\mathcal{C}_{f,c}|$, is polynomial in n .
2. For every $\alpha \in [0, 1]$ there is an efficient algorithm for computing the agent's demand.
3. For every $\alpha \in [0, 1]$ there is an efficient algorithm for finding the next critical value.

This paradigm is used in [13] to find the optimal contract when rewards are gross substitutes and costs are additive. For establishing property 3, the authors of [13] rely on the greedy implementation of the agent's demand oracle [6, 36]. For the larger class of submodular rewards, they exploit that the demand oracle problem is NP-hard (i.e., property 2. is violated), to show that the problem of finding the optimal linear contract is NP-hard. They also give a submodular function for which there are exponentially many critical values (i.e., property 1. is violated).

3 Enumerating all Critical Values

In this section we present an efficient algorithm for finding all critical values. That is, given oracle access to the agent's demand, the algorithm returns all values of α at which the agent's demand changes. The running time of the algorithm is linear in the number of critical values. The algorithm doesn't assume anything regarding the structure of the agent's demand problem, nor the implementation of its oracle.

This algorithm has been discovered several times in different contexts. In particular, it has been used by Gusfield [22] in the context of sensitivity analysis, where it is referred to as the Eisner-Severance technique [18]. Here the problem is, given a parameterized combinatorial optimization problem, find all the breakpoints where the optimal solution changes. These breakpoints correspond to the critical values in our contract problem.

We thus show:

PROPOSITION 3.1. *Let $c : 2^A \rightarrow \mathbb{R}_{\geq 0}$ and $f : 2^A \rightarrow [0, 1]$ be two functions that describe the cost and reward of a given set of actions, such that $|\mathcal{C}_{f,c}|$ is polynomial in $|A| = n$. Assume further that there exists an efficient algorithm to compute the agent's demand, then the optimal contract can be computed efficiently.*

Note how this theorem reduces the requirements for the general scheme of [13] from three to two conditions by showing that properties 1. and 2. implies property 3. Before proving Proposition 3.1, we describe the algorithm for enumerating all critical values and prove its correctness.

Recall the agent's utility function $u_a : [0, 1] \rightarrow \mathbb{R}$, given its best response to the contract α , where we break ties towards the set with the higher reward and in some consistent manner:

$$u_a(\alpha) = \max_{S \subseteq A} \alpha f(S) - c(S) = \max_{S \subseteq A} u(\alpha, S)$$

It is clear that u_a is piece-wise linear as a function of α . Moreover, as the maximum over linear functions, it also exhibits convexity, as established in the following lemma

LEMMA 3.1. *Let $S_\alpha, S_\beta \subseteq A$ be two different sets that maximize the agent's utility for two different contracts $0 \leq \alpha < \beta \leq 1$. Then,*

1. $f(S_\alpha) < f(S_\beta)$
2. $c(S_\alpha) < c(S_\beta)$

Proof. The fact that S_α is superior for contract α and S_β is superior for β implies that

$$\begin{aligned}\alpha f(S_\beta) - c(S_\beta) &\leq \alpha f(S_\alpha) - c(S_\alpha) \\ \beta f(S_\beta) - c(S_\beta) &\geq \beta f(S_\alpha) - c(S_\alpha)\end{aligned}$$

Rearranging,

$$\alpha[f(S_\beta) - f(S_\alpha)] \leq \beta[f(S_\beta) - f(S_\alpha)]$$

Together with the fact that $\alpha < \beta$ we get that $f(S_\alpha) \leq f(S_\beta)$. Due to consistent tie breaking, it must be that $f(S_\alpha) < f(S_\beta)$. For the second part of the claim, observe that if $c(S_\alpha) \geq c(S_\beta)$, then

$$u(\alpha, S_\beta) = \alpha f(S_\beta) - c(S_\beta) \geq \alpha f(S_\alpha) - c(S_\alpha) = u(\alpha, S_\alpha)$$

Which contradicts the optimality of S_α . \square

Our algorithm for enumerating the set of critical values for a given segment, $[\alpha, \beta]$, makes use of the utility's piece-wise linear structure to query the agent's demand oracle with contracts that may be critical. Because critical values are contracts for which the agent's demand change, those are points of indifference between two sets of actions. Namely, for any critical point γ there exists at least two sets of actions $S, T \subseteq A$ with $u_a(\gamma, S) = u_a(\gamma, T)$. This observation, together with the linear structure, allow us to identify all critical points using a recursive strategy:

Starting with the end-points of the segment $[\alpha, \beta]$, we query the oracle. If $S_\alpha = S_\beta$, there are no critical values in this segment. If the sets S_α and S_β differ, then we have a critical value in this segment. If there is a single critical value, it can only be the point of indifference between those two sets: $\gamma = \frac{c(S_\beta) - c(S_\alpha)}{f(S_\beta) - f(S_\alpha)}$ and the oracle will return $S_\gamma = S_\beta$ (by tie breaking consistency). If there is more than one critical value, we will get $S_\gamma \neq S_\beta$, and the critical values can be found by solving for both sub-segments $[\alpha, \gamma]$ and $[\gamma, \beta]$. The full description of the algorithm is presented below.

Algorithm 1 Find All Critical Values in a Segment. $\mathbf{CV}(\alpha, \beta)$

```

1: if  $S_\alpha = S_\beta$  then
2:   return  $\emptyset$ 
3: end if
4:  $\gamma \leftarrow \frac{c(S_\beta) - c(S_\alpha)}{f(S_\beta) - f(S_\alpha)}$   $\triangleright \gamma$  is the intersection between  $u_a(\cdot, S_\alpha)$  and  $u_a(\cdot, S_\beta)$ 
5: if  $S_\gamma = S_\beta$  then
6:   return  $\{\gamma\}$ 
7: end if
8:  $C_1 \leftarrow \mathbf{CV}(\alpha, \gamma)$ 
9:  $C_2 \leftarrow \mathbf{CV}(\gamma, \beta)$ 
10: return  $C_1 \cup C_2$ 

```

CLAIM 3.1. *For any segment $[\alpha, \beta]$, $\mathbf{CV}(\alpha, \beta)$ returns all the critical values in that segment.*

Proof. We prove the claim by induction over the number of critical values in the segment.

For the basis of the induction, assume there are no critical values in $[\alpha, \beta]$. Thus, it must be that the agent's utility is linear in that segment, by the consistency of tie breaking $S_\alpha = S_\beta$ — and the algorithm returns \emptyset .

If there exists a single critical value, since the agent's utility is piece-wise linear, $S_\alpha \neq S_\beta$ and the critical value is at the intersection of the linear functions $u_a(x, S_\alpha) = xf(S_\alpha) - c(S_\alpha)$ and $u_a(x, S_\beta) = xf(S_\beta) - c(S_\beta)$, which is exactly $\gamma = \frac{c(S_\beta) - c(S_\alpha)}{f(S_\beta) - f(S_\alpha)}$. Note that this implies that $S_\gamma = S_\beta$, as ties are broken in favor of the principal.

Assume the correctness of the algorithm for any a segment with $k > i \geq 1$ critical values and consider the segment $[\alpha, \beta]$ with k critical values. The inductive step will show that there is at least one critical value on each side of γ , so each of the sub-segments $[\alpha, \gamma]$ and $[\gamma, \beta]$ have strictly less than k critical values. Thus, when calling $\text{CV}(\alpha, \gamma)$ and $\text{CV}(\gamma, \beta)$, by the induction hypothesis, the critical values for each of these sub-segments are returned. As implied by the base of the induction, if γ is a critical value it will emerge from the $\text{CV}(\alpha, \gamma)$ branch in the calls tree.

Let δ_{\min} be the smallest critical value in $(\alpha, \beta]$. As $S_{\delta_{\min}}$ is the indifference point between S_α and $S_{\delta_{\min}}$, $\delta_{\min} = \frac{c(S_{\delta_{\min}}) - c(S_\alpha)}{f(S_{\delta_{\min}}) - f(S_\alpha)}$. Aiming for contradiction, assume $\delta_{\min} \in [\gamma, \beta]$. Thus, S_α dominates $S_{\delta_{\min}}$ in the segment $[\alpha, \gamma]$ and at γ :

$$u_a(S_\beta, \gamma) = u_a(S_\alpha, \gamma) \geq u_a(S_{\delta_{\min}}, \gamma)$$

By Lemma 3.1, $f(S_\beta) > f(S_{\delta_{\min}})$ and so S_β dominates $S_{\delta_{\min}}$ in $[\gamma, \beta]$, contradicting the optimality of $S_{\delta_{\min}}$.

An analogous argument can be applied to show the existence of a critical value in the segment $(\gamma, \beta]$: let δ_{\max} be the largest critical value in $[\alpha, \beta]$. Observe that $S_{\delta_{\max}} = S_\beta$, and let S be the optimal set just preceding S_β (that is, for every sufficiently small $\varepsilon > 0$, $S_{\delta_{\max} - \varepsilon} = S$). Aiming for contradiction, assume $\delta_{\max} \in [\alpha, \gamma]$. As δ_{\max} is the indifference point between S and S_β , and $f(S_\beta) > f(S)$, it holds that S_β dominates S in the segment $[\gamma, \beta]$. Under contract γ :

$$u_a(S_\alpha, \gamma) = u_a(S_\beta, \gamma) \geq u_a(S, \gamma)$$

Combined with $f(S) > f(S_\alpha)$ (by Lemma 3.1), we get that for any $\delta < \gamma$:

$$\begin{aligned} u_a(S_\alpha, \delta) &= u_a(S_\alpha, \gamma) - (\gamma - \delta)f(S_\alpha) \\ &\geq u_a(S, \gamma) - (\gamma - \delta)f(S_\alpha) \\ &> u_a(S, \gamma) - (\gamma - \delta)f(S) \\ &= u_a(S, \delta) \end{aligned}$$

This implies that S is never optimal, a contradiction. \square

CLAIM 3.2. *The above algorithm takes linear time in the number of critical values in $[\alpha, \beta]$.*

Proof. This can also be shown using an inductive argument. If the segment $[\alpha, \beta]$ has $n > 1$ critical values in it, by the description of the algorithm the recurrence relation describing the time complexity is of the form

$$T(n) = c_0 + T(n - k) + T(k)$$

For some $1 \leq k < n$, where c_0 is the number of operation made in lines 1 – 7 of Algorithm 1. The relation is true because each critical value is found in exactly one of the two sides of γ^3 , and there is at least one critical value on each side.

For the basis of the induction, observe that $T(0), T(1) \leq c_0$. Thus, if we assume for any $l < n$ that $T(l) \leq c_0(2l - 1)$, then by the above recurrence relation

$$T(n) = c_0 + T(n - k) + T(k) \leq c_0 + c_0(2(n - k) - 1) + c_0(2k - 1) = c_0(2n - 1),$$

as claimed. \square

Proof. [of Proposition 3.1] By claims 3.1 and 3.2, given an oracle access to the agent's demand and a polynomial number of critical values, we can efficiently find all critical values and by observation 2.1 one if these values constitutes the optimal contract. \square

The sensitivity literature has also studied the number of breakpoints of several parameterized combinatorial optimization problems, typically with *affine* parametrization [e.g., 8, 22, 29]. This is closely related to determining the number of critical values of α in linear contracts for combinatorial contracting problems. Our result (Proposition 3.1) gives an immediate algorithmic implication of a small (polynomial) number of critical values, namely a polynomial-time algorithm for finding the optimal linear contract.

³if γ is a critical point, it is processed in $\text{CV}(\alpha, \gamma)$, as discussed in the proof for correctness.

4 Submodular Costs and Supermodular Rewards

In this section, we use Proposition 3.1 to show that the optimal contract problem for complementary actions can be found in polynomial time. We model complementary actions using monotone submodular cost functions and monotone supermodular reward functions. In this case, the marginal cost of an action is non-increasing in the set of actions already taken and the marginal reward is non-decreasing.

THEOREM 4.1. *Let $c : 2^A \rightarrow \mathbb{R}$ be a monotone submodular function and let $f : 2^A \rightarrow [0, 1]$ be a monotone supermodular function, then there exists an efficient algorithm for computing the optimal contract.*

We will use Proposition 3.1 and show the two sufficient conditions specified to guarantee the existence of an efficient algorithm for this problem: polynomially-many critical values and an efficient algorithm for the agent's demand.

Tractability of the Agent's Demand Problem. When the agent's utility function, $u(\alpha, S) = \alpha f(S) - c(S)$, is such that c is submodular, and f is supermodular it can be easily seen that for any fixed α , $u(\alpha, \cdot)$ is supermodular. As a demand oracle for the agent requires finding the set $S_\alpha \subseteq A$ which maximizes $u(\alpha, \cdot)$, this is equivalent to finding S_α which minimizes the submodular function $-u(\alpha, \cdot)$. It is well known that there exists a polynomial time algorithm for submodular function minimization, commonly known as SFM (e.g. [28]). Thus, the agent's problem can be solved in polynomial time and a demand oracle for the agent can be implemented efficiently.

CLAIM 4.1. *There exists an efficient algorithm that computes the demand of the agent for any contract α .*

Linear Number of Critical Values. The supermodularity of the agent's utility does not only imply the existence of an efficient demand oracle, but also that there are no more than n critical values. We show that in fact, the set of actions which is the agent's best response may only be extended as we increase the contract α . This immediately implies that the number of optimal sets cannot surpass n .

CLAIM 4.2. *Let c be a monotone submodular cost function and f a monotone supermodular reward function, then for any two contracts $\alpha < \alpha'$ and two corresponding sets in the agent's demand $S_\alpha, S_{\alpha'}$ it holds that $S_\alpha \subseteq S_{\alpha'}$.*

Proof. If $S_\alpha = S_{\alpha'}$ the claim obviously hold. Otherwise, assume that $S_{\alpha'}$ is a maximal best-response for contract α' (this is in line with our tie-breaking assumption), and also that $S_\alpha \setminus S_{\alpha'} = R$ is such that $R \neq \emptyset$, we will show that a contradiction is reached. By the fact that S_α is optimal for α , it must be that

$$u_a(\alpha, R \mid S_\alpha \cap S_{\alpha'}) = u_a(\alpha, S_\alpha) - u_a(\alpha, S_\alpha \cap S_{\alpha'}) \geq 0$$

By the supermodularity of f and submodularity of c it holds that $f(R \mid S_\alpha \cap S_{\alpha'}) \leq f(R \mid S_{\alpha'})$ and $c(R \mid S_\alpha \cap S_{\alpha'}) \geq c(R \mid S_{\alpha'})$. Putting everything together we get

$$\begin{aligned} u(\alpha', R \mid S_{\alpha'}) &= \alpha' f(R \mid S_{\alpha'}) - c(R \mid S_{\alpha'}) \\ &\geq \alpha' f(R \mid S_\alpha \cap S_{\alpha'}) - c(R \mid S_\alpha \cap S_{\alpha'}) \\ &\geq \alpha f(R \mid S_\alpha \cap S_{\alpha'}) - c(R \mid S_\alpha \cap S_{\alpha'}) \\ &= u(\alpha, R \mid S_\alpha) \\ &\geq 0, \end{aligned}$$

where the second inequality follows from the monotonicity of f , which imply $f(R \mid S_\alpha \cap S_{\alpha'}) \geq 0$. Thus, we can add R to $S_{\alpha'}$ while not losing utility, contradicting its maximality. \square

It is worth mentioning that it is possible to prove Theorem 4.1 via the blueprint in [13], which requires an efficient algorithm for solving the agent's demand problem, a polynomial number of critical values, and a polynomial time procedure to compute the next critical value.

The existence of the first two conditions is established in Claim 4.1 and Claim 4.2 above. The task of finding the next critical value can be solved by adopting ideas from Nagano [35]. Nagano's goal is to solve the minimum ratio problem, a natural question when considering parametric submodular minimization, and they accomplish this using the parametric search technique of Meggido [33].

Using Nagano’s approach, finding the next critical value requires following a submodular function minimization algorithm which is fully-combinatorial (i.e. only uses additions, subtractions and comparisons), and for every comparison it makes run another instance of any SFM algorithm. The best known fully combinatorial SFM algorithm runs in $\tilde{O}(n^8)$ steps [26]. Our approach is both simpler and more efficient, since we compute all critical values by calling $O(n)$ times to a demand oracle, which can be implemented using any SFM algorithm (e.g. [27], which runs in $\tilde{O}(n^7)$).

5 Matching-Based Rewards and Additive Costs

In this section we consider a settings in which an agent matches between tasks and resources. This is captured by a bipartite graph $G = (V \cup U, A)$ where V is the set of tasks and U is the set of resources. An edge $(v, u) \in A$ indicates that resource u can be used to complete task v . Each such edge is associated with cost $c_{v,u} \geq 0$ and reward $f_{v,u} \geq 0$. For a set $S \subseteq A$ of edges (actions) chosen by the agent, its cost $c(S)$ is additive and its reward $f(S)$ is the max weight matching in the subgraph induced by S , with respect to weights $\{f_{v,u}\}_{(v,u) \in A}$. We call such a function f *matching-based reward function*.

It is easy to see that every matching based reward function is XOS, a superset of submodular functions. Example 5.1 shows that some matching-based functions are strictly XOS.

EXAMPLE 5.1. Consider the bipartite graph with vertices $V = \{v_1, v_2\}$ and $U = \{u_1, u_2\}$ and edges $A = \{(v_1, u_1), (v_1, u_2), (v_2, u_1)\}$, with rewards (see Figure 2):

$$f(v_1, u_1) = f(v_1, u_2) = 2 \quad f(v_2, u_1) = 1$$

Observe that $f(v_2, u_1) = 1$ and $f((v_2, u_1) \mid \{(v_1, u_1)\}) = 0$, which suggests f is not supermodular. However, since the edges $\{(v_1, u_2), (v_2, u_1)\}$ make up a more rewarding matching than $\{(v_1, u_1)\}$, we have that $f((v_2, u_1) \mid \{(v_1, u_1), (v_1, u_2)\}) = 1$, so f is not submodular either.




Figure 2: Matching-based instance with f not submodular nor supermodular

In this setting, the agent’s utility for a set of edges S and a contract α is given by

$$u_a(\alpha, S) = \max_{\substack{S' \subseteq S, \\ S' \text{ is a matching}}} \sum_{(v,u) \in S'} \alpha f_{v,u} - \sum_{(v,u) \in S} c_{v,u}$$

and its best response for a contract α is without loss of generality a set $S_\alpha \subseteq A$ which constitutes a matching. Thus, the agent’s demand problem corresponds to finding a maximum weight matching where the weight of an edge (v, u) is $w_\alpha(v, u) = \alpha f_{v,u} - c_{v,u}$ (note that the weights may be negative). As such, the demand can be computed efficiently using any algorithm for maximum weighted matching, so the first condition mentioned in Proposition 3.1 is met.

In Section 5.1 we show that the second condition of Proposition 3.1 is not always fulfilled, and present family of matching-based instances for which the number of critical values is super-polynomial. In Section 5.2, we give two special cases for matching-based rewards and costs in which there are polynomially-many critical values and the optimal contract is tractable.

5.1 A Super-Polynomial Lower Bound on the Number of Critical Values We derive a super-polynomial lower bound on the number of critical values from the construction made by Mulmuley and Shah [34], based on the work of Carstensen [8], for the parameterized shortest $s - t$ path problem. In this problem,

we are given a directed graph with two special nodes, $D = (X \cup \{s, t\}, \tilde{E})$, and a set of affine weight functions, $\{\tilde{w}_\alpha(e) = \tilde{c}_e - \alpha \tilde{f}_e\}_{e \in E}$, parameterized by $\alpha \in \mathbb{R}_{\geq 0}$. The weights are such that for every $e \in E$, $\tilde{c}_e \geq 0$. The goal is to find all the shortest paths between s and t for a given range of values of α .

THEOREM 5.1. (THEOREM 1.3, [34]) *There is a family of graphs $D_n = (X_n \cup \{s, t\}, \tilde{E}_n)$ with $O(n^4)$ vertices and affine weight functions $\{\tilde{w}_\alpha(e) = \tilde{c}_e - \alpha \tilde{f}_e \mid e \in \tilde{E}_n\}$, together with an upper bound $Q_n \in \mathbb{R}_{\geq 0}$, such that*

1. For every $e \in \tilde{E}_n$, $\tilde{c}_e \geq 0$
2. For every $\alpha \in [0, Q_n]$ and for every $e \in \tilde{E}_n$, $\tilde{w}_e = \tilde{c}_e - \alpha \tilde{f}_e \geq 0$
3. There are $m = 2^{\Omega(\log^2 n)}$ values, $0 \leq \alpha_1 < \dots < \alpha_m < Q_n$ such that for any two $\alpha_j \neq \alpha_i$, the shortest $s - t$ path with respect to weights \tilde{w}_{α_i} is different from the shortest $s - t$ path with respect to \tilde{w}_{α_j} .

In order to apply the above lower bound to the number of critical values, we present a reduction from the parametric shortest $s - t$ path to the agent’s demand problem for matching-based rewards and additive costs. We begin by reducing an instance of the shortest $s - t$ path problem for a fixed value of α to the min-weight perfect matching problem in a bipartite graph. We will show that the shortest $s - t$ path has a corresponding min-weight perfect matching of the same weight. This reduction will preserve the parametric structure of the weights, so we will get m different perfect weight matchings for $\alpha_1, \dots, \alpha_m$.

This will be followed by another reduction from the min-weight perfect matching to the agent’s demand problem. The affine structure of the weights will remain, but we will manipulate the coefficients in a way that will make this max-weight matching problem a valid instance of the agent’s demand problem.

Shortest $s - t$ Path to Minimum Weight Perfect Matching. Given an instance of the $s - t$ path problem (D, \tilde{w}_α) , where $D = (X \cup \{s, t\}, \tilde{E})$ is a directed graph with non-negative weights $\{\tilde{w}_e = \tilde{c}_e - \alpha \tilde{f}_e\}_{e \in \tilde{E}}$ on its edges, we construct an instance of the min-weight perfect matching problem, (G', w'_α) . In this problem, we need to find the min-weight perfect matching in the undirected bipartite graph $G' = (V' \cup U', E')$ with respect to weights defined over the edges $w'_\alpha : E' \rightarrow \mathbb{R}$. The instance we construct is defined as follows (see illustration in Figure 3).

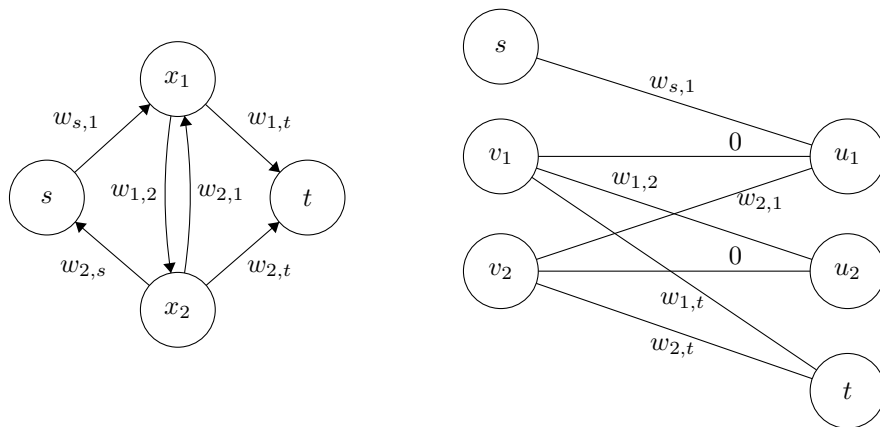


Figure 3: A reduction from a directed graph for the $s - t$ problem (left) to a bipartite graph for the minimum perfect weight matching (right).

Each vertex $x_i \in X$ is associated with two vertices in G' , $v_i \in V'$ and $u_i \in U'$ and the vertices s and t in D correspond to $s \in V'$ and $t \in U'$ respectively. The edges, E' , are defined as follows: For every directed edge $(x_i, x_j) \in \tilde{E}$, where $x_i, x_j \in X$, we have a corresponding $(v_i, u_j) \in E'$. For every edge that leaves s in the graph D , we have a corresponding one in E' . That is, $(s, x_i) \in \tilde{E}$ implies $(s, u_i) \in E'$ and $(s, t) \in \tilde{E}$ implies $(s, t) \in E'$. We also define corresponding edges to those that enter t , namely if $(x_i, t) \in \tilde{E}$, then $(v_i, t) \in E'$. We ignore any edge that enters s or leaves t , as those are never part of the optimal $s - t$ path. We also add an edge $(v_i, u_i) \in E'$

for every $i \in \{1, \dots, |X|\}$. For each $(v, u) \in E'$ we define the weights w' as follows:

$$w'_\alpha(v, u) = \begin{cases} \tilde{w}_\alpha(x_i, x_j) = \tilde{c}_{i,j} - \alpha \tilde{f}_{i,j} & v = v_i, u = u_j, j \neq i \\ \tilde{w}_\alpha(s, x_i) = \tilde{c}_{s,i} - \alpha \tilde{f}_{s,i} & v = s, u = x_i \\ \tilde{w}_\alpha(x_i, t) = \tilde{c}_{i,t} - \alpha \tilde{f}_{i,t} & v = v_i, u = t \\ 0 = 0 - \alpha 0 & v = v_i, u = u_i \end{cases}$$

We claim that this construction is such that for any α , the shortest $s - t$ path in (D, \tilde{w}_α) has a corresponding min-weight perfect matching in (G, w'_α) with the same weight.

CLAIM 5.1. *For any $\alpha \in [0, Q_n]$, the shortest $s - t$ path in (D, \tilde{w}_α) , P_α has a corresponding min-weight perfect matching in (G', w'_α) , M_α such that both have the same weight.*

Proof. Fix α and assume without loss of generality that $P_\alpha = (s, x_1, \dots, x_k, t)$ is the shortest $s - t$ path in D , then by definition of G' there exists a perfect matching of the same weight, $M = \{(s, u_1), (v_1, u_2), \dots, (v_{k-1}, u_k), (v_k, t)\} \cup \{(v_j, u_j) \mid j \notin [k]\}$.

Let M'_α be the min-weight perfect matching in (G', w'_α) and assume that its weight is strictly lesser than the weight of M above. We will construct an $s - t$ path, P , in (D, \tilde{w}_α) , which is shorter than P_α and reach contradiction. Let us rename the vertices such that $(s, u'_1), (v'_1, u'_2), \dots, (v'_{l-1}, u'_l), (v'_l, t) \in M'_\alpha$. We can assume the rest of the edges in M'_α are of the form (v'_j, u'_j) , as those are edges with minimal weights. Observe that the path $P = (s, x'_1, \dots, x'_l, t)$ is an $s - t$ path in D with the same weight as M_α . Thus, by assumption

$$\tilde{w}_\alpha(P) = w'_\alpha(M'_\alpha) < w'_\alpha(M) = \tilde{w}_\alpha(P_\alpha),$$

which completes the proof. \square

Claim 5.1 above together with Theorem 5.1, imply that there exists a family of bipartite graphs with affine weight functions such that the set of min-weight perfect matchings, for all values of the parameter $\alpha \in [0, Q_n]$, corresponds exactly to the set of all shortest $s - t$ paths.

COROLLARY 5.1. *There exists a family of bipartite graphs $G'_n = (V'_n \cup U'_n, E'_n)$ with $O(n^4)$ vertices and affine weight functions $\{w'_\alpha(e) = c'_e - \alpha f'_e \mid e \in E'_n\}$, together with an upper bound $Q_n \in \mathbb{R}_{\geq 0}$, such that*

1. *For every $e \in E'_n$, $c'_e \geq 0$*
2. *For every $\alpha \in [0, Q_n]$ and for every $e \in E_n$, $w'_e = c'_e - \alpha f'_e \geq 0$*
3. *There are $m = 2^{\Omega(\log^2 n)}$ values, $0 \leq \alpha_1 < \dots < \alpha_m < Q_n$ such that for any two $\alpha_j \neq \alpha_i$, the min-weight perfect matching with respect to weights w'_{α_i} is different from the min-weight perfect matching with respect to w'_{α_j} .*

Minimum Weight Perfect Matching to the Agent's Demand Problem. Given a min-weight perfect matching instance (G', w') as above, we construct a bipartite graph $G = (V \cup U, E)$ with weights $w : E \rightarrow \mathbb{R}_{\geq 0}$ that correspond to the agent's demand problem. The vertices and edges of G are the same as G' , so $V = V'$, $U = U'$ and $E = E'$, we only modify the weights.

Observe that if we just negate the weights on edges, namely setting them to be $\alpha f'_e - c'_e$, then those will not constitute a valid instance of the agent's demand problem, as we still need to guarantee that every subset of edges S in the new graph has a valid success probability, namely $f(S) \in [0, 1]$. Moreover, we need increase the weights of all edges, so that they are non-negative for every $\alpha \in [0, Q_n]$. This will allow us to map every min-weight perfect matching in (G', w') to a max-weight matching in (G, w) , which will give us the super-polynomial lower bound on the number of different max-weight matchings.

We modify the weight of an edge $e \in E$ to be

$$w_\alpha(e) = \alpha f_e - c_e = \alpha \left(\frac{f'_e + K}{2K|E|} \right) - \frac{c'_e}{2K|E|},$$

where

$$K = \max_{e \in E} |f'_e| + \max_{e \in E} \frac{c'_e}{\alpha_2} > 0$$

(recall that $\alpha_2 > 0$). Applying the above weights to the family of graphs G'_n from Corollary 5.1, we obtain the super-polynomial lower bound on the number of critical values for matching-based rewards and additive costs.

THEOREM 5.2. *There is a family of bipartite graphs $G_n = (V_n \cup U_n, E_n)$, with $O(n^4)$ vertices and affine weight functions $\{w_\alpha(e) = \alpha f_e - c_e \mid e \in E_n\}$, together with a lower bound $\alpha_2 > 0$ and an upper bound $Q_n \in (\alpha_2, \infty)$, such that:*

1. For every $e \in E_n$, $f_e, c_e \geq 0$.
2. For every $S \subseteq E_n$, $f(S) \in [0, 1]$.
3. For every $\alpha \in [\alpha_2, Q_n]$ and for every $e \in E_n$, $w_\alpha(e) \geq 0$
4. There are $m = 2^{\Omega(\log^2 n)}$ critical values, $0 < \alpha_2 < \dots < \alpha_m < Q_n$ such that for any $i \in \{2, \dots, m\}$, each critical value α_i has a distinct max-weight matching.

Proof. One can easily verify that the definition of K guarantees that the first two conditions hold. For condition 3 (i.e., non-negativity of the edges), consider some $\alpha \geq \alpha_2$:

$$\begin{aligned} w_\alpha(e) &= \alpha \left(\frac{f'_e + K}{2K|E|} \right) - \frac{c'_e}{2K|E|} \\ &= \frac{1}{2K|E|} (\alpha(f'_e + K) - c'_e) \\ &\geq \frac{1}{2K|E|} (\alpha(f'_e + \max_{e \in E} |f'_e|) + \alpha \max_{e \in E} \frac{c'_e}{\alpha_2} - c'_e) \\ &\geq 0 \end{aligned}$$

It remains to show that the fourth condition holds. First note that for any contract $\alpha \in [\alpha_2, Q_n]$, since $w_\alpha(e) \geq 0$ for every $e \in E$, every max-weight perfect matching is a max-weight matching.

Moreover, for every $\alpha \in [\alpha_2, Q_n]$, a matching M_α is a min-weight perfect matching in (G'_n, w'_α) if and only if it is a max-weight perfect matching in (G_n, w_α) . This is because for any other perfect matching M

$$\begin{aligned} w'_\alpha(M_\alpha) \leq w'_\alpha(M) &\Leftrightarrow \frac{1}{2K|E|} (|V_n|(K\alpha) - w'_\alpha(M_\alpha)) \geq \frac{1}{2K|E|} (|V_n|(K\alpha) - w'_\alpha(M)) \\ &\Leftrightarrow \frac{1}{2K|E|} \sum_{e \in M_\alpha} (\alpha K + \alpha f'_e - c'_e) \geq \frac{1}{2K|E|} \sum_{e \in M} (\alpha K + \alpha f'_e - c'_e) \\ &\Leftrightarrow \sum_{e \in M_\alpha} \left(\alpha \left(\frac{f'_e + K}{2K|E|} \right) - \frac{c'_e}{2K|E|} \right) \geq \sum_{e \in M} \left(\alpha \left(\frac{f'_e + K}{2K|E|} \right) - \frac{c'_e}{2K|E|} \right) \\ &\Leftrightarrow \sum_{e \in M_\alpha} (\alpha f_e - c_e) \geq \sum_{e \in M} (\alpha f_e - c_e) \\ &\Leftrightarrow w_\alpha(M_\alpha) \geq w_\alpha(M) \end{aligned}$$

This one-to-one correspondence, together with Corollary 5.1, shows that that every α_i for $i \in \{2, \dots, m\}$ is a distinct critical value. \square

5.2 Polynomial Number of Critical Values for Special Cases In this section we show the tracability of the optimal contract for restricted cases of the matching based setting. Besides the two cases below, one can easily show that when the space of costs or rewards is of constant size, the number of critical values is small. We give the proof in Appendix A.

Pseudo-Polynomial Bound for Integer f and c . When the costs and rewards take integer values, a result by Carstensen [8] implies that the optimal contract can be found in time polynomial in the number of vertices n , and the maximal reward $\max_{u,v} f_{u,v}$ and cost $\max_{u,v} c_{u,v}$.

THEOREM 5.3. (CARSTENSEN [8]) *The number of breakpoints of an integer linear program with a parameterized objective function $\max_x \alpha f(x) - c(x)$, for $f, c \in \mathbb{Z}^n$, and a fixed set of constraints $Ax \leq b$, is polynomial in n , F and C , where n is the dimension of x , $F = \max_{u,v} f_{u,v}$ and $C = \max_{u,v} c_{u,v}$.*

COROLLARY 5.2. *For matching based rewards and additive costs, with $f, c \in \mathbb{Z}^{|A|}$, where $C = \max_{e \in E} c_i$ and $F = \max_{e \in E} f_i$ are both polynomial in n , the optimal contract can be computed efficiently.*

One-Sided Costs. If the costs of the edges depends only on the U side of the edge, namely $c_{v,u} = c_u$ for every $(v, u) \in A$ (the case for V side is symmetric), then the number of critical values can be bounded by reducing the agent’s utility problem to an instance in which costs are additive and rewards are gross-substitutes, utilizing the result of [13].

THEOREM 5.4. *For a matching-based instance in which the costs are one-sided, the number of critical values is $O(n^2)$, where n is the number of vertices.*

Recall that for a set of edges $S \subseteq A$, when costs are one-sided the agent’s utility under contract α is

$$u_a(\alpha, S) = \max_{\substack{S' \subseteq S, \\ S' \text{ is a matching}}} \sum_{(v,u) \in S'} \alpha f_{v,u} - \sum_{(v,u) \in S} c_u$$

We show that the agent’s utility problem can be reformulated in terms of the vertices of U , with additive costs and an OXS reward function (as defined in Section 2).

Let $c' : 2^U \rightarrow \mathbb{R}_{\geq 0}$ be an additive cost function. That is, for any set $T \subseteq U$, $c'(T) = \sum_{u \in T} c_u$. Let $f' : 2^U \rightarrow \mathbb{R}$ be OXS with respect to the graph G with weights $\{f_{v,u}\}_{(v,u) \in A}$.

CLAIM 5.2. *For every $\alpha \in [0, 1]$, if S_α is a matching that maximizes the agent’s utility among all subsets of A , then the set of vertices in U that are matched in S_α maximize $\alpha f'(T) - c'(T)$ among all subsets $T \subseteq U$. In the other direction, if $T_\alpha \subseteq U$ maximizes $\alpha f'(T) - c'(T)$ among all subsets $T \subseteq U$, then the max weight matching in the subgraph induced by T_α with respect to weights $\{f_{v,u}\}_{(v,u) \in A}$, maximizes the agent’s utility under contract α .*

Proof. For every matching $S \subseteq A$, denote by $U(S) \subseteq U$ the vertices from U which that are matched in S . For every set $T \subseteq U$, denote by $A(T)$ the max weight matching in the subgraph induced by T w.r.t weights $\{f_{v,u}\}_{(v,u) \in A}$. Observe that since the weights are non-negative and costs are one-sided:

$$\begin{aligned} \max_{S \subseteq A} \alpha f(S) - c(S) &= \alpha f(S_\alpha) - c(S_\alpha) \\ &= \alpha \sum_{(v,u) \in S_\alpha} f_{v,u} - \sum_{(v,u) \in S_\alpha} c_{v,u} \\ &= \alpha \sum_{(v,u) \in S_\alpha} f_{v,u} - \sum_{u \in U(S_\alpha)} c_u \\ &= \alpha \sum_{(v,u) \in A(U(S_\alpha))} f_{v,u} - \sum_{u \in U(S_\alpha)} c_u \\ &= \alpha f'(U(S_\alpha)) - c'(U(S_\alpha)) \\ &\leq \max_{T \subseteq U} \alpha f'(T) - c'(T), \end{aligned}$$

where the fourth equality holds because S_α is a matching with respect to the vertices $U(S_\alpha)$, and because the existence of a matching with respect to $U(S_\alpha)$ with a higher reward would contradict the fact that S_α is a best-response.

On the other hand, we can use similar arguments to conclude that

$$\begin{aligned}
\max_{T \subseteq U} \alpha f'(T) - c'(T) &= \alpha f'(T_\alpha) - c'(T_\alpha) \\
&= \alpha \sum_{(v,u) \in A(T_\alpha)} f_{v,u} - \sum_{u \in T_\alpha} c_u \\
&= \alpha \sum_{(v,u) \in A(T_\alpha)} f_{v,u} - \sum_{(v,u) \in A(T_\alpha)} c_u \\
&= \alpha f(A(T_\alpha)) - c(A(T_\alpha)) \\
&\leq \alpha f(S_\alpha) - c(S_\alpha),
\end{aligned}$$

where the third and fourth equality transition hold because $A(T_\alpha)$ is a matching. \square

Proof. [of Theorem 5.4] Claim 5.2 shows that under one-sided costs the agent's utility maximization problem can be formulated in terms of the vertices of U and the reward and cost functions $f' : 2^U \rightarrow \mathbb{R}_{\geq 0}$ and $c' : 2^U \rightarrow \mathbb{R}_{\geq 0}$ described above. As c' is additive and f' is an OXS function, a sub-class of gross-substitutes, it follows from [13] that the number of critical values is at most $O(n^2)$, where n refers to the number of vertices in U . \square

References

- [1] T. ALON, P. DÜTTING, Y. LI, AND I. TALGAM-COHEN, *Bayesian analysis of linear contracts*, in EC 2023, 2023, p. 66.
- [2] T. ALON, P. DÜTTING, AND I. TALGAM-COHEN, *Contracts with private cost per unit-of-effort*, in EC 2021, 2021, pp. 52–69.
- [3] M. BABAI OFF, M. FELDMAN, AND N. NISAN, *Combinatorial agency*, in EC 2006, 2006, pp. 18–28.
- [4] ———, *Mixed strategies in combinatorial agency*, in WINE 2006, 2006, pp. 353–364.
- [5] ———, *Free-riding and free-labor in combinatorial agency.*, in SAGT 2009, 2009, pp. 109–121.
- [6] A. BERTELSEN, *Substitutes Valuations and m -concavity*, PhD thesis, The Hebrew University, 2005.
- [7] G. CARROLL, *Robustness and linear contracts*, Am. Econ. Rev., 105 (2015), pp. 536–563.
- [8] P. J. CARSTENSEN, *Complexity of some parametric integer and network programming problems*, Math. Program., 26 (1983), pp. 64–75.
- [9] M. CASTIGLIONI, A. MARCHESI, AND N. GATTI, *Bayesian agency: Linear versus tractable contracts*, in EC 2021, 2021, pp. 285–286.
- [10] M. CASTIGLIONI, A. MARCHESI, AND N. GATTI, *Designing menus of contracts efficiently: the power of randomization*, in EC 2022, 2022, pp. 705–735.
- [11] M. CASTIGLIONI, A. MARCHESI, AND N. GATTI, *Multi-agent contract design: How to commission multiple agents with individual outcomes*, in EC 2023, 2023, pp. 412–448.
- [12] R. DEO-CAMPO VUONG, S. DUGHMI, N. PATEL, AND A. PRASAD, *On supermodular contracts and dense subgraphs*, CoRR, abs/2308.07473 (2023).
- [13] P. DÜTTING, T. EZRA, M. FELDMAN, AND T. KESSELHEIM, *Combinatorial contracts*, in FOCS 2021, 2022, pp. 815–826.
- [14] P. DÜTTING, T. EZRA, M. FELDMAN, AND T. KESSELHEIM, *Multi-agent contracts*, in STOC 2023, 2023, pp. 1311–1324.
- [15] P. DÜTTING, G. GURUGANESH, J. SCHNEIDER, AND J. WANG, *Optimal no-regret learning for one-side lipschitz functions*, in ICML 2023, 2023. Forthcoming.

- [16] P. DÜTTING, T. ROUGHGARDEN, AND I. TALGAM-COHEN, *Simple versus optimal contracts*, in EC 2019, 2019, pp. 369–387.
- [17] P. DÜTTING, T. ROUGHGARDEN, AND I. TALGAM-COHEN, *The complexity of contracts*, SIAM J. Comput., 50 (2021), pp. 211–254.
- [18] M. J. EISNER AND D. G. SEVERANCE, *Mathematical techniques for efficient record segmentation in large shared databases*, J. ACM, 23 (1976), pp. 619–635.
- [19] S. J. GROSSMAN AND O. D. HART, *An analysis of the principal-agent problem*, Springer, 1992.
- [20] G. GURUGANESH, J. SCHNEIDER, AND J. R. WANG, *Contracts under moral hazard and adverse selection*, in EC 2021, 2021, pp. 563–582.
- [21] G. GURUGANESH, J. SCHNEIDER, J. R. WANG, AND J. ZHAO, *The power of menus in contract design*, in EC 2023, 2023, pp. 818–848.
- [22] D. M. GUSFIELD, *Sensitivity analysis for combinatorial optimization*, University of California, Berkeley, 1980.
- [23] C.-J. HO, A. SLIVKINS, AND J. W. VAUGHAN, *Adaptive contract design for crowdsourcing markets: Bandit algorithms for repeated principal-agent problems*, in EC 2014, 2014, pp. 359–376.
- [24] B. HOLMSTRÖM, *Moral hazard and observability*, The Bell Journal of Economics, (1979), pp. 74–91.
- [25] B. HOLMSTROM AND P. MILGROM, *Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design*, J. Law Econ. Organ., 7 (1991), pp. 24–52.
- [26] S. IWATA, *A fully combinatorial algorithm for submodular function minimization*, J. Comb., Series B, 84 (2002), pp. 203–212.
- [27] ———, *A faster scaling algorithm for minimizing submodular functions*, SIAM J. Comput., 32 (2003), pp. 833–840.
- [28] S. IWATA, L. FLEISCHER, AND S. FUJISHIGE, *A combinatorial strongly polynomial algorithm for minimizing submodular functions*, J. ACM, 48 (2001), pp. 761–777.
- [29] R. M. KARP AND J. B. ORLIN, *Parametric shortest path algorithms with an application to cyclic staffing*, Discret. Appl. Math., 3 (1981), pp. 37–45.
- [30] J. KLEINBERG AND R. KLEINBERG, *Delegated search approximates efficient search*, in EC 2018, 2018, pp. 287–302.
- [31] J. KLEINBERG AND M. RAGHAVAN, *How do classifiers induce agents to invest effort strategically?*, Trans. Econ. Comput., 8 (2020), pp. 1–23.
- [32] B. LEHMANN, D. LEHMANN, AND N. NISAN, *Combinatorial auctions with decreasing marginal utilities*, in Proceedings of the 3rd ACM conference on Electronic Commerce, 2001, pp. 18–28.
- [33] N. MEGIDDO, *Combinatorial optimization with rational objective functions*, in STOC 1978, 1978, pp. 1–12.
- [34] K. MULMULEY AND P. SHAH, *A lower bound for the shortest path problem*, in CCC 2000, 2000, pp. 14–21.
- [35] K. NAGANO, *A strongly polynomial algorithm for line search in submodular polyhedra*, Discret. Optim., 4 (2007), pp. 349–359.
- [36] R. PAES LEME, *Gross substitutability: An algorithmic survey*, Games and Economic Behavior, 106 (2017), pp. 294–316.
- [37] R. P. STANLEY, *Enumerative combinatorics volume 1 second edition*, Cambridge studies in advanced mathematics, (2011).
- [38] B. ZHU, S. BATES, Z. YANG, Y. WANG, J. JIAO, AND M. I. JORDAN, *The sample complexity of online contract design*, in EC 2023, 2023, p. 1188.

A Matching-based Rewards and Additive Costs with a Constant-Size Reward Space

This section refers to a small number of different edge rewards, the case of edge costs is symmetric.

CLAIM A.1. *For matching-based rewards and additive costs, if the number of different values $f_{v,u}$ takes is $k = O(1)$, i.e. $|\{x \mid f_{v,u} = x, v \in V, u \in U\}| = k$, then the number of critical values is $O(n^{k+1})$, where n is the number of edges.*

Proof. Recall that by Observation 2.1, for any two critical values $\alpha < \alpha'$, the best responses satisfy $f(S_\alpha) < f(S_{\alpha'})$. So the number of critical values is upper bounded by the number of different values $f(S)$ may take, where $S \subseteq A$ is a matching in G . Since $f(S) = \sum_{(v,u) \in S} f_{v,u}$, for any $j \in \{1, \dots, n\}$, the number of different values $f(S)$ may take when $|S| = j$ is upper bounded by the number of distinct multisets of size j with elements taken from a set of cardinality k . This is exactly multiset coefficient, $\binom{j}{k} = \binom{j+k-1}{j}$ [37]. The total number of different values f may take is thus, $\sum_{j=1}^n \binom{j+k-1}{j}$, which one can easily bound with $O(n^{k+1})$. \square