Theoretical Computer Science 452 (2012) 56–74

Contents lists available at SciVerse ScienceDirect

Theoretical Computer Science

journal homepage: www.elsevier.com/locate/tcs

Computing optimal contracts in combinatorial agencies☆

Yuval Emek[a], Michal Feldman[b,*]

[a] Computer Engineering and Networks Laboratory (TIK), ETH Zurich, Zurich, Switzerland

[b] School of Business Administration and the Center for the Study of Rationality, The Hebrew University of Jerusalem, Jerusalem, Israel

ARTICLE INFO

Article history:
Received 30 May 2010
Received in revised form 3 May 2012
Accepted 12 May 2012
Communicated by A. Fiat

Keywords:
Game theory
Mechanism design
Hidden action

ABSTRACT

We study an economic setting in which a principal motivates a team of strategic agents to exert costly effort toward the success of a joint *project*. The action of each agent is *hidden* and affects the outcome of the agent's individual *task* in a stochastic manner. A Boolean *technology* function maps the outcomes of the individual tasks to the project's outcome. The principal induces a Nash equilibrium on the agents' actions through payments which are conditioned on the project's outcome and the main challenge is that of determining the Nash equilibrium that maximizes the principal's net utility, namely, the *optimal contract*.

Babaioff, Feldman and Nisan study a basic *combinatorial agency* model for this setting, and provide a full analysis of the AND technology. Here, we concentrate mainly on OR technologies that, surprisingly, turn out to be much more complex. We provide a complete analysis of the computational complexity of the optimal contract problem in OR technologies which resolves an open question and disproves a conjecture raised by Babaioff et al. While the AND case admits a polynomial time algorithm, we show that computing the optimal contract in an OR technology is NP-hard. On the positive side, we devise an FPTAS for OR technologies. We also study *series–parallel (SP)* technologies, which are constructed inductively from AND and OR technologies. We establish a scheme that given any SP technology, provides a $(1+\epsilon)$-approximation for all but an $\hat{\epsilon}$-fraction of the relevant instances in time polynomial in the size of the technology and in the reciprocals of ϵ and $\hat{\epsilon}$.

1. Introduction

We consider the setting in which a principal motivates a team of rational agents to exert costly effort toward the success of a joint project, where the actions of the agents are hidden from the principal. The outcome (usually, success or failure of the project) is stochastically determined by the set of actions taken by the agents and is visible to all. As agents' actions are invisible, their compensation depends on the outcome and the principal's challenge is to design contracts (conditional payments to the agents) as to maximize her net utility, given the payoff that she obtains from a successful outcome.

The problem of hidden actions in production teams has been extensively studied in the economics literature [11,14,19,12,20]. More recently, the problem has been examined from a computational perspective [9,1,4,2,13,6,3,5]. This line of research complements the field of Algorithmic Mechanism Design (AMD) [16,15,18,8,17] that received much attention in the past decade. While AMD studies the design of mechanisms in scenarios characterized by private information held by the individual agents, our focus is on the complementary problem, that of hidden-action taken by the individual agents.

☆ A preliminary version has appeared in the Workshop on Internet and Network Economics, 2009.

* Corresponding author. Tel.: +972 54 3020222.
E-mail addresses: yuval.emek@tik.ee.ethz.ch (Y. Emek), michal.feldman@huji.ac.il (M. Feldman).

0304-3975/$ – see front matter © 2012 Elsevier B.V. All rights reserved.
doi:10.1016/j.tcs.2012.05.018

The model. We use the model presented in [1] (which is an extension of the model devised in [21]). In this model, a principal employs a set[1] N of agents in a joint *project*. Each agent i takes an action $a_i \in \{0, 1\}$, which is known only to him, and succeeds or fails in his own *task* probabilistically and independently. The individual outcome of agent i is denoted by $x_i \in \{0, 1\}$. If the agent shirks ($a_i = 0$), he succeeds in his individual task ($x_i = 1$) with probability $0 < \gamma_i < 1$ and incurs no cost. If, however, he decides to exert effort ($a_i = 1$), he succeeds with probability $0 < \delta_i < 1$, where $\delta_i > \gamma_i$, but incurs some positive real *cost* $c_i > 0$.

A key component of the model is the way in which the individual outcomes determine the outcome of the whole project. We assume a monotone Boolean function $\varphi : \{0, 1\}^n \to \{0, 1\}$ that determines whether the project succeeds as a function of the individual outcomes of the n agents' tasks (and is not determined by any set of $n-1$ agents). Two fundamental examples of such Boolean functions are AND and OR. The AND function is the logical conjunction of x_i ($\varphi(x_1, \ldots, x_n) = \bigwedge_{i\in N} x_i$), representing the case in which the project succeeds only if *all* agents succeed in their tasks. In this case, we say that the agents *complement* each other. The OR function represents the other extreme, in which the project succeeds if *at least one* of the agents succeeds in his task. This function is the logical disjunction of x_i ($\varphi(x_1, \ldots, x_n) = \bigvee_{i\in N} x_i$), and we say that the agents *substitute* each other.

Given the action profile $a = (a_1, \ldots, a_n) \in \{0, 1\}^n$ and a monotone Boolean function $\varphi : \{0, 1\}^n \to \{0, 1\}$, the *effectiveness* of the action profile a, denoted by $f(a)$, is the probability that the whole project succeeds under a and φ according to the distribution specified above. That is, the effectiveness $f(a)$ is defined as the probability that $\varphi(x_1, \ldots, x_n) = 1$, where $x_i \in \{0, 1\}$ is determined probabilistically (and independently) by a_i: if $a_i = 0$, then $x_i = 1$ with probability γ_i; if $a_i = 1$, then $x_i = 1$ with probability δ_i. The monotonicity of φ and the assumption that $\delta_i > \gamma_i$ for every $i \in N$ imply the monotonicity of the effectiveness function f, i.e., if we denote by $a_{-i} \in \{0, 1\}^{n-1}$ the vector of actions taken by all agents excluding agent i (namely, $a_{-i} = (a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n)$), then the effectiveness function must satisfy $f(1, a_{-i}) > f(0, a_{-i})$ for every $i \in N$ and $a_{-i} \in \{0, 1\}^{n-1}$. Note that it is inherent to our model (in fact, the model of [1]) that the effectiveness $f(a)$ consists of a "probabilistic component" that determines the individual outcomes $x_1, \ldots, x_n$ and a "deterministic component" that maps these individual outcomes to success or failure of the whole project.

The agents' success probabilities, the costs of exerting effort, and the monotone Boolean function that determines the final outcome define the *technology*, formally defined as the five-tuple $t = \langle N, \{\gamma_i\}_{i=1}^n, \{\delta_i\}_{i=1}^n, \{c_i\}_{i=1}^n, \varphi\rangle$, where N is a (finite) set of agents, γ_i (respectively, δ_i) is the probability that $x_i = 1$ when agent i shirks (resp., when agent i exerts effort), where $\delta_i > \gamma_i$, c_i is the cost incurred on agent i for exerting effort, and $\varphi : \{0, 1\}^n \to \{0, 1\}$ is the monotone Boolean function that maps the individual outcomes $x_1, \ldots, x_n$ to the outcome of the whole project. We sometimes abuse notation and refer to the Boolean function φ as the technology. It is important to emphasize that the technology is assumed to be known by the principal and the agents.

Since exerting effort entails some positive cost, an agent will not exert effort unless induced to do so by appropriately designed incentives. The principal can motivate the agents by offering them individual *payments*. However, due to the non-visibility of the agents' actions, the individual payments cannot be directly contingent on the actions of the agents, but rather only on the success of the whole project. The *conditional payment* to agent i is thus given by a real value $p_i \geq 0$ that is granted to agent i by the principal if the project succeeds (otherwise, the agent receives 0 payment).[2]

The expected *utility* of agent i under the profile of actions $a = (a_1, \ldots, a_n)$ and the conditional payment p_i is $p_i \cdot f(a)$ if $a_i = 0$, and $p_i \cdot f(a) - c_i$ if $a_i = 1$. Given a real *payoff* $v > 0$ that the principal obtains from a successful outcome of the project, the principal wishes to design the payments p_i as to maximize her own expected *utility* defined as $U_a(v) = f(a) \cdot \left(v - \sum_{i\in N} p_i\right)$, where the action profile a is assumed to be at Nash-equilibrium with respect to the payments p_i (i.e., no agent can improve his utility by a unilateral deviation). As multiple Nash equilibria may (and actually do) exist, we focus on the one that maximizes the utility of the principal. This is as if we let the principal choose the desired Nash equilibrium, and "suggest" it to the agents. The following observation is established in [1].

Observation. *The best conditional payments (from the principal's point of view) that induce the action profile $a \in \{0, 1\}^n$ as a Nash equilibrium are $p_i = 0$ for agent i who shirks ($a_i = 0$), and $p_i = \frac{c_i}{\Delta_i(a_{-i})}$ for agent i who exerts effort ($a_i = 1$), where $\Delta_i(a_{-i}) = f(1, a_{-i}) - f(0, a_{-i})$. (Note that the monotonicity of the effectiveness function guarantees that $\Delta_i(a_{-i})$ is always positive.)*

The last observation implies that once the principal chooses the action profile $a \in \{0, 1\}^n$, her (maximum) expected utility is determined to be $U_a(v) = f(a) \cdot (v - p(a))$, where $p(a)$ is the total *payment* (in the case of a successful outcome of the project), given by $p(a) = \sum_{i|a_i=1} \frac{c_i}{\Delta_i(a_{-i})}$. Therefore the principal's goal is merely to choose a subset $S \subseteq N$ of agents that exert effort (the rest of the agents shirk) so that her expected utility is maximized. The agent subset S is referred to as a *contract* and we say that the principal *contracts with agent i* if $i \in S$. We sometimes abuse notation and denote $f(S)$, $p(S)$ and $U_S(v)$ instead of $f(a)$, $p(a)$ and $U_a(v)$, respectively, where $a_i = 1$ if $i \in S$ and $a_i = 0$ if $i \notin S$. Given the principal's payoff $v > 0$, a contract $T \subseteq N$ is said to be *optimal* if $U_T(v) \geq U_S(v)$ for every contract $S \subseteq N$.

While finding the optimal set of payments that induces a particular set of agents to exert effort is a straightforward task (and can be efficiently computed), finding an optimal contract for a given payoff $v > 0$ is the main challenge addressed in this

1 Unless stated otherwise, we assume that $N = [n]$, where $[n]$ denotes the set $\{1, \ldots, n\}$.

2 We impose the *limited liability* constraint, implying that the principal can pay the agents but not fine them. Thus, all payments must be non-negative.

paper. Given a technology $t = \langle N, \{\gamma_i\}_{i=1}^n, \{\delta_i\}_{i=1}^n, \{c_i\}_{i=1}^n, \varphi\rangle$, we refer to the collection of contracts that can be obtained as an optimal contract for some payoff as the *orbit* of t (ties between different contracts are broken according to a lexicographic order[3]). Once the contract $S \subseteq N$ is chosen, the expected utility of the principal $U_S(v) = f(S)(v - p(S))$ becomes a linear function of the payoff v. Therefore each contract S corresponds to some line in $\mathbb{R}^2$. It follows that computing the orbit of t is equivalent to identifying the (positive) top envelope of the line collection $\{U_S(\cdot) \mid S \subseteq N\}$ in $\mathbb{R}^2$.

It is easy to see that for sufficiently low payoffs, no agent will ever be contracted while for sufficiently high payoffs, all agents will always be contracted. Therefore the *trivial contracts* $\emptyset$ and N are always in the orbit. Let $v^* = \inf\{v > 0 \mid N \text{ is optimal for } v\}$. Clearly, the trivial contract N is optimal for every $v > v^*$ and the infinite interval (v^*, ∞) does not exhibit any transitions in the orbit. We refer to the payoffs in the interval $(0, v^*]$ as the *relevant* payoffs.

Our results. Multi-agent projects may exhibit delicate combinatorial structures of dependencies between the agents' actions, which can be represented by a wide range of monotone Boolean functions. In the two extremes of this range reside two simple and natural functions, namely AND and OR, which correspond to the respective cases of pure complementarities and pure substitutabilities. These are arguably the two most fundamental interrelations between agents' actions.

The AND case was fully analyzed in [1]. In particular, it was (implicitly) shown that the optimal contract of any AND technology can be computed in polynomial time. In contrast, the OR case was left unresolved to the most part. Specifically, one of the main open questions raised in [1] asked whether the optimal contract problem on OR technologies can be solved in polynomial time.

We provide a complete analysis of the computational complexity of the optimal contract problem on OR technologies. Our first theorem, established in Section 2, addresses the hardness of this problem. Note that aside from establishing the computational hardness of the problem, our analysis implies the existence of OR technologies which admit exponential-size orbits, thus refuting a conjecture raised in [1].

Theorem 1. *The problem of computing the optimal contract in OR technologies is NP-hard.*[4] *The problem remains NP-hard even for the special case in which $c_i = 1$ and $\delta_i = 1 - \gamma_i$ for every $i \in N$.*

This hardness result is complemented by an approximation scheme in Section 3.2.

Theorem 2. *The problem of computing the optimal contract in OR technologies admits a fully polynomial-time approximation scheme (FPTAS).*

In fact, the FPTAS devised in Theorem 2 is based on a more general scheme, developed in Section 3.1. This scheme can be applied to a generalization of AND and OR technologies, referred to as *series–parallel* (SP) technologies that we now turn to define.

SP functions are monotone Boolean functions defined inductively as follows. The uni-argument identity function is considered SP. Consider two SP functions $\varphi_l : \{0,1\}^{n_l} \to \{0,1\}$ and $\varphi_r : \{0,1\}^{n_r} \to \{0,1\}$. The Boolean functions $\varphi_l \wedge \varphi_r : \{0,1\}^{n_l+n_r} \to \{0,1\}$, defined as the logical conjunction of φ_l and φ_r, and $\varphi_l \vee \varphi_r : \{0,1\}^{n_l+n_r} \to \{0,1\}$, defined as the logical disjunction of φ_l and φ_r, are also considered SP. We refer to the former (respectively, the latter) as a *series composition* (resp., a *parallel composition*) of φ_l and φ_r, hence the name series–parallel. Since series and parallel compositions are associative, it follows that the class of SP Boolean functions is indeed a generalization of both AND and OR Boolean functions. SP Boolean functions are of great interest to computer science. For instance, they play an important role in combinatorial games due to their equivalence to game trees (AND-OR trees).

General SP technologies are considerably more involved and the approximability of the optimal contract problem on such technologies remains an open question. However, an interesting insight into this question is provided by a scheme that approximates all but a small fraction of the relevant payoffs. The following theorem is established in Section 3.3.

Theorem 3. *Given an SP technology t and two real parameters $0 < \epsilon, \hat{\epsilon} \le 1$, there exists a scheme that on input payoff $v > 0$, either returns a $(1+\epsilon)$-approximate solution for v or outputs a failure message, in time $\text{poly}(|t|, 1/\epsilon, 1/\hat{\epsilon})$. Assuming that $F \subseteq \mathbb{R}_{>0}$ is the set of reals on which the scheme outputs a failure message, it is guaranteed that*[5] $\int_0^\infty 1_F(v)dv \le \hat{\epsilon} v^*$*, where 1_F is the characteristic function of F.*

It may be the case that the hardness of the optimal contract problem on SP technologies is somehow "concentrated" exactly in those payoffs which cannot be reached by the scheme of Theorem 3. However, if an instance of the problem is chosen uniformly at random out of the "relevant instances", then with high probability our scheme provides a good approximation for this instance. (Recall that the trivial contract N is optimal for any non-relevant payoff.) In fact, the payoffs v on which the scheme of Theorem 3 outputs a failure message belong to a small (polynomial) number of sub-intervals of $(0, v^*]$; by making the parameter $\hat{\epsilon}$ smaller, we decrease the guaranteed bound on the size of each such sub-interval.

It is interesting to contrast the aforementioned results with the *observable-action* case, where the agents' actions are not hidden and may be contracted on, which admits a polynomial time algorithm for SP technologies [7].

3 This implies that there are no two contracts with the same effectiveness in the orbit.

4 The hardness is established via a polynomial time Turing reduction.

5 Recall that $v^* = \inf\{v > 0 \mid N \text{ is optimal for } v\}$.

Finally, we obtain a positive result regarding the general case. Consider an arbitrary technology t and let $\mathcal{S}$ be a collection of contracts. Given some real $\alpha > 1$, we say that $\mathcal{S}$ is an *α-approximation* of t's orbit if for every payoff v, there exists a contract $S \in \mathcal{S}$ such that $U_S(v) \geq \frac{U_T(v)}{\alpha}$, where T is optimal for v. The following theorem, which guarantees the *existence* of a polynomial size collection approximating t, is established in Section 3.4.

Theorem 4. *For every technology $t = \langle N, \{\gamma_i\}_{i=1}^n, \{\delta_i\}_{i=1}^n, \{c_i\}_{i=1}^n, \varphi\rangle$ and for any $\epsilon > 0$, the orbit of t admits a $(1+\epsilon)$-approximation of size* poly$(|t|, 1/\epsilon)$.

Unfortunately, in the case of arbitrary technologies (as opposed to OR technologies) we do not know how to construct the approximating collection efficiently.

2. NP-hardness of OR technologies

In this section, we establish the NP-hardness of the optimal contract problem on OR technologies.

Theorem 1. *The problem of computing the optimal contract in OR technologies is NP-hard.*

We start with a high-level description.

Overview. We present a polynomial time Turing reduction from X3SAT (Problem LO4 in [10]) to the problem of computing an optimal contract for an OR technology. A 3-CNF formula ϕ is solvable under X3SAT if there exists a truth assignment for the variables of ϕ that assigns true to exactly one literal in every clause. The X3SAT problem is known to be NP-hard even if the literals in ϕ are all positive. Given a 3-CNF formula ϕ with m clauses and n variables in which all literals are positive, we construct an OR technology $t = \langle N, \{\gamma_j\}_{j=1}^{n+5}, \{\delta_j\}_{j=1}^{n+5}, \{c_j\}_{j=1}^{n+5}, \varphi\rangle$ such that (1) the agent set N contains $n+5$ agents, (2) the cost incurred on agent j for exerting effort is $c_j = 1$ for every $j \in N$, and (3) $\gamma_j = 1 - \delta_j$ for every $j \in N$. The construction is designed to guarantee that by performing $O(n)$ queries, each revealing the optimal contract for some carefully chosen payoff, we can decide whether ϕ is solvable under X3SAT.

Let $\mathcal{W} = \{0, 1, 2, 3\}^{m+2} \times \{0, 1\}^2$. Each agent $j \in N$ is assigned with a *vector* $\mathbf{u}^j = (u_0^j, \ldots, u_{m+3}^j) \in \mathcal{W}$. The first n agents correspond to the n variables of the 3-CNF formula ϕ and affect coordinates $1, \ldots, m$ in a manner that reflects the appearance of their corresponding variables in the m clauses. The additional 5 agents affect coordinates $0, m+1, m+2, m+3$ and are provided for the sake of analysis. We extend the assignment of vectors to sets of agents (a.k.a. contracts) in a natural way: given a contract $S \subseteq N$, we define the vector $\mathbf{u}^S = \sum_{j \in S} \mathbf{u}^j$. (Note that different contracts may be assigned with the same vector.) The assignment of vectors to contracts guarantees that the formula ϕ is solvable under X3SAT if and only if there exists a contract S with vector $\mathbf{u}^S = (1, \ldots, 1)$.

The parameters of $\{\gamma_j\}_{j=1}^{n+5}$ and $\{\delta_j\}_{j=1}^{n+5}$ are defined as follows. Consider the vector $\mathbf{x} = (x_0, \ldots, x_{m+3})$ in $\mathcal{W}$. Let $\sigma(\mathbf{x}) = \sum_{i=0}^{m+1} x_i 4^i$ and fix $\mu = 4^{5(m+2)}$. The *evaluation* of $\mathbf{x}$ is defined to be $\tau(\mathbf{x}) = \left(1 + \frac{1}{\mu}\right)^{\sigma(\mathbf{x})} \cdot \mu^{2x_{m+2}} \cdot \mu^{5x_{m+3}}$. Let $\epsilon = \mu^{-\kappa}$, where κ is a sufficiently large constant.

We would have wanted to fix $\gamma_j = 1 - \delta_j = \tau(\mathbf{u}^j) \cdot \epsilon$ for every $j \in N$. Unfortunately, the standard binary representation of $\tau(\mathbf{u}^j)$ may be much larger than the binary representation of ϕ for some j, and in particular, exponential in m. To overcome this obstacle, we use a carefully chosen estimation of $\tau(\mathbf{u}^j)$, so that on the one hand, the desired properties of the evaluation function are preserved, and on the other hand, the binary representation of γ_j (and δ_j) is polynomial in m. In particular, the choice of $\{\gamma_j\}_{j=1}^{n+5}$ and $\{\delta_j\}_{j=1}^{n+5}$ guarantees that for every two contracts $S, T \subseteq N, f(S) > f(T)$ if and only if $|S| > |T|$ or $|S| = |T|$ and $\mathbf{u}^S$ is lexicographically smaller than $\mathbf{u}^T$.

We argue that if some contracts S with $\mathbf{u}^{\mathbf{S}} = (1, \ldots, 1)$ exist, then at least one of them is in the orbit. This is done as follows. A vector $\mathbf{x} = (x_0, \ldots, x_{m+3})$ is said to be *protected* if $x_{m+2} = x_{m+3} = 1$. The key lemma of our proof asserts that any contract assigned with a protected vector $\mathbf{x}$ cannot be dominated by any two contracts assigned with different vectors. Following some standard geometric arguments, we conclude that the contracts assigned with $\mathbf{x}$ cannot be dominated by any set of (other) contracts. More formally, for every $0 \leq k \leq n+5$, we denote $\Psi_k(\mathbf{x}) = \{S \subseteq N \mid \mathbf{u}^S = \mathbf{x} \text{ and } |S| = k\}$, and show that for any protected vector $\mathbf{x}$, if $\Psi_k(\mathbf{x})$ is not empty, then at least one contract in $\Psi_k(\mathbf{x})$ is in the orbit. In particular, assuming that $\mathbf{x} = (1, \ldots, 1)$, if $\Psi_k(\mathbf{x}) \neq \emptyset$, then there exist a contract $S \in \Psi_k(\mathbf{x})$ and a payoff v_k^* such that S is optimal for v_k^*.

Computing the payoff v_k^* for every $1 \leq k \leq n+5$ remains our ultimate challenge. To achieve this goal, we define two additional vectors $\mathbf{w} = (2, 1, 1, \ldots, 1) \in \mathcal{W}$ and $\mathbf{y} = (0, 1, 1, \ldots, 1) \in \mathcal{W}$. The choice of the additional vectors guarantees that if $\Psi_k(\mathbf{x})$ is not empty, then neither are $\Psi_k(\mathbf{y})$ and $\Psi_k(\mathbf{w})$. Suppose that $\Psi_k(\mathbf{x}) \neq \emptyset$ and fix $\lambda_k^{w,x} = \max\{v[S, T] \mid S \in \Psi_k(\mathbf{w}) \text{ and } T \in \Psi_k(\mathbf{x})\}$ and $\lambda_k^{x,y} = \min\{v[S, T] \mid S \in \Psi_k(\mathbf{x}) \text{ and } T \in \Psi_k(\mathbf{y})\}$, where $v[S, T]$ is the intersection payoff of S and T, i.e., $U_S(v[S, T]) = U_T(v[S, T])$. We show that the optimal contract for every $\lambda_k^{w,x} < v < \lambda_k^{x,y}$ must be in $\Psi_k(\mathbf{x})$. The analysis is completed by identifying some payoff $\lambda_k^{w,x} < v_k^* < \lambda_k^{x,y}$ such that the binary representation of v_k^* is polynomial in m.

The decision whether the formula ϕ is solvable under X3SAT is now carried out as follows. For $k = 1, \ldots, n+5$, we query on the optimal contract S_k for the payoff v_k^*. If $\mathbf{u}^{S_k}$ is of the form $(1, \ldots, 1)$ for some k, then ϕ must be solvable. Otherwise, there does not exist any such contract and ϕ is not solvable.

i	0	1	2	…	m	$m+1$	$m+2$	$m+3$
$u^1,\dots,u^n$	0	$u^j_i=1$ if $i\in\Gamma_j$ and $u^j_i=0$ otherwise				0	0	0
u^α	0	0				0	1	0
u^β	0	0				0	0	1
u^A	0	0				1	0	0
u^B	1	0				1	0	0
u^C	2	0				1	0	0

Fig. 1. The $n+5$ vectors representing the $n+5$ agents of the technology. The first n agents correspond to the n variables of the 3-CNF formula ϕ, and the additional 5 agents are assigned with the vectors u^α, u^β, u^A, u^B and u^C.

The reduction. We now turn to describe the reduction in detail. Let $\mathcal{W} = \{0,1,2,3\}^{m+2} \times \{0,1\}^2$. Each agent $j \in N$ is assigned with a *vector* $\mathbf{u}^j = (u^j_0, \dots, u^j_{m+3}) \in \mathcal{W}$. The first n agents correspond to the n variables of the 3-CNF formula ϕ. Assuming that variable j appears in clauses $\Gamma_j \subseteq \{1,\dots,m\}$ (always as a positive literal), the vector $\mathbf{u}^j$ is defined so that $u^j_i = 1$ if $i \in \Gamma_j$ and $u^j_i = 0$ if $i \notin \Gamma_j$ for every $0 \le i \le m+3$ (thus $u^j_0 = u^j_{m+1} = u^j_{m+2} = u^j_{m+3} = 0$).

Agents $n+j$ for $j = 1,\dots,5$ are provided for the sake of the analysis. To avoid cumbersome indexing, we denote $n+1$ and $n+2$ by α and β, respectively, and $n+3$, $n+4$ and $n+5$ by A, B and C, respectively. Agents α and β are assigned with the vectors $\mathbf{u}^\alpha = (0,\dots,0,1,0) \in \mathcal{W}$ and $\mathbf{u}^\beta = (0,\dots,0,0,1) \in \mathcal{W}$, respectively. Agents A, B and C are assigned with the vectors $\mathbf{u}^A = (0,\dots,0,1,0,0) \in \mathcal{W}$, $\mathbf{u}^B = (1,0,0,\dots,0,1,0,0) \in \mathcal{W}$ and $\mathbf{u}^C = (2,0,0,\dots,0,1,0,0) \in \mathcal{W}$, respectively (see Fig. 1). Observe that the first n agents affect coordinates $1,\dots,m$, agents α and β affect coordinates $m+2$ and $m+3$, and agents A, B and C affect coordinates 0 and $m+1$.

We extend the assignment of vectors to sets of agents (a.k.a. contracts) in a natural way. Given a contract $S \subseteq N$, we define the vector $\mathbf{u}^S = \sum_{j\in S} \mathbf{u}^j$. As each clause in ϕ contains (at most) three variables, and by the definition of the vectors $\mathbf{u}^\alpha$, $\mathbf{u}^\beta$, $\mathbf{u}^A$, $\mathbf{u}^B$ and $\mathbf{u}^C$, it follows that $\mathbf{u}^S \in \mathcal{W}$ for every contract $S \subseteq N$. Observe that different contracts may be assigned with the same vector in $\mathcal{W}$.

The reduction relies on the following fact: the formula ϕ is solvable under X3SAT if and only if there exists a contract S with vector $\mathbf{u}^S = (1,\dots,1) \in \mathcal{W}$. To justify this fact, note that there exists a truth assignment that assigns true to exactly one variable in every clause of ϕ if and only if there exists a contract $S \subseteq [n]$ such that $\mathbf{u}^S = (u^S_0, 1, 1, \dots, 1, u^S_{m+1}, u^S_{m+2}, u^S_{m+3})$, where $u^S_0 = u^S_{m+1} = u^S_{m+2} = u^S_{m+3} = 0$. Agents α, β and B can be added to S, thus setting $u^S_0 = u^S_{m+1} = u^S_{m+2} = u^S_{m+3} = 1$, without affecting any other coordinate. We will show that if such a contract exists, then it is optimal for some payoff v^* which will be determined later on.

Vector evaluations. We now turn to define the parameters γ_i and δ_i of the agents. For that purpose, we first have to define a couple of functions that map the vectors in $\mathcal{W}$ to the reals. Consider the vector $\mathbf{x} = (x_0,\dots,x_{m+3})$ in $\mathcal{W}$. Let

$$\sigma(\mathbf{x}) = \sum_{i=0}^{m+1} x_i 4^i$$

and fix $\mu = 4^{5(m+2)}$. The *partial evaluation* of $\mathbf{x}$ is defined to be

$$\tau_p(\mathbf{x}) = \left(1 + \frac{1}{\mu}\right)^{\sigma(\mathbf{x})}$$

and the *full evaluation* of $\mathbf{x}$ is defined to be

$$\tau(\mathbf{x}) = \tau_p(\mathbf{x}) \cdot \mu^{2x_{m+2}} \cdot \mu^{5x_{m+3}}.$$

Observe that $\tau(\mathbf{x}) = \tau_p(\mathbf{x})$ if $x_{m+2} = x_{m+3} = 0$.

Let $\epsilon = \mu^{-\kappa}$, where κ is a sufficiently large constant (independent of m and n) that will be determined later on. We would have wanted to define the effectiveness factors of the OR technology by fixing $\gamma_j = 1 - \delta_j = \tau(\mathbf{u}^j) \cdot \epsilon$ for every $j \in N$. Unfortunately, the standard binary representation of $\tau(\mathbf{u}^j)$ may be much larger than the binary representation of ϕ for some j, and in particular, exponential in m. We handle this obstacle by estimating the vector evaluations as follows.

Note that since $x_i \le 3$ for every $0 \le i \le m+1$, and since $\mu > \left(\sum_{i=0}^{m+1} 3 \cdot 4^i\right)^5$, it follows that $\mu > (\sigma(\mathbf{x}))^5$ for every $\mathbf{x} \in \mathcal{W}$. The partial evaluation of $\mathbf{x}$ can be rewritten as $\tau_p(\mathbf{x}) = \sum_{j=0}^{\sigma(\mathbf{x})} \binom{\sigma(\mathbf{x})}{j} \mu^{-j}$, thus

$$\tau_p(\mathbf{x}) = \sum_{j=0}^{k-1} \binom{\sigma(\mathbf{x})}{j} \mu^{-j} + O\left(\mu^{-4k/5}\right) \tag{1}$$

for any $0 < k \le \sigma(\mathbf{x})$. Moreover, $\tau(\mathbf{x}) \le (1 + O(\mu^{-4/5}))\mu^7$. Fix $\chi = 2\mu^7$ (so that $\chi > \tau(\mathbf{x})$ for every vector $\mathbf{x} \in \mathcal{W}$).

Given a vector $\mathbf{x} = (x_0,\dots,x_{m+3}) \in \mathcal{W}$, let $\widetilde{\tau}_p(\mathbf{x}) = \sum_{j=0}^{\lceil 5(\kappa+7)/4 \rceil - 1} \binom{\sigma(\mathbf{x})}{j} \mu^{-j} = \tau_p(\mathbf{x}) - O(\mu^{-\kappa-7}) = \tau_p(\mathbf{x}) - O(\epsilon\mu^{-7})$ and $\widetilde{\tau}(\mathbf{x}) = \widetilde{\tau}_p(\mathbf{x}) \cdot \mu^{2x_{m+2}} \cdot \mu^{5x_{m+3}} = \tau(\mathbf{x}) - O(\epsilon)$. Note that the size of the binary representation of $\widetilde{\tau}(\mathbf{x})$ is polynomial (linear actually) in m. The technology t is now determined by fixing

$$\gamma_j = 1 - \delta_j = \widetilde{\tau}(\mathbf{u}^j)\epsilon \; = \; \tau(\mathbf{u}^j)\epsilon - O(\epsilon^2) \tag{2}$$

for all $j \in N$.

Next, we establish some important properties of the vector evaluations. We impose a lexicographic order on the vectors in $\mathcal{W}$: the vector $\mathbf{x} = (x_0, \ldots, x_{m+3})$ is lexicographically greater than the vector $\mathbf{y} = (y_0, \ldots, y_{m+3})$ if there exists a coordinate $0 \le j \le m+3$ such that $x_i = y_i$ for every $i > j$ and $x_j > y_j$. Clearly, for any two vectors $\mathbf{x}, \mathbf{y} \in \mathcal{W}$, the full evaluation of $\mathbf{x}$ is greater than the full evaluation of $\mathbf{y}$ if and only if $\mathbf{x}$ is lexicographically greater than $\mathbf{y}$.

Proposition 2.1. *Let $\mathbf{x} = (x_0, \ldots, x_{m+3})$ and $\mathbf{y} = (y_0, \ldots, y_{m+3})$ be two vectors in $\mathcal{W}$ such that $\mathbf{x}$ is lexicographically greater than $\mathbf{y}$. The difference $\tau(\mathbf{x}) - \tau(\mathbf{y})$ satisfies* (i) *if $x_{m+2} \neq y_{m+2}$ or $x_{m+3} \neq y_{m+3}$, then $\tau(\mathbf{x}) - \tau(\mathbf{y}) = (1 + o(1))\mu^{2x_{m+2}+5x_{m+3}}$, and* (ii) *if $x_{m+2} = y_{m+2}$ and $x_{m+3} = y_{m+3}$, then $\mu^{2x_{m+2}+5x_{m+3}-1} \le \tau(\mathbf{x}) - \tau(\mathbf{y}) \le O(\mu^{2x_{m+2}+5x_{m+3}-(4/5)})$.*

Proof. The bound in (i) follows immediately from the definition of full evaluation as the partial evaluation is $1 + o(1)$. To establish (ii), note that since $\tau(\mathbf{x}) > \tau(\mathbf{y})$ although $x_{m+2} = y_{m+2}$ and $x_{m+3} = y_{m+3}$, we must have $\tau_p(\mathbf{x}) > \tau_p(\mathbf{y})$. By the definition of partial evaluation, it follows that $\frac{\tau_p(\mathbf{x})}{\tau_p(\mathbf{y})} = \left(1 + \mu^{-1}\right)^{\sigma(\mathbf{x})-\sigma(\mathbf{y})}$, hence $1 + \mu^{-1} \le \frac{\tau_p(\mathbf{x})}{\tau_p(\mathbf{y})} \le 1 + O(\mu^{-4/5})$. Therefore

$$\mu^{-1} \le \tau_p(\mathbf{y})(1 + \mu^{-1} - 1) \;\le\; \tau_p(\mathbf{x}) - \tau_p(\mathbf{y}) \;\le\; \tau_p(\mathbf{y})(1 + O(\mu^{-4/5}) - 1) \;\le\; O(\mu^{-4/5})$$

The proof is completed as $\tau(\mathbf{x}) - \tau(\mathbf{y}) = \mu^{2x_{m+2}+5x_{m+3}}(\tau_p(\mathbf{x}) - \tau_p(\mathbf{y}))$. $\square$

Let $S \subseteq N$ be some contract and assume that $|S| = k > 0$. Let ν be the maximum among all constants hidden in the O notation of (2), that is, $\tau(\mathbf{u}^j)\epsilon - \gamma_j \le \nu\epsilon^2$ for every $j \in N$. By the definition of OR technologies, we have

$$\begin{aligned}
f(S) &= 1 - \prod_{j \in S}(1 - \delta_j) \prod_{j \in N-S}(1 - \gamma_j) \\
&= 1 - \prod_{j \in S} \epsilon \left(\tau(\mathbf{u}^j) - O(\epsilon)\right) \prod_{j \in N-S} \left(1 - \epsilon\left(\tau(\mathbf{u}^j) - O(\epsilon)\right)\right) \\
&= 1 - \epsilon^k \prod_{j \in S} \tau(\mathbf{u}^j) - \sum_{l=1}^{n+5} (-1)^l \epsilon^{k+l} \cdot O\left(\nu^l \chi \binom{n+5}{l}\right) \\
&= 1 - \tau(\mathbf{u}^S)\epsilon^k - \sum_{l=1}^{n+5} (-1)^l \epsilon^{k+l} \cdot O\left(\nu^l \chi \binom{n+5}{l}\right).
\end{aligned}$$

Taking $\epsilon < \left(\frac{1}{\nu\chi(n+5)}\right)^2$ guarantees that

$$f(S) = 1 - \tau(\mathbf{u}^S)\epsilon^k \pm O(\epsilon^{k+(1/2)}). \tag{3}$$

Following a similar line of arguments, we conclude that $f(\emptyset) = O(\epsilon^{1/2})$. The next proposition can now be established.

Proposition 2.2. *Let $S, S' \subseteq N$ be two contracts and let $k = |S|$, $k' = |S'|$. Then $f(S) < f(S')$ if and only if* (i) $k < k'$; or (ii) $k = k'$ and $\tau(\mathbf{u}^S) > \tau(\mathbf{u}^{S'})$.

Proof. The first claim follows immediately from (3) by taking $\epsilon \ll \chi^{-1}$. For the second claim, note that by (3), it is sufficient to prove that $\tau(\mathbf{u}^S) - \tau(\mathbf{u}^{S'}) = \omega(\epsilon^{1/2})$. This is guaranteed due to Proposition 2.1 by taking $\epsilon \ll \mu^{-2}$. $\square$

A direct consequence of Proposition 2.2 is that $f(S) = f(S')$ if and only if $|S| = |S'|$ and $\mathbf{u}^S = \mathbf{u}^{S'}$. The conditional payment to the agents in S, where $|S| = k$, can now be expressed as

$$\begin{aligned}
p(S) &= \sum_{j \in S} \frac{1}{f(S) - f(S-j)} \\
&= \sum_{j \in S} \left[1 - \tau(\mathbf{u}^S)\epsilon^k \pm O(\epsilon^{k+(1/2)}) - 1 + \tau(\mathbf{u}^{S-j})\epsilon^{k-1} \pm O(\epsilon^{k-(1/2)})\right]^{-1} \\
&= \sum_{j \in S} \left[\tau(\mathbf{u}^{S-j})\epsilon^{k-1} \pm O(\epsilon^{k-(1/2)})\right]^{-1} \\
&= \sum_{j \in S} \frac{1}{\tau(\mathbf{u}^{S-j})}\epsilon^{1-k} \pm O(\epsilon^{(3/2)-k}) \\
&= \sum_{j \in S} \frac{1}{\tau(\mathbf{u}^{S-j})}\epsilon^{1-k} \;\pm\; O(\epsilon^{(5/4)-k}),
\end{aligned}$$

where $S-j$ denotes the contract $S-\{j\}$ and the last equation follows by taking $\epsilon < (n+5)^{-4}$. Define $\pi(S) = \sum_{j\in S} \tau^{-1}(\mathbf{u}^{S-j})$, so that

$$p(S) = \pi(S)\epsilon^{1-k} \pm O(\epsilon^{(5/4)-k}). \tag{4}$$

Note that $\pi(S) < |S|$ for every contract $S \subseteq N$ since each term in the sum is smaller than 1.

Let $S \subseteq N$ be some contract and assume that $|S| = k > 0$. By plugging (3) and (4) into the definition of utility, we get

$$\begin{aligned} U_S(v) &= \left(1 - \tau(\mathbf{u}^S)\epsilon^k \pm O(\epsilon^{k+(1/2)})\right)\left(v - \pi(S)\epsilon^{1-k} \pm O(\epsilon^{(5/4)-k})\right) \\ &= v - \pi(S)\epsilon^{1-k} \pm O(\epsilon^{(5/4)-k}) - \tau(\mathbf{u}^S)v\epsilon^k + \pi(S)\tau(\mathbf{u}^S)\epsilon \pm O(\tau(\mathbf{u}^S)\epsilon^{5/4}) \\ &\pm O(v\epsilon^{k+(1/2)}) \pm O(\pi(S)\epsilon^{3/2}) \pm O(\epsilon^{7/4}) \\ &= v - \pi(S)\epsilon^{1-k} - \tau(\mathbf{u}^S)v\epsilon^k \pm O(\epsilon^{(5/4)-k}) \pm O(v\epsilon^{k+(1/2)}), \end{aligned}$$

where the last equation is guaranteed by taking $\epsilon < ((n+5)\chi)^{-4/3}$. For the empty contract, we have $p(\emptyset) = 0$ and $U_\emptyset(v) = v \cdot O(\epsilon^{1/2})$.

Consider two contracts $S, T \subseteq N$. Assuming that $f(S) \neq f(T)$, we refer to the payoff on which the lines $U_S(\cdot)$ and $U_T(\cdot)$ intersect as the *intersection payoff* of S and T, denoted $v[S,T]$, namely, $U_S(v[S,T]) = U_T(v[S,T])$. The next lemma correlates the intersection payoffs to the size of the contracts and to the vectors representing the contracts.

Lemma 2.3. *Let $S, S' \subseteq N$ be two contracts such that $f(S) \neq f(S')$. Define $k = |S|$ and $k' = |S'|$. The intersection payoff $v[S,S']$ satisfies (i) if $0 < k = k'$, then*

$$v[S,S'] = \epsilon^{1-2k}\frac{\pi(S') - \pi(S) \pm O(\epsilon^{1/4})}{\tau(\mathbf{u}^S) - \tau(\mathbf{u}^{S'}) \pm O(\epsilon^{1/2})};$$

and (ii) *if $k \neq k'$, $k, k' \geq 0$, then*

$$\Omega(\epsilon^{(5/4)-k-k'}) \leq v[S,S'] \leq O(\epsilon^{(3/4)-k-k'}).$$

(Observe that the case $0 = k = k'$ is irrelevant as there is only one empty contract.)

Proof. Assume without loss of generality that $k \leq k'$. Suppose first that $k > 0$. By comparing the utilities of S and S' on payoff $v[S,S']$, we get

$$\begin{aligned} &\pi(S)\epsilon^{1-k} + \tau(\mathbf{u}^S)v[S,S']\epsilon^k \pm O(\epsilon^{(5/4)-k}) \pm O(v[S,S']\epsilon^{k+(1/2)}) \\ &\quad = \pi(S')\epsilon^{1-k'} + \tau(\mathbf{u}^{S'})v[S,S']\epsilon^{k'} \pm O(\epsilon^{(5/4)-k'}) \pm O(v[S,S']\epsilon^{k'+(1/2)}), \end{aligned}$$

hence

$$v[S,S'] = \frac{\pi(S')\epsilon^{1-k'} - \pi(S)\epsilon^{1-k} \pm O(\epsilon^{(5/4)-k'})}{\tau(\mathbf{u}^S)\epsilon^k - \tau(\mathbf{u}^{S'})\epsilon^{k'} \pm O(\epsilon^{k+(1/2)})}.$$

By setting $k = k'$, (i) is established. Otherwise, if $0 < k < k'$, then, by taking $\epsilon < \min\{(n+5)^{-2}, \chi^{-2}\}$, we get

$$v[S,S'] = \frac{\pi(S')\epsilon^{1-k'} \pm O(\epsilon^{(5/4)-k'})}{\tau(\mathbf{u}^S)\epsilon^k \pm O(\epsilon^{k+(1/2)})} = \epsilon^{1-k'-k}\frac{\pi(S') \pm O(\epsilon^{1/4})}{\tau(\mathbf{u}^S) \pm O(\epsilon^{1/2})}. \tag{5}$$

It remains to consider the case $0 = k < k'$. Once again by comparing the utilities of S and S' on payoff $v[S,S']$, we have

$$v[S,S'] - \pi(S')\epsilon^{1-k'} - \tau(\mathbf{u}^{S'})v[S,S']\epsilon^{k'} \pm O(\epsilon^{(5/4)-k'}) \pm O(v[S,S']\epsilon^{k'+(1/2)}) = v[S,S'] \cdot O(\epsilon^{1/2}),$$

hence, by taking $\epsilon < \chi^{-2}$, we get

$$v[S,S'] = \frac{\pi(S')\epsilon^{1-k'} \pm O(\epsilon^{(5/4)-k'})}{1 - O(\epsilon^{1/2})}. \tag{6}$$

The bounds in (ii) are established by taking $\epsilon < (\max\{(n+5), \chi\})^{-4}$. $\square$

Protected vectors. Let $\mathbf{x} = (x_0, \ldots, x_{m+3})$ be a vector in $\mathcal{W}$. We say that $\mathbf{x}$ is *protected* if $x_{m+2} = x_{m+3} = 1$. For every $0 \le k \le n+5$, let $\Psi_k(\mathbf{x}) = \{S \subseteq N \mid \mathbf{u}^S = \mathbf{x} \text{ and } |S| = k\}$. We argue that if $\mathbf{x}$ is a protected vector in $\mathcal{W}$, and if $\Psi_k(\mathbf{x}) \neq \emptyset$, then at least one contract in $\Psi_k(\mathbf{x})$ is in the top envelope of the line collection $\{U_S(\cdot) \mid S \subseteq N\}$. We first establish some bounds related to $\pi(\cdot)$.

Proposition 2.4. *Let $S \subseteq N$ be a contract. If $\mathbf{u}^S$ is protected, then $\pi(S) = \Theta(\mu^{-2})$ and in particular, $\tau^{-1}(\mathbf{u}^{S-\beta}) \le \pi(S) \le (1 + O(\mu^{-3}))\tau^{-1}(\mathbf{u}^{S-\beta})$. If $\mathbf{u}^S$ is not protected, then $1 - o(1) \le \pi(S) \le |S|$.*

Proof. Suppose that $\mathbf{u}^S$ is protected. First observe that since $\alpha \in S - \beta$, it follows that $\tau(\mathbf{u}^{S-\beta}) = \Theta(\mu^2)$. Therefore if $\tau^{-1}(\mathbf{u}^{S-\beta}) \le \pi(S) \le (1 + O(\mu^{-3}))\tau^{-1}(\mathbf{u}^{S-\beta})$, then $\pi(S)$ is indeed $\Theta(\mu^{-2})$. Recall that $\pi(S) = \sum_{j \in S} \tau^{-1}(\mathbf{u}^{S-j}) = \sum_{j \in S - \{\alpha,\beta\}} \tau^{-1}(\mathbf{u}^{S-j}) + \tau^{-1}(\mathbf{u}^{S-\alpha}) + \tau^{-1}(\mathbf{u}^{S-\beta})$. For every $j \in S - \{\alpha, \beta\}$, we have $\frac{\tau^{-1}(\mathbf{u}^{S-j})}{\tau^{-1}(\mathbf{u}^{S-\beta})} = \frac{\tau(\mathbf{u}^j)}{\tau(\mathbf{u}^\beta)} = \frac{1+O(\mu^{-4/5})}{\mu^5}$, and $\frac{\tau^{-1}(\mathbf{u}^{S-\alpha})}{\tau^{-1}(\mathbf{u}^{S-\beta})} = \frac{\tau(\mathbf{u}^\alpha)}{\tau(\mathbf{u}^\beta)} = \frac{1}{\mu^3}$. Therefore $\frac{\pi(S)}{\tau^{-1}(\mathbf{u}^{S-\beta})} = \frac{(k-2)(1+O(\mu^{-4/5}))}{\mu^5} + \frac{1}{\mu^3} + 1$. Since $k - 2 \le n + 3 \le 3m + 3 \ll \mu$, we have $\pi(S) = (1 + \frac{O(1)}{\mu^3})\tau^{-1}(\mathbf{u}^{S-\beta})$.

Now suppose that $\mathbf{u}^S$ is not protected. We choose agent j' as follows. If $\alpha \in S$ or $\beta \in S$, then let j' be the (sole) agent in $S \cap \{\alpha, \beta\}$. (Recall that $\{\alpha, \beta\} \nsubseteq S$ as S is not protected.) Otherwise, let j' be any agent in S. Denote $\mathbf{u}^{S-j'} = (u_0, \ldots, u_{m+3})$. Since $\mathbf{u}^S$ is not protected, it follows that $u_{m+2} = u_{m+3} = 0$. Therefore $\tau(\mathbf{u}^{S-j'}) = \tau_p(\mathbf{u}^{S-j'}) = 1 + O(\mu^{-4/5})$, and $\pi(S) \ge \tau^{-1}(\mathbf{u}^{S-j'}) = 1 - o(1)$. $\square$

Proposition 2.5. *Let $S, S' \subseteq N$ be two contracts such that $\mathbf{u}^S$ is protected and $\tau(\mathbf{u}^S) > \tau(\mathbf{u}^{S'})$. Then $\pi(S') - \pi(S) = \Omega(\mu^{-3})$.*

Proof. If $\mathbf{u}^{S'}$ is not protected, then Proposition 2.4 guarantees that $\pi(S') - \pi(S) = \Omega(1)$. Assume that $\mathbf{u}^{S'}$ is protected. Since coordinate $m + 2$ is set in both $\mathbf{u}^S$ and $\mathbf{u}^{S'}$, we have $\frac{\tau(\mathbf{u}^{S-\beta})}{\tau(\mathbf{u}^{S'-\beta})} = \frac{\tau_p(\mathbf{u}^S)}{\tau_p(\mathbf{u}^{S'})} \ge 1 + \mu^{-1}$. By Proposition 2.4, we have $\pi(S') \ge \tau^{-1}(\mathbf{u}^{S'-\beta})$ and $\pi(S) \le (1 + O(\mu^{-3}))\tau^{-1}(\mathbf{u}^{S-\beta})$. Therefore $\pi(S') - \pi(S) \ge \tau^{-1}(\mathbf{u}^{S-\beta})(1 + \mu^{-1} - 1 - O(\mu^{-3}))$. As $\tau^{-1}(\mathbf{u}^{S-\beta}) = \Theta(\mu^{-2})$, it follows that $\pi(S') - \pi(S) = \Omega(\mu^{-3})$. $\square$

Geometric interpretation. Consider the collection $\mathcal{F}$ of all continuous functions $g : \mathbb{R} \to \mathbb{R}$. Let H be a finite subset of $\mathcal{F}$ and let g be a function in $\mathcal{F}$. We say that g is *dominated* by the functions in H if for every $v \in \mathbb{R}$, there exists a function $g' \in H$ such that $g(v) \le g'(v)$. Suppose that g and the functions in H are linear. Following some standard geometric arguments, one can show that if g is not dominated by any two functions in H, then g is not dominated by all functions in H. Given a contract $S \subseteq N$ and a subset of contracts $H \subseteq 2^N$, we say that S is *dominated* by the contracts in H if $U_S(\cdot)$ is dominated by the functions in $\{U_T(\cdot) \mid T \in H\}$.

We now turn to state the main lemma of this section, namely, that a contract assigned with a protected vector cannot be dominated by any two contracts assigned with different vectors.

Lemma 2.6. *Let $S \subseteq N$ be a contract such that $\mathbf{u}^S$ is protected and let $k = |S|$. Consider two contracts $R, T \notin \Psi_k(\mathbf{u}^S)$. Then there exists a payoff v for which $U_S(v) > \max\{U_R(v), U_T(v)\}$.*

Proof. Assume without loss of generality that $f(R) \le f(T)$. Proposition 2.2 implies that $f(S) \neq f(R)$ and $f(S) \neq f(T)$, hence it is sufficient to consider the case $f(R) < f(S) < f(T)$ (otherwise, S cannot be dominated by R and T). We prove that $v[R, S] < v[S, T]$. This establishes the lemma as it implies that $U_S(v) > \max\{U_R(v), U_T(v)\}$ for all $v[R, S] < v < v[S, T]$.

Let $k^R = |R|$ and $k^T = |T|$. We know, due to Proposition 2.2, that $k^R \le k \le k^T$. Lemma 2.3 is employed in order to show that it is sufficient to consider the case $k^R = k^T = k$. First if $k^R < k < k^T$, then $v[R, S] = O(\epsilon^{(3/4)-k^R-k})$ and $v[S, T] = \Omega(\epsilon^{(5/4)-k-k^T})$, thus $\frac{v[S,T]}{v[R,S]} = \Omega(\epsilon^{(1/2)-k^T+k^R}) \gg 1$, so the assertion holds. If $k^R < k = k^T$, then, by Proposition 2.2, we have $\tau(\mathbf{u}^S) > \tau(\mathbf{u}^T)$. By taking $\epsilon \ll \mu^{-12}$, Proposition 2.5 implies that $v[S, T] = \epsilon^{1-2k}\Omega(\mu^{-11})$. Hence, taking $\epsilon < \mu^{-22}$ guarantees that $v[S, T] = \Omega(\epsilon^{(3/2)-2k})$. As $v[R, S] = O(\epsilon^{(3/4)-k^R-k})$, we have $\frac{v[S,T]}{v[R,S]} = \Omega(\epsilon^{(3/4)-k+k^R}) \gg 1$, so the assertion holds. If $k^R = k < k^T$, then, by Proposition 2.2, we have $\tau(\mathbf{u}^R) > \tau(\mathbf{u}^S)$. By Propositions 2.4 and 2.1, it follows that $v[R, S] = O(\epsilon^{1-2k})$. As $v[S, T] = \Omega(\epsilon^{(5/4)-k-k^T})$, we have $\frac{v[S,T]}{v[R,S]} = \Omega(\epsilon^{(1/4)-k^T+k}) \gg 1$, so the assertion holds.

In what follows we assume that $k^R = k^T = k$ and $\tau(\mathbf{u}^R) > \tau(\mathbf{u}^S) > \tau(\mathbf{u}^T)$. We have to show that $\frac{\pi(S)-\pi(R)\pm O(\epsilon^{1/4})}{\tau(\mathbf{u}^R)-\tau(\mathbf{u}^S)\pm O(\epsilon^{1/2})} < \frac{\pi(T)-\pi(S)\pm O(\epsilon^{1/4})}{\tau(\mathbf{u}^S)-\tau(\mathbf{u}^T)\pm O(\epsilon^{1/2})}$. By taking $\epsilon < \chi^{-4}$, it is sufficient to prove that $(\pi(T) - \pi(S))(\tau(\mathbf{u}^R) - \tau(\mathbf{u}^S)) - (\pi(S) - \pi(R))(\tau(\mathbf{u}^S) - \tau(\mathbf{u}^T)) > \epsilon^{1/8}$. Instead, we take $\epsilon \ll \mu^{-8}$ and establish the stronger bound

$$\pi(T)\left(\tau(\mathbf{u}^R) - \tau(\mathbf{u}^S)\right) + \pi(R)\left(\tau(\mathbf{u}^S) - \tau(\mathbf{u}^T)\right) - \pi(S)\left(\tau(\mathbf{u}^R) - \tau(\mathbf{u}^T)\right) = \Omega(\mu^{-1})\,. \tag{7}$$

Since $\mathbf{u}^S$ is protected, and since $\tau(\mathbf{u}^R) > \tau(\mathbf{u}^S)$, we conclude that $\mathbf{u}^R$ must be protected too. As for $\mathbf{u}^T$, we have to consider both cases (protected or not). If $\mathbf{u}^T$ is not protected, then we establish Eq. (7) by proving that $\pi(T)(\tau(\mathbf{u}^R) - \tau(\mathbf{u}^S)) - \pi(S)\tau(\mathbf{u}^R) = \Omega(\mu^6)$. Propositions 2.4 and 2.1 guarantee that $\pi(T)(\tau(\mathbf{u}^R) - \tau(\mathbf{u}^S)) = \Omega(\mu^6)$ and $\pi(S)\tau(\mathbf{u}^R) = O(\mu^5)$, thus the assertion holds. In the remainder of this proof we assume that $\mathbf{u}^R$, $\mathbf{u}^S$ and $\mathbf{u}^T$ are all protected.

We will soon show that

$$\frac{\tau_p(\mathbf{u}^R)-\tau_p(\mathbf{u}^S)}{\tau_p(\mathbf{u}^T)}+\frac{\tau_p(\mathbf{u}^S)-\tau_p(\mathbf{u}^T)}{\tau_p(\mathbf{u}^R)}-\frac{\tau_p(\mathbf{u}^R)-\tau_p(\mathbf{u}^T)}{\tau_p(\mathbf{u}^S)} \;=\; \Omega(\mu^{-3}), \tag{8}$$

thus, by the definition of full evaluation, it follows that

$$\tau^{-1}(\mathbf{u}^{T-\beta})(\tau(\mathbf{u}^R)-\tau(\mathbf{u}^S))+\tau^{-1}(\mathbf{u}^{R-\beta})(\tau(\mathbf{u}^S)-\tau(\mathbf{u}^T))-\tau^{-1}(\mathbf{u}^{S-\beta})(\tau(\mathbf{u}^R)-\tau(\mathbf{u}^T))=\Omega(\mu^2).$$

As Proposition 2.1 guarantees that $\tau^{-1}(\mathbf{u}^{S-\beta})(\tau(\mathbf{u}^R)-\tau(\mathbf{u}^T))=o(\mu^5)$, we conclude that

$$\tau^{-1}(\mathbf{u}^{T-\beta})(\tau(\mathbf{u}^R)-\tau(\mathbf{u}^S))+\tau^{-1}(\mathbf{u}^{R-\beta})(\tau(\mathbf{u}^S)-\tau(\mathbf{u}^T))-(1+O(\mu^{-3}))\tau^{-1}(\mathbf{u}^{S-\beta})(\tau(\mathbf{u}^R)-\tau(\mathbf{u}^T)) \;=\; \Omega(\mu^2).$$

Eq. (7) follows due to Proposition 2.4 and the assertion holds.

To establish Eq. (8), let $a=\sigma(\mathbf{u}^R)-\sigma(\mathbf{u}^S)$ and $b=\sigma(\mathbf{u}^S)-\sigma(\mathbf{u}^T)$. Eq. (8) can be rewritten as

$$(1+\mu^{-1})^{a+b}+(1+\mu^{-1})^{-a}+(1+\mu^{-1})^{-b}-(1+\mu^{-1})^{-a-b}-(1+\mu^{-1})^{a}-(1+\mu^{-1})^{b}=\Omega(\mu^{-3})\,,$$

which is equivalent to

$$\begin{aligned}&\sum_{j=0}^{a+b}\binom{a+b}{j}\mu^{-j}+\sum_{j=0}^{\infty}(-1)^j\binom{a+j-1}{j}\mu^{-j}+\sum_{j=0}^{\infty}(-1)^j\binom{b+j-1}{j}\mu^{-j}\\&\quad-\sum_{j=0}^{\infty}(-1)^j\binom{a+b+j-1}{j}\mu^{-j}-\sum_{j=0}^{a}\binom{a}{j}\mu^{-j}-\sum_{j=0}^{b}\binom{b}{j}\mu^{-j}=\Omega(\mu^{-3}),\end{aligned} \tag{9}$$

due to the Taylor expansions

$$(1+z)^q=\sum_{j=0}^{q}\binom{q}{j}z^j \quad\text{and}\quad (1+z)^{-q}=\sum_{j=0}^{\infty}(-1)^j\binom{q+j-1}{j}z^j.$$

It is easy to verify that the j^{th} terms of the six sums on the left hand side of Eq. (9) cancel each other for $j=0,1,2$. For $j=3$, the terms on the left hand side of Eq. (9) sums up to

$$\left(\binom{a+b}{3}-\binom{a+2}{3}-\binom{b+2}{3}+\binom{a+b+2}{3}-\binom{a}{3}-\binom{b}{3}\right)\mu^{-3}=(a^2b+ab^2)\mu^{-3}=\Omega(\mu^{-3}).$$

It remains to show that the absolute value of the sums on the left hand side of Eq. (9) for $j=4,5,\ldots$ is $o(\mu^{-3})$. Instead we bound the larger expression

$$\begin{aligned}&\sum_{j=4}^{a+b}\binom{a+b}{j}\mu^{-j}+\sum_{j=4}^{\infty}\binom{a+j-1}{j}\mu^{-j}+\sum_{j=4}^{\infty}\binom{b+j-1}{j}\mu^{-j}\\&\quad+\sum_{j=4}^{\infty}\binom{a+b+j-1}{j}\mu^{-j}+\sum_{j=4}^{a}\binom{a}{j}\mu^{-j}+\sum_{j=4}^{b}\binom{b}{j}\mu^{-j}\\&\quad\le 6\cdot\sum_{j=4}^{\infty}\binom{a+b+j-1}{j}\mu^{-j}.\end{aligned}$$

As $\binom{a+b+j}{j+1}/\binom{a+b+j-1}{j}\le a+b$ for every positive j, we have

$$\sum_{j=4}^{\infty}\binom{a+b+j-1}{j}\mu^{-j}\le\binom{a+b+3}{4}\mu^{-4}\sum_{j=0}^{\infty}\left(\frac{a+b}{\mu}\right)^j=O\left(\mu^{\frac{4}{5}-4}\right)\cdot O(1)=o(\mu^{-3}),$$

where the equality in the middle follows from $\mu=\Omega\left((a+b)^5\right)$. Therefore Eq. (9) is satisfied and the assertion holds. □

The next corollary follows.

Corollary 2.7. *If* $\mathbf{x}$ *is a protected vector in* $\mathcal{W}$*, then for every* $0\le k\le n+5$*, either* $\Psi_k(\mathbf{x})=\emptyset$ *or there exist a contract* $S\in\Psi_k(\mathbf{x})$ *and a payoff* v *such that* S *is optimal for* v*.*

Consider the vector $\mathbf{x}=(1,\ldots,1)\in\mathcal{W}$. Recall that our goal is to decide whether there exists a contract S with $\mathbf{u}^S=\mathbf{x}$. Note that S is of size at least 4 as it must contain agents α,β,B and at least one more agent. For every $4\le k\le n+5$, Corollary 2.7 guarantees that if $\Psi_k(\mathbf{x})$ is not empty, then such a contract S is optimal for some payoff v_k^*. If we know the payoffs v_k^* for all $4\le k\le n+5$, then we can query all of them, thus deciding whether or not there exists a contract S with $\mathbf{u}^S=\mathbf{x}$.

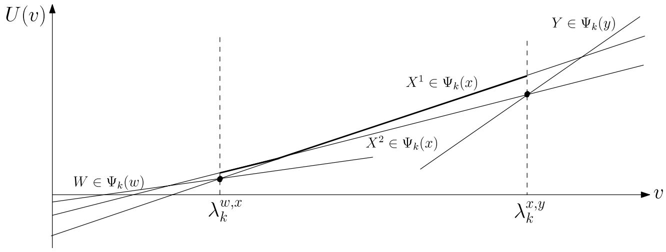

Fig. 2. The contracts $X^1 \in \Psi_k(\mathbf{x})$ and $W \in \Psi_k(\mathbf{w})$ realize $\lambda_k^{w,x}$; the contracts $X^2 \in \Psi_k(\mathbf{x})$ and $Y \in \Psi_k(\mathbf{y})$ realize $\lambda_k^{x,y}$. For every payoff $\lambda_k^{w,x} \le v \le \lambda_k^{x,y}$, there exists a contract in $\Psi_k(\mathbf{x})$ which is optimal for v (bold lines).

Consider some $4 \le k \le n+5$ and assume that $\Psi_k(\mathbf{x})$ is not empty. Recall that $\mathbf{u}^A = (0, \ldots, 0, 1, 0, 0)$, $\mathbf{u}^B = (1, 0, 0, \ldots, 0, 1, 0, 0)$ and $\mathbf{u}^C = (2, 0, 0, \ldots, 0, 1, 0, 0)$. Let $\mathbf{w} = (2, 1, 1, \ldots, 1) \in \mathcal{W}$ and let $\mathbf{y} = (0, 1, 1, \ldots, 1) \in \mathcal{W}$. Since $\mathbf{u}^A$, $\mathbf{u}^B$ and $\mathbf{u}^C$ determine the value of coordinates 0 and $m+1$ in $\mathcal{W}$ without affecting any other coordinate, and since $B \in S$ and $A, C \notin S$ for every contract S such that $\mathbf{u}^S = \mathbf{x}$, it follows that $\Psi_k(\mathbf{w}) \neq \emptyset$ and $\Psi_k(\mathbf{y}) \neq \emptyset$ (as $\Psi_k(\mathbf{x}) \neq \emptyset$ and agent B can be replaced by agent A or C in S).

Let $\lambda_k^{w,x} = \max\{v[S,T] \mid S \in \Psi_k(\mathbf{w}) \text{ and } T \in \Psi_k(\mathbf{x})\}$ and let $\lambda_k^{x,y} = \min\{v[S,T] \mid S \in \Psi_k(\mathbf{x}) \text{ and } T \in \Psi_k(\mathbf{y})\}$ (see Fig. 2). Note that $\lambda_k^{w,x}$ and $\lambda_k^{x,y}$ are well defined as $\Psi_k(\mathbf{w})$, $\Psi_k(\mathbf{x})$ and $\Psi_k(\mathbf{y})$ are not empty. Define $v_k^* = \frac{\epsilon^{1-2k}}{(1+\xi\mu^{-1})\mu^9}$, where $\xi = 2 \cdot \sum_{j=0}^{m+1} 4^j$. Observe that the binary representation of v_k^* is polynomial in m.

Lemma 2.8. *The payoff v_k^* satisfies $\lambda_k^{w,x} < v_k^* < \lambda_k^{x,y}$.*

Proof. Define $\mathbf{w}' = (2, 1, 1, \ldots, 1, 0) \in \mathcal{W}$, $\mathbf{x}' = (1, \ldots, 1, 0) \in \mathcal{W}$ and $\mathbf{y}' = (0, 1, 1, \ldots, 1, 0) \in \mathcal{W}$. By Lemma 2.3 and by Proposition 2.4, we have $\lambda_k^{w,x} \le \epsilon^{1-2k} \frac{(1+O(\mu^{-3}))\tau^{-1}(\mathbf{x}')-\tau^{-1}(\mathbf{w}')+O(\epsilon^{1/4})}{\tau(\mathbf{w})-\tau(\mathbf{x})-O(\epsilon^{1/2})}$ and $\lambda_k^{x,y} \ge \epsilon^{1-2k} \frac{\tau^{-1}(\mathbf{y}')-(1+O(\mu^{-3}))\tau^{-1}(\mathbf{x}')-O(\epsilon^{1/4})}{\tau(\mathbf{x})-\tau(\mathbf{y})+O(\epsilon^{1/2})}$. Propositions 2.1 and 2.5 imply that

$$\lambda_k^{w,x} \le \epsilon^{1-2k} \left(\frac{(1+O(\mu^{-3}))\tau^{-1}(\mathbf{x}') - \tau^{-1}(\mathbf{w}')}{\tau(\mathbf{w}) - \tau(\mathbf{x})} + o(\epsilon^{1/4}) \right)$$

and

$$\lambda_k^{x,y} \ge \epsilon^{1-2k} \left(\frac{\tau^{-1}(\mathbf{y}') - (1+O(\mu^{-3}))\tau^{-1}(\mathbf{x}')}{\tau(\mathbf{x}) - \tau(\mathbf{y})} - o(\epsilon^{1/4}) \right).$$

As $\frac{\tau^{-1}(\mathbf{y}')}{\tau^{-1}(\mathbf{x}')} = \frac{\tau^{-1}(\mathbf{x}')}{\tau^{-1}(\mathbf{w}')} = \frac{\tau(\mathbf{x})}{\tau(\mathbf{y})} = \frac{\tau(\mathbf{w})}{\tau(\mathbf{x})} = 1 + \mu^{-1}$, it follows that

$$\lambda_k^{w,x} \le \epsilon^{1-2k} \left(\frac{\tau^{-1}(\mathbf{w}')}{\tau(\mathbf{x})} (1 + O(\mu^{-2})) + o(\epsilon^{1/4}) \right)$$

and

$$\lambda_k^{x,y} \ge \epsilon^{1-2k} \left(\frac{\tau^{-1}(\mathbf{x}')}{\tau(\mathbf{y})} (1 - O(\mu^{-2})) - o(\epsilon^{1/4}) \right).$$

By the definition of full evaluation, we have $\frac{\tau^{-1}(\mathbf{w}')}{\tau(\mathbf{x})} = (1+\mu^{-1})^{-(\xi+1)}\mu^{-9}$ and $\frac{\tau^{-1}(\mathbf{x}')}{\tau(\mathbf{y})} = (1+\mu^{-1})^{-(\xi-1)}\mu^{-9}$, thus taking $\epsilon < \mu^{-44}$ guarantees that $\lambda_k^{w,x} \le \epsilon^{1-2k}(1+\mu^{-1})^{-(\xi+1)}\mu^{-9}(1+O(\mu^{-2}))$ and $\lambda_k^{x,y} \ge \epsilon^{1-2k}(1+\mu^{-1})^{-(\xi-1)}\mu^{-9}(1-O(\mu^{-2}))$. Since $\mu > \xi^5$, it follows that $(1+\mu^{-1})^{\xi+1} = (1+\mu^{-1})(1+\xi\mu^{-1}+O(\mu^{-8/5})) \ge (1+\mu^{-1})(1+\xi\mu^{-1})$ and $(1+\mu^{-1})^{\xi-1} = 1 + (\xi-1)\mu^{-1} + O(\mu^{-8/5}) \le 1 + \xi\mu^{-1} - \mu^{-1}/2$, hence

$$\frac{\lambda_k^{w,x}}{v_k^*} \le \frac{(1+\xi\mu^{-1})(1+O(\mu^{-2}))}{(1+\mu^{-1})^{\xi+1}} \le \frac{1+O(\mu^{-2})}{1+\mu^{-1}} < 1$$

and

$$\frac{\lambda_k^{x,y}}{v_k^*} \ge \frac{(1+\xi\mu^{-1})(1-O(\mu^{-2}))}{(1+\mu^{-1})^{\xi-1}} \ge \frac{1+\xi\mu^{-1}-O(\mu^{-2})}{1+\xi\mu^{-1}-\mu^{-1}/2} > 1.$$

The assertion follows. □

The analysis is completed with the following lemma, which together with Lemma 2.8 derives Theorem 1.

Lemma 2.9. *The optimal contract for the payoff v is in $\Psi_k(\mathbf{x})$ for every $\lambda_k^{w,x} < v < \lambda_k^{x,y}$.*

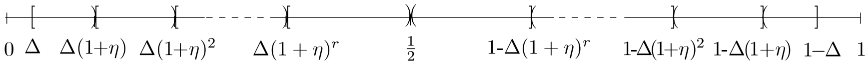

Fig. 3. A scale of precision $1 + \eta$.

Proof. Consider an arbitrary payoff $\lambda_k^{w,x} < \bar{v} < \lambda_k^{x,y}$ and suppose toward deriving contradiction that there exists a contract $T \notin \Psi_k(\mathbf{x})$ such that T is optimal for $\bar{v}$. Recall that Proposition 2.2 implies that $f(S^w) < f(S^x) < f(S^y)$ for every three contracts $S^w \in \Psi_k(\mathbf{w})$, $S^x \in \Psi_k(\mathbf{x})$ and $S^y \in \Psi_k(\mathbf{y})$. Therefore by the definition of $\lambda_k^{w,x}$ and $\lambda_k^{x,y}$, it follows that $T \notin \Psi_k(\mathbf{w})$ and $T \notin \Psi_k(\mathbf{y})$. Let $R^w \in \Psi_k(\mathbf{w})$ and $R^x \in \Psi_k(\mathbf{x})$ be the contracts that realize $\lambda_k^{w,x}$ and let $S^x \in \Psi_k(\mathbf{x})$ and $S^y \in \Psi_k(\mathbf{y})$ be the contracts that realize $\lambda_k^{x,y}$, i.e., $v[R^w, R^x] = \lambda_k^{w,x}$ and $v[S^x, S^y] = \lambda_k^{x,y}$.

We argue that T must satisfy $f(R^w) \leq f(T) \leq f(S^y)$. This can be justified as follows. If $f(T) < f(R^w)$, then since $U_T(\bar{v}) > U_{R^w}(\bar{v})$, we have $U_T(v) > U_{R^w}(v)$ for every $v < \bar{v}$. As $U_{R^x}(v) > U_{R^w}(v)$ for every $v > \lambda_k^{w,x}$, and since $\bar{v} > \lambda_k^{w,x}$, it follows that R^w is dominated by T and R^x, in contradiction to Lemma 2.6. The case where $f(T) > f(S^y)$ is analogous. Proposition 2.2 implies that $|T| = k$ and $\tau(\mathbf{y}) < \tau(\mathbf{u}^T) < \tau(\mathbf{w})$ as otherwise, we get $f(T) < f(R^w)$ or $f(T) > f(S^y)$. But this implies that $\mathbf{u}^T = \mathbf{x}$, in contradiction to the assumption, as $\mathbf{x}$ is the only vector in $\mathcal{W}$ which is lexicographically smaller than $\mathbf{w}$ and greater than $\mathbf{y}$. The assertion follows. □

3. Approximations

In this section we provide approximation results for the optimal contract problem. The section begins with Section 3.1, which provides a general approximation scheme for arbitrary SP technologies. This scheme is carried through Algorithm `Calibrate`, presented in this section. Based on the general scheme presented in Section 3.1, in Section 3.2, we construct an FPTAS for arbitrary OR technologies. Another use of Algorithm `Calibrate` is presented in Section 3.3, where we provide a scheme that approximates all but a small fraction of the relevant payoffs for any SP technology. Finally, Section 3.4 deals with the most general case, and shows that every technology admits a polynomial size collection that approximates the optimal contract.

3.1. A general scheme

Consider some technology $t = \langle N, \{\gamma_i\}_{i=1}^n, \{\delta_i\}_{i=1}^n, \{c_i\}_{i=1}^n, \varphi\rangle$ and let $S \subseteq N$ be an arbitrary contract. We first observe that if φ is the AND function, then the effectiveness of S is given by $f(S) = \prod_{i\in S}\delta_i \prod_{i\in N-S}\gamma_i$. For the OR function, we have $f(S) = 1 - \prod_{i\in S}(1-\delta_i)\prod_{i\in N-S}(1-\gamma_i)$. Therefore if all agents shirk, then the effectiveness under an AND technology is $\prod_{i\in N}\gamma_i$. On the other hand, if all agents exert effort, then the effectiveness under an OR technology is $1 - \prod_{i\in N}(1-\delta_i)$. Fix $\Delta = \min\left\{\prod_{i\in N}\gamma_i, \prod_{i\in N}(1-\delta_i)\right\}$. It is easy to verify that if t is an AND technology or an OR technology, then $f(S) \in [\Delta, 1-\Delta]$. The following lemma generalizes this property to the whole range of technologies.

Lemma 3.1. *The effectiveness $f(S)$ satisfies $f(S) \in [\Delta, 1-\Delta]$ regardless of the choice of the monotone Boolean function $\varphi : \{0, 1\}^n \to \{0, 1\}$.*

Proof. Consider the underlying n-variables truth table of the Boolean function $\varphi(x_1, \ldots, x_n)$. Since φ is not a function of any $n-1$ variables, it cannot assign 0 to all rows of the table. Therefore, the minimum possible effectiveness is achieved when φ assigns 1 to exactly one row (otherwise, it can achieve a lower value by replacing a single 1 value with 0). By the monotonicity of φ, this single row must correspond to $x_1 = \cdots = x_n = 1$. (This is exactly the truth table of the AND function.) Clearly, the minimum possible effectiveness is achieved when all agents shirk. Combined together, the minimum possible effectiveness is simply $\mathbf{P}(x_1 = 1 \wedge \cdots \wedge x_n = 1 \mid a = (0, \ldots, 0)) = \prod_{i\in N}\gamma_i$. The proof that the maximum possible effectiveness is $\mathbf{P}(x_1 = 1 \vee \cdots \vee x_n = 1 \mid a = (1, \ldots, 1)) = 1 - \prod_{i\in N}(1-\delta_i)$ is analogous. □

Our scheme is executed by an algorithm, referred to as Algorithm `Calibrate`, which we shall soon present. The description of Algorithm `Calibrate` requires some preparation.

Consider an SP technology $t = \langle N, \{\gamma_i\}_{i=1}^n, \{\delta_i\}_{i=1}^n, \{c_i\}_{i=1}^n, \varphi\rangle$ input to Algorithm `Calibrate` and let $0 < \rho \leq 1$ be the *performance parameter* of the algorithm. Algorithm `Calibrate` generates a collection $\mathcal{C}$ of contracts in time $O\left(\frac{n^3\log^2(1/\Delta)}{\rho^2}\right)$. (Note that the binary representation of $\{\gamma_i\}_{i=1}^n$ and $\{\delta_i\}_{i=1}^n$ requires $\Omega(\log(1/\Delta))$ bits.) We will soon prove that for every contract $T \subseteq N$, there exists a contract $S \in \mathcal{C}$ such that $f(S) \geq \frac{f(T)}{(1+\rho)}$, and $p(S) \leq (1+\rho)p(T)$.

Let $\eta = \frac{\rho\ln 2}{2n-1}$, and let $r = \max\left\{k \in \mathbb{Z}_{\geq 0} \mid \Delta(1+\eta)^k < \frac{1}{2}\right\}$. Since $r < \log_{1+\eta}\left(\frac{1}{2\Delta}\right) = \log\frac{1}{2\Delta}\cdot\log_{1+\eta}(2)$, and since $\log_{1+\eta}(2) \leq \frac{1}{\eta}$, we conclude that $r < \frac{1}{\eta}\log\frac{1}{\Delta}$. We partition the interval $[\Delta, 1-\Delta]$ into $2r+3$ smaller intervals $[\Delta, \Delta(1+\eta)), [\Delta(1+\eta), \Delta(1+\eta)^2), \ldots, [\Delta(1+\eta)^{r-1}, \Delta(1+\eta)^r), [\Delta(1+\eta)^r, \frac{1}{2}), [\frac{1}{2}, \frac{1}{2}], (\frac{1}{2}, 1-\Delta(1+\eta)^r], (1-\Delta(1+\eta)^r, 1-\Delta(1+\eta)^{r-1}], \ldots, (1-\Delta(1+\eta)^2, 1-\Delta(1+\eta)], (1-\Delta(1+\eta), 1-\Delta]$. The collection of these smaller intervals is called the *scale*. Refer to Fig. 3 for an illustration of the scale. The *precision* of the scale is defined as $1+\eta$. We say that contract S is *calibrated* to interval $\mathcal{I}$ in the scale if $f(S) \in \mathcal{I}$. (Recall that Lemma 3.1 implies that every contract is calibrated to some interval in the scale.)

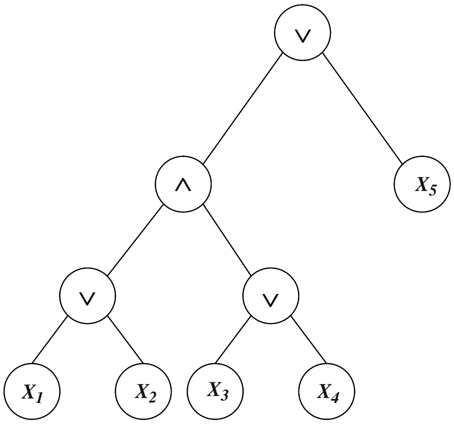

Fig. 4. The composition tree of the SP Boolean function $\varphi(x_1, x_2, x_3, x_4, x_5) = ((x_1 \vee x_2) \wedge (x_3 \vee x_4)) \vee x_5$.

Observation 3.2. *Let $S, S' \in N$ be two contracts. The scale is designed to ensure that if S and S' are calibrated to the same interval, then $\frac{f(S')}{1+\eta} \leq f(S) \leq (1+\eta)f(S')$ and $\frac{1-f(S')}{1+\eta} \leq 1 - f(S) \leq (1+\eta)(1-f(S'))$.*

Throughout the execution, Algorithm `Calibrate` maintains a collection $\mathcal{C}$ of contracts. The algorithm guarantees that no two contracts in $\mathcal{C}$ are calibrated to the same interval, thus $|\mathcal{C}| \leq 2r + 3$ at any given moment.

Every SP function φ is constructed inductively from two simpler SP functions by either a series composition or by a parallel composition. Therefore the function φ can be represented by a full binary tree $\mathcal{T}$, referred to as the *composition tree* of φ. The leaves of $\mathcal{T}$ represents the identity functions of φ's arguments. An internal node is said to be an $\wedge$*-node* (respectively, an $\vee$*-node*) if it represents a series (resp., parallel) composition of the functions represented by its children. (Refer to Fig. 4 for an illustration.)

Consider the SP technology $t = \langle N, \{\gamma_i\}_{i=1}^n, \{\delta_i\}_{i=1}^n, \{c_i\}_{i=1}^n, \varphi\rangle$ and let $\mathcal{T}$ be the tree that represents the Boolean function φ. Let x be some node in $\mathcal{T}$ and consider the subtree $\mathcal{T}_x$ of $\mathcal{T}$ rooted at x. The subtree $\mathcal{T}_x$ corresponds to some (SP) subtechnology t_x of t. Let N_x denote the set of agents in t_x (corresponding to the leaves of $\mathcal{T}_x$) and let m_x denote the number of nodes in $\mathcal{T}_x$ (as $\mathcal{T}$ is a full binary tree, we have $m_x = 2|N_x| - 1$). Given some contract $S \subseteq N_x$, we denote the effectiveness and payment of S under t_x by $f_x(S)$ and $p_x(S)$, respectively.

Suppose that x is an internal node in $\mathcal{T}$ with left child l and right child r. Let $S = L \cup R$ be some contract in t_x, where $L \subseteq N_l$ and $R \subseteq N_r$. Clearly, if x is an $\wedge$-node, then $f_x(S) = f_l(L) \cdot f_r(R)$, and if x is an $\vee$-node, then $f_x(S) = 1 - (1 - f_l(L))(1 - f_r(R))$. It is simple to verify that if x is an $\wedge$-node, then $p_x(S) = \frac{p_l(L)}{f_r(R)} + \frac{p_r(R)}{f_l(L)}$, and if x is an $\vee$-node, then $p_x(S) = \frac{p_l(L)}{1-f_r(R)} + \frac{p_r(R)}{1-f_l(L)}$.

With this we are ready to present Algorithm `Calibrate`.

Algorithm `Calibrate` traverses the composition tree $\mathcal{T}$ in a postorder fashion. Consider some leaf x in $\mathcal{T}$ that corresponds to agent $i \in N$. The algorithm calibrates the contracts $\emptyset$ and $\{i\}$ to a (fresh) scale according to their effectiveness under the technology t_x, that is, $f_x(\emptyset) = \gamma_i$ and $f_x(\{i\}) = \delta_i$. If both $\emptyset$ and $\{i\}$ are calibrated to the same interval $\mathcal{I}$, then $\{i\}$ is removed from the scale. The resulting contract(s) in the scale constitutes the collection $\mathcal{C}_x$.

Now, consider some internal node x in $\mathcal{T}$ with left child l and right child r and suppose that the algorithm has already constructed the collections $\mathcal{C}_l$ and $\mathcal{C}_r$ for the technologies t_l and t_r, respectively. The collection $\mathcal{C}_x$ for the technology t_x is constructed as follows. Let $\mathcal{S} = \{L \cup R \mid L \in \mathcal{C}_l \text{ and } R \in \mathcal{C}_r\}$. (Note that $\mathcal{S}$ contains $|\mathcal{C}_l| \cdot |\mathcal{C}_r| = O(r^2)$ contracts of the technology t_x.) The contracts in $\mathcal{S}$ are calibrated to a (fresh) scale according to the effectiveness function $f_x(\cdot)$. Consequently, there may exist some interval in the new scale to which two (or more) contracts are calibrated (a conflict).

Let $\mathcal{I}$ be an interval in the scale and suppose that $S_1, \ldots, S_k \in \mathcal{S}$ were all calibrated to $\mathcal{I}$ ($k > 1$), that is, $f_x(S_i) \in \mathcal{I}$ for every $1 \leq i \leq k$. Assume without loss of generality that S_k admits a minimum payment under t_x, i.e., $p_x(S_k) \leq p_x(S_i)$ for every $1 \leq i < k$. The algorithm then resolves the conflict by removing the contracts $S_1, \ldots, S_{k-1}$ from the scale so that S_k remains the only contract calibrated to $\mathcal{I}$. In that case we say that the contracts $S_1, \ldots, S_{k-1}$ were *compensated* by the contract S_k. The contracts that remain in the scale constitutes the collection $\mathcal{C}_x$. Thus the new collection $\mathcal{C}_x$ contains at most one contract for every interval and we may proceed with the next stage of the algorithm. At the end of this postorder process, when Algorithm `Calibrate` reaches the root z of $\mathcal{T}$, it returns the collection $\mathcal{C} = \mathcal{C}_z$.

We turn to the analysis of Algorithm `Calibrate`. The running time of the algorithm is determined by the number of nodes in $\mathcal{T}$ (which is $2n - 1$) and by the size of the collection $\mathcal{C}_x$ for every node x in the tree. The latter cannot exceed the number of intervals in the scale which is $O\left(\frac{1}{\eta}\log\frac{1}{\Delta}\right)$. In order to analyze the performance guarantee of the algorithm, we first define the following notion. Given two contracts $S, S' \subseteq N$ and some real $\alpha > 1$, we say that S is an α*-estimation* of S' under the technology t if the following three conditions hold:

$$\frac{f(S')}{\alpha} \leq f(S) \leq \alpha f(S'); \tag{10}$$

$$\frac{1-f(S')}{\alpha} \leq 1-f(S) \leq \alpha(1-f(S')); \quad \text{and} \tag{11}$$
$$p(S) \leq \alpha p(S'). \tag{12}$$

We say that a collection $\mathcal{S}$ of contracts is an α-*estimation* of the technology t if for every contract $S' \subseteq N$ there exists a contract $S \in \mathcal{S}$ such that S is an α-estimation of S' under t. The following observation serves as a key ingredient in the proof of Lemma 3.4, which is the main lemma in this section.

Observation 3.3. *For any choice of reals $0 < a, b, a', b' < 1$ and $\alpha, \beta > 1$, if $\frac{a'}{\alpha} \leq a \leq \alpha a'$ and $\frac{b'}{\beta} \leq b \leq \beta b'$, then $\frac{1-(1-a')(1-b')}{\alpha\beta} \leq 1-(1-a)(1-b) \leq \alpha\beta(1-(1-a')(1-b'))$.*

We are now ready to establish Lemma 3.4.

Lemma 3.4. *The collection $\mathcal{C}_x$ is a $(1+\eta)^{m_x}$-estimation of the technology t_x for every node x in the composition tree $\mathcal{T}$.*

Proof. The proof is by induction on the height of x in $\mathcal{T}$. The assertion trivially holds if x is a leaf. (Recall that Observation 3.2 guarantees that if the contracts $\emptyset$ and $\{i\}$ are calibrated to the same interval under t_x, then $\emptyset$ is a $(1+\eta)$-estimation of $\{i\}$.) Consider some internal node x in $\mathcal{T}$ and assume that the assertion holds for x's left child l and right child r. Let $S' = L' \cup R'$ be some contract in t_x, where $L' \subseteq N_l$ and $R' \subseteq N_r$. By the inductive hypothesis, there exist some contracts $L \in \mathcal{C}_l$ and $R \in \mathcal{C}_r$ such that L is a $(1+\eta)^{m_l}$-estimation of L' under t_l and R is a $(1+\eta)^{m_r}$-estimation of R' under t_r.

We argue that the contract $L \cup R$ is a $(1+\eta)^{m_l+m_r}$-estimation of $S' = L' \cup R'$ under the technology t_x. If x is an $\wedge$-node, then $f_x(L \cup R) = f_l(L) \cdot f_r(R)$ and Condition (10) holds trivially. Condition (11) holds by plugging $a = 1 - f_l(L)$, $b = 1 - f_r(R)$, $a' = 1 - f_l(L')$, and $b' = 1 - f_r(R')$ into Observation 3.3 with $\alpha = (1+\eta)^{m_l}$ and $\beta = (1+\eta)^{m_r}$. If x is an $\vee$-node, then $f_x(L \cup R) = 1 - (1 - f_l(L))(1 - f_r(R))$ and Condition (11) holds trivially. Condition (10) holds by plugging $a = f_l(L)$, $b = f_r(R)$, $a' = f_l(L')$, and $b' = f_r(R')$ into Observation 3.3 with $\alpha = (1+\eta)^{m_l}$ and $\beta = (1+\eta)^{m_r}$.

It remains to prove that Condition (12) holds. If x is an $\wedge$-node, then

$$\begin{aligned} p_x(L \cup R) &= \frac{1}{f_r(R)} p_l(L) + \frac{1}{f_l(L)} p_r(R) \\ &\leq \frac{(1+\eta)^{m_r}}{f_r(R')} (1+\eta)^{m_l} p_l(L') + \frac{(1+\eta)^{m_l}}{f_l(L')} (1+\eta)^{m_r} p_r(R') \\ &= (1+\eta)^{m_l+m_r} p_x(L' \cup R') \, . \end{aligned}$$

On the other hand, if x is an $\vee$-node, then

$$\begin{aligned} p_x(L \cup R) &= \frac{1}{1-f_r(R)} p_l(L) + \frac{1}{1-f_l(L)} p_r(R) \\ &\leq \frac{(1+\eta)^{m_r}}{1-f_r(R')} (1+\eta)^{m_l} p_l(L') + \frac{(1+\eta)^{m_l}}{1-f_l(L')} (1+\eta)^{m_r} p_r(R') \\ &= (1+\eta)^{m_l+m_r} p_x(L' \cup R') \, . \end{aligned}$$

The argument follows.

The contract $L \cup R$ is considered by the algorithm in the scale that corresponds to node x. If $L \cup R$ survives and finds its way to $\mathcal{C}_x$, then the proof is completed. Assume that $L \cup R$ is compensated by some contract $S \in \mathcal{C}_x$. We prove that S is a $(1+\eta)^{m_x}$-estimation of S'. Condition (12) holds as $p_x(S) \leq p_x(L \cup R)$. Conditions (10) and (11) follow from Observation 3.2 since $L \cup R$ is a $(1+\eta)^{m_l+m_r}$-estimation of S', and since $m_x = m_l + m_r + 1$. $\square$

Lemma 3.4 implies that $\mathcal{C}$ serves as a $(1+\eta)^{2n-1}$-estimation of t. By the definition of $\eta = \frac{\rho \ln 2}{2n-1}$, we have $(1+\eta)^{2n-1} \leq e^{\rho \ln 2} = 2^\rho \leq 1 + \rho$, which establishes the following corollary.

Corollary 3.5. *Given an* SP *technology $t = \langle N, \{\gamma_i\}_{i=1}^n, \{\delta_i\}_{i=1}^n, \{c_i\}_{i=1}^n, \varphi \rangle$ and a performance parameter $0 < \rho \leq 1$, it is guaranteed that Algorithm* `Calibrate` *generates a collection $\mathcal{C} \subseteq 2^N$ that serves as a $(1+\rho)$-estimation of t in time $O\left(\frac{n^3 \log^2(1/\Delta)}{\rho^2}\right)$.*

3.2. An FPTAS for OR *technologies*

In this section we establish an FPTAS for OR technologies.

Theorem 2. *The problem of computing the optimal contract in* OR *technologies admits a fully polynomial-time approximation scheme (FPTAS).*

The proof of Theorem 2 requires several preparations. We first establish several properties of OR technologies which will be used to present the desired FPTAS starting with the sub-modularity of OR technologies. We say that a function $h : 2^N \to \mathbb{R}$ is *strictly sub-modular* if $h(S) + h(T) \geq h(S \cup T) + h(S \cap T)$ for every $S, T \subseteq N$, where equality holds (if and) only if $S \subseteq T$ or $T \subseteq S$.

Lemma 3.6. *The effectiveness function of every OR technology is strictly sub-modular.*

Proof. Consider an arbitrary OR technology $t = \langle N, \{\gamma_i\}_{i=1}^n, \{\delta_i\}_{i=1}^n, c, \varphi\rangle$. We need to show that $f(S) + f(T) > f(S \cup T) + f(S \cap T)$ for every two contracts $S, T \subseteq N$ such that $S - T \neq \emptyset$ and $T - S \neq \emptyset$. By definition, we have

$$f(S) + f(T) = 2 - \prod_{i \in S}(1 - \delta_i) \prod_{i \in N-S}(1 - \gamma_i) - \prod_{i \in T}(1 - \delta_i) \prod_{i \in N-T}(1 - \gamma_i)$$

and

$$f(S \cup T) + f(S \cap T) = 2 - \prod_{i \in S \cup T}(1 - \delta_i) \prod_{i \in N-(S \cup T)}(1 - \gamma_i) - \prod_{i \in S \cap T}(1 - \delta_i) \prod_{i \in N-(S \cap T)}(1 - \gamma_i).$$

Dividing both equations by $\prod_{i \in S \cap T}(1 - \delta_i) \prod_{i \in N-(S \cup T)}(1 - \gamma_i)$, we conclude that it is sufficient to prove that

$$\begin{aligned} &\prod_{i \in S-T}(1 - \delta_i) \prod_{i \in T-S}(1 - \gamma_i) + \prod_{i \in T-S}(1 - \delta_i) \prod_{i \in S-T}(1 - \gamma_i) \\ &\quad - \prod_{i \in S-T}(1 - \delta_i) \prod_{i \in T-S}(1 - \delta_i) - \prod_{i \in S-T}(1 - \gamma_i) \prod_{i \in T-S}(1 - \gamma_i) < 0. \end{aligned}$$

The last inequality holds if and only if

$$\begin{aligned} &\prod_{i \in S-T}(1 - \delta_i) \left(\prod_{i \in T-S}(1 - \gamma_i) - \prod_{i \in T-S}(1 - \delta_i) \right) \\ &\quad + \prod_{i \in S-T}(1 - \gamma_i) \left(\prod_{i \in T-S}(1 - \delta_i) - \prod_{i \in T-S}(1 - \gamma_i) \right) < 0, \end{aligned}$$

which in turn, can be rewritten as

$$\left(\prod_{i \in T-S}(1 - \gamma_i) - \prod_{i \in T-S}(1 - \delta_i) \right) \left(\prod_{i \in S-T}(1 - \delta_i) - \prod_{i \in S-T}(1 - \gamma_i) \right) < 0.$$

The assertion follows as $\delta_i > \gamma_i$ for every $i \in N$. □

Consider an arbitrary OR technology $t = \langle N, \{\gamma_i\}_{i=1}^n, \{\delta_i\}_{i=1}^n, \{c_i\}_{i=1}^n, \varphi\rangle$. Let $T \subseteq N$ be some contract, $|T| \geq 2$, and consider some partition $T = R_1 \cup R_2$, $R_1 \cap R_2 = \emptyset$, such that $|R_1|, |R_2| \geq 1$. A direct consequence of Lemma 3.6 is that $f(R_1) + f(R_2) > f(T)$. Another consequence is that $p_j(R_i) < p_j(T)$ for every $i = 1, 2$ and every agent $j \in R_i$ (this consequence is derived by considering the sets R_i and $T - j$), thus $p(R_1) + p(R_2) < p(T)$. These consequences of Lemma 3.6 are employed to establish the following key property.

Lemma 3.7. *Let $v > 0$ be some payoff and let T be an optimal contract for v under the OR technology t. If $v < (1 + \hat{\sigma})p(T)$ for some positive real $\hat{\sigma} \leq 1/n$, then there exists some agent $j \in T$ such that $f(\{j\}) > (1 - \hat{\sigma})f(T)$.*

Proof. The assertion trivially holds if $|T| = 1$. Assume that $|T| \geq 2$ and consider some partition $T = R_1 \cup R_2$, $R_1 \cap R_2 = \emptyset$, such that $|R_1|, |R_2| \geq 1$. As T is optimal for v, we have $f(T)(v - p(T)) \geq f(R_i)(v - p(R_i))$, which can be rewritten as

$$(f(T) - f(R_i))v + f(R_i)p(R_i) \geq f(T)p(T).$$

Since $\frac{v}{p(T)} < 1 + \hat{\sigma}$, it follows that

$$(f(T) - f(R_i))p(T)(1 + \hat{\sigma}) + f(R_i)p(R_i) > f(T)p(T),$$

hence

$$\hat{\sigma}p(T)(f(T) - f(R_i)) + f(R_i)p(R_i) > f(T)p(T).$$

By summing the last inequality for $i = 1, 2$, we obtain

$$\hat{\sigma}p(T)(2f(T) - (f(R_1) + f(R_2))) + f(R_1)p(R_1) + f(R_2)p(R_2) > (f(R_1) + f(R_2))p(T).$$

Since $f(T) < f(R_1) + f(R_2)$, it follows that

$$\hat{\sigma}f(T)p(T) + f(R_1)p(R_1) + f(R_2)p(R_2) > (f(R_1) + f(R_2))p(T).$$

Suppose toward deriving contradiction that $p(R_i) \le (1-\hat{\sigma})p(T)$ for both $i = 1, 2$. Therefore

$$\hat{\sigma} f(T)p(T) + (f(R_1) + f(R_2))(1-\hat{\sigma})p(T) > (f(R_1) + f(R_2))p(T)$$

and $f(T) > f(R_1) + f(R_2)$, in contradiction to Lemma 3.6. We conclude that for every $j \in T$, either

$$p(\{j\}) > (1-\hat{\sigma})p(T) \quad \text{or} \quad p(T - \{j\}) > (1-\hat{\sigma})p(T).$$

Assume by way of contradiction that $p(T - \{j\}) > (1-\hat{\sigma})p(T)$ for every $j \in T$. Summing the last inequality for all $j \in T$ yields

$$\sum_{j\in T} p(T - \{j\}) > m(1-\hat{\sigma})p(T),$$

where $m = |T|$. Substituting for $p(T - \{j\})$, we get

$$\sum_{j\in T} \sum_{k\in T-\{j\}} \frac{c_k}{f(T-\{j\}) - f(T-\{j\}-\{k\})} > m(1-\hat{\sigma})p(T).$$

By Lemma 3.6, we obtain

$$\sum_{j\in T} \sum_{k\in T-\{j\}} \frac{c_k}{f(T) - f(T-\{k\})} > m(1-\hat{\sigma})p(T),$$

which can be rewritten as

$$(m-1)\sum_{k\in T} \frac{c_k}{f(T) - f(T-\{k\})} = (m-1)p(T) > m(1-\hat{\sigma})p(T).$$

Therefore $m - 1 > m(1-\hat{\sigma})$, in contradiction to $\hat{\sigma} \le 1/n \le 1/m$. It follows that there exists some agent $j \in T$ such that $p(\{j\}) > (1-\hat{\sigma})p(T)$.

As T is optimal for v, we have

$$f(T)(v - p(T)) > f(T - \{j\})(v - p(T - \{j\}))$$

and since $f(T) < f(\{j\}) + f(T - \{j\})$, it follows that

$$f(\{j\})(v - p(T)) > f(T - \{j\})(p(T) - p(T - \{j\})).$$

By the assumption that $v - p(T) < \hat{\sigma} p(T)$, we conclude that

$$\hat{\sigma} f(\{j\})p(T) > f(T - \{j\})(p(T) - p(T - \{j\})).$$

Lemma 3.6 implies that $p(\{j\}) < p(T) - p(T - \{j\})$, thus

$$\hat{\sigma} f(\{j\})p(T) > f(T - \{j\})p(\{j\}).$$

By the choice of j, we have

$$\hat{\sigma} f(\{j\}) > (1-\hat{\sigma})f(T - \{j\}).$$

Another application of Lemma 3.6 deduces that

$$\hat{\sigma} f(\{j\}) > (1-\hat{\sigma})(f(T) - f(\{j\})),$$

and hence $f(\{j\}) > (1-\hat{\sigma})f(T)$, which completes the proof. □

We are now ready to establish an FPTAS for the optimal contract problem on OR technologies. Let $\epsilon > 0$ be the performance parameter of the FPTAS. (Recall that for every $\epsilon > 0$, the FPTAS returns a solution which is at most $1+\epsilon$ times worse than the optimal solution in time $\mathrm{poly}(|t|, 1/\epsilon)$.) Subsequently, we assume that $\epsilon \le 1/n$ at the price of incurring an extra additive $\mathrm{poly}(|t|)$ term on the running time.

Fix $\sigma = \epsilon$ and $\hat{\sigma} = \frac{\epsilon}{1+\epsilon}$, and let $\mathcal{C}$ be the collection generated by Algorithm `Calibrate` when invoked on t with performance parameter $\rho = \frac{\sigma\hat{\sigma}}{1+2\hat{\sigma}}$. The FPTAS will consider the contracts in $\mathcal{C} \cup \{\{j\} \mid j \in N\}$, namely, the contracts in $\mathcal{C}$ and all the singleton contracts. Consider an arbitrary payoff $v > 0$ and let $T \subseteq N$ be an optimal contract for v. In order to establish Theorem 2, we have to prove that there exists a contract $S \in \mathcal{C} \cup \{\{j\} \mid j \in N\}$ such that $U_T(v)/U_S(v) \le 1+\epsilon$.

Assume first that $v < (1+\hat{\sigma})p(T)$. Since $\hat{\sigma} < \sigma \le 1/n$, we may apply Lemma 3.7 and conclude that there exists some agent $j \in N$ such that $f(\{j\}) > (1-\hat{\sigma})f(T)$. By Lemma 3.6, we have $p(\{j\}) \le p(T)$, hence $\frac{U_T(v)}{U_{\{j\}}(v)} = \frac{f(T)(v-p(T))}{f(\{j\})(v-p(\{j\}))} \le \frac{f(T)}{f(\{j\})} < \frac{1}{1-\hat{\sigma}}$. The assertion follows by the choice of $\hat{\sigma} = \frac{\epsilon}{1+\epsilon}$.

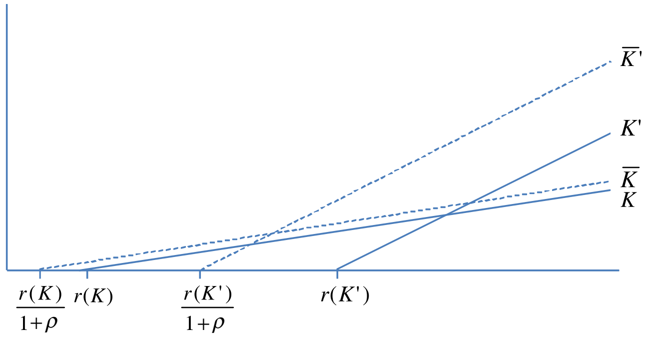

Fig. 5. Two lines in $\mathcal{K}$ and their shadows.

Now, assume that $v \geq (1+\hat{\sigma})p(T)$. Let S be the contract in $\mathcal{C}$ that serves as a $(1+\rho)$-estimation of T. Since $f(S) \geq f(T)/(1+\rho)$ and $p(S) \leq (1+\rho)p(T)$, we have

$$\begin{aligned}\frac{U_T(v)}{U_S(v)} &= \frac{f(T)(v-p(T))}{f(S)(v-p(S))} \\ &\leq (1+\rho)\frac{v-p(T)}{v-(1+\rho)p(T)} \\ &\leq (1+\rho)\frac{(1+\hat{\sigma})p(T)-p(T)}{(1+\hat{\sigma})p(T)-(1+\rho)p(T)} = \frac{(1+\rho)\hat{\sigma}}{\hat{\sigma}-\rho}.\end{aligned}$$

The requirement $\frac{(1+\rho)\hat{\sigma}}{\hat{\sigma}-\rho} \leq 1+\epsilon = 1+\sigma$ is guaranteed by the choice of the performance parameter $\rho = \frac{\sigma\hat{\sigma}}{1+2\hat{\sigma}}$ as $\frac{(1+\rho)\hat{\sigma}}{\hat{\sigma}-\rho} \leq 1+\sigma \iff \hat{\sigma}+\rho\hat{\sigma} \leq \hat{\sigma}+\sigma\hat{\sigma}-\rho-\rho\hat{\sigma} \iff \rho(1+2\hat{\sigma}) \leq \sigma\hat{\sigma}$.

3.3. Approximation for almost all relevant instances of SP technologies

Our goal in this section is to establish Theorem 3.

Theorem 3. *Given an SP technology t and two real parameters $0 < \epsilon, \hat{\epsilon} \leq 1$, there exists a scheme that on input payoff $v > 0$, either returns a $(1+\epsilon)$-approximate solution for v or outputs a failure message, in time* $\text{poly}(|t|, 1/\epsilon, 1/\hat{\epsilon})$. *Assuming that $F \subseteq \mathbb{R}_{>0}$ is the set of reals on which the scheme outputs a failure message, it is guaranteed that $\int_0^\infty 1_F(v)dv \leq \hat{\epsilon}v^*$, where 1_F is the characteristic function of F.*

In order to prove Theorem 3, we shall develop some (general purpose) insights regarding the geometric representation of combinatorial agency. Recall that the principal's expected utility for contract S is an increasing linear function of the payoff $v \in \mathbb{R}_{>0}$. In the scope of this section we will often represent it as such by considering a linear function (or line) L that assigns a real $L(v)$ to every real v. We denote the (positive) slope of L by $s(L)$ and the (unique) root of L by $r(L)$. (Under combinatorial agency terms, we have $s(L) = f(S)$ and $r(L) = p(S)$. Since $p(S) \geq 0$ for every contract S, our attention is restricted to lines with non-negative roots).

Consider some (finite) line collection $\mathcal{L}$. We denote the maximum real to which v is assigned under $\mathcal{L}$ by $\mathcal{L}(v) = \max\{L(v) \mid L \in \mathcal{L}\}$. A minimal subset $\mathcal{L}'$ of $\mathcal{L}$ that satisfies $\mathcal{L}'(v) = \mathcal{L}(v)$ for every $v \in \mathbb{R}$ is called an *orbit* of $\mathcal{L}$. Clearly, the upper envelope of $\mathcal{L}'$ is identical to that of $\mathcal{L}$. Moreover, a line $L \in \mathcal{L}$ that minimizes $r(L)$ and a line $L \in \mathcal{L}$ that maximizes $s(L)$ must be in the orbit. A typical real $v \in \mathbb{R}_{>0}$ admits a unique line $L \in \mathcal{L}$ that satisfies $L(v) = \mathcal{L}(v)$, but there is a finite number of reals $v \in \mathbb{R}_{>0}$ that admit two such lines, and we refer to these reals as the *transition numbers* of $\mathcal{L}$. (Actually, $\mathcal{L}$ has exactly $|\mathcal{L}'| - 1$ transition numbers.) The largest transition number of $\mathcal{L}$ is denoted by $v^*(\mathcal{L})$.

Consider some SP technology $t = \langle N, \{\gamma_i\}_{i=1}^n, \{\delta_i\}_{i=1}^n, \{c_i\}_{i=1}^n, \varphi\rangle$. Let $\mathcal{L}$ be the line collection which contains the line corresponding to U_S for every $S \subseteq N$. Fix $\Delta = \min\{\prod_{i\in N}\gamma_i, \prod_{i\in N}(1-\delta_i)\}$ and recall that Lemma 3.1 guarantees that $\Delta \leq s(L) \leq 1-\Delta$ for every line $L \in \mathcal{L}$.

Let $0 < \epsilon, \hat{\epsilon} \leq 1$ be the (real) parameters of Theorem 3. Fix $\sigma = \epsilon/3$ and $\hat{\sigma} = \hat{\epsilon}/(4n\ln(1/\Delta)+6)$. We invoke Algorithm `Calibrate` on t with performance parameter $\rho = \frac{\sigma\hat{\sigma}}{1+\sigma}$ to generate the contract collection $\mathcal{C}$ and append the contracts $\emptyset$ and N to $\mathcal{C}$ (if they are not already there). Consider the line collection $\mathcal{K}$ which contains the line corresponding to U_S for every $S \in \mathcal{C}$. Given some line $K \in \mathcal{K}$, we define its *shadow* line $\bar{K}$ by setting $s(\bar{K}) = s(K)$ and $r(\bar{K}) = r(K)/(1+\rho)$ (see Fig. 5). The *shadow* line collection is defined to be $\bar{\mathcal{K}} = \{\bar{K} \mid K \in \mathcal{K}\}$.

Corollary 3.5 guarantees that for every line $L \in \mathcal{L}$, there exists some line $K \in \mathcal{K}$ such that $s(K) \geq s(L)/(1+\rho)$ and $r(K) \leq r(L)(1+\rho)$. By the definition of $\bar{\mathcal{K}}$, we conclude that for every line $L \in \mathcal{L}$, there exists some line $\bar{K} \in \bar{\mathcal{K}}$ such that $s(\bar{K}) \geq s(L)/(1+\rho)$ and $r(\bar{K}) \leq r(L)$. It follows that $\bar{\mathcal{K}}(v) \geq \mathcal{L}(v)/(1+\rho)$ for every $v \in \mathbb{R}_{>0}$. Given some $v \in \mathbb{R}_{>0}$, if $\mathcal{K}(v) \geq \bar{\mathcal{K}}(v)/(1+\sigma)$, then $\mathcal{K}(v) \geq \mathcal{L}(v)/((1+\rho)(1+\sigma))$. Since $\rho \leq \sigma = \epsilon/3$, we have $(1+\rho)(1+\sigma) = 1+\rho+\sigma+\rho\sigma \leq 1+\epsilon$, which gives rise to the following corollary.

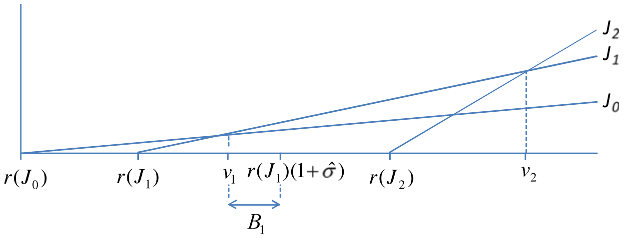

Fig. 6. The lines J_0, J_1, and J_2.

Corollary 3.8. *For every $v \in \mathbb{R}_{>0}$, if $\mathcal{K}$ provides a $(1+\sigma)$-approximation for v with respect to $\bar{\mathcal{K}}$, then $\mathcal{K}$ provides a $(1+\epsilon)$-approximation for v with respect to $\mathcal{L}$.*

On payoff $v \in \mathbb{R}_{>0}$ given as input, our algorithm works as follows. We first test whether $\mathcal{K}(v) \geq \bar{\mathcal{K}}(v)/(1+\sigma)$ (both $\mathcal{K}$ and $\bar{\mathcal{K}}$ are computed in time $\text{poly}(|t|, 1/\epsilon, 1/\hat{\epsilon})$ and are available whenever we wish to perform this test). If this test is positive, then we return the contract $S \in \mathcal{C}$ which corresponds to a line $K \in \mathcal{K}$ that realizes $\mathcal{K}(v)$. Corollary 3.8 guarantees that $U_S(v) \geq U_T(v)/(1+\epsilon)$ for any contract $T \subseteq N$ as promised by Theorem 3. Otherwise (if the above test is negative), we output a failure message. It remains to bound the fraction of payoffs v out of all relevant payoffs for which $\mathcal{K}$ does not provide a $(1+\sigma)$-approximation with respect to $\bar{\mathcal{K}}$. Note that there may be some non-relevant payoffs on which our algorithm outputs a failure message and we will account for them as well.

Let $\mathcal{J}$ be an orbit of $\bar{\mathcal{K}}$. Let $J_0, \ldots, J_{k+1}$ be the lines in $\mathcal{J}$ ordered such that $\Delta \leq s(J_0) < \cdots < s(J_{k+1}) < 1$ (recall that $s(\bar{K}) = s(K)$ for every $\bar{K} \in \bar{\mathcal{K}}$, hence the slopes in $\mathcal{J}$ are bounded between Δ and $1-\Delta$). It is easy to verify that $0 = r(J_0) < \cdots < r(J_{k+1})$. Fix $v_i = \inf\{v > 0 \mid J_i \text{ realizes } \mathcal{J}(v)\}$ for every $0 \leq i \leq k+1$ (this is well defined since $\mathcal{J}$ is an orbit). Refer to Fig. 6 for illustration. Clearly, $v_0 = 0$. For $1 \leq i \leq k+1$, v_i is actually the i^{th} transition real of $\mathcal{J}$ and it is easy to verify that

$$v_i = \frac{s(J_i)r(J_i) - s(J_{i-1})r(J_{i-1})}{s(J_i) - s(J_{i-1})}. \tag{13}$$

Consider some $v \in \mathbb{R}_{>0}$ and let J_i, $0 \leq i \leq k+1$, be a line that realizes $\mathcal{J}(v) = \bar{\mathcal{K}}(v)$. If v is not a transition real (which means that v is neither v_i nor v_{i+1}), then we say that J_i is *optimal* for v. If $i = 0$ (i.e., if $r(J_i) = 0$), then $\mathcal{K}(v) = \mathcal{J}(v)$ since by definition, there exists some line $K \in \mathcal{K}$ such that $s(K) = s(J_0)$ and $r(K) = r(J_0) = 0$ (K corresponds to the contract $\emptyset$). Now, consider some $v \in \mathbb{R}_{>0}$ and let J_i, $1 \leq i \leq k+1$, be an optimal line for v. We say that v is a *bad real* if $v \in (v_i, r(J_i)(1+\hat{\sigma}))$; otherwise, v is said to be a *good real*. The following proposition covers the good reals.

Proposition 3.9. *Consider some line J_i, $1 \leq i \leq k+1$, and let $v \geq r(J_i)(1+\hat{\sigma})$ be some real such that J_i is optimal for v. Then $\mathcal{K}(v) \geq \bar{\mathcal{K}}(v)/(1+\sigma)$.*

Proof. By definition, there exists some line $K \in \mathcal{K}$ such that $s(K) = s(J_i)$ and $r(K) = r(J_i)(1+\rho)$. We have

$$\begin{aligned}
\frac{\mathcal{J}(v)}{\mathcal{K}(v)} &\leq \frac{J_i(v)}{K(v)} \\
&= \frac{s(J_i)(v - r(J_i))}{s(K)(v - r(K))} \\
&= \frac{v - r(J_i)}{v - r(J_i)(1+\rho)} \\
&\leq \frac{r(J_i)(1+\hat{\sigma}) - r(J_i)}{r(J_i)(1+\hat{\sigma}) - r(J_i)(1+\rho)} \\
&= \frac{\hat{\sigma}}{\hat{\sigma} - \rho}.
\end{aligned}$$

By the choice of $\rho = \frac{\sigma\hat{\sigma}}{1+\sigma}$, we conclude that $\mathcal{J}(v)/\mathcal{K}(v) \leq (1+\sigma)\frac{\hat{\sigma}}{\hat{\sigma}(1+\sigma) - \sigma\hat{\sigma}} = 1+\sigma$, thus establishing the proposition. □

We define the *bad interval* exhibited by the line J_i to be $\mathcal{B}_i = (v_i, \min\{r(J_i)(1+\hat{\sigma}), v_{i+1}\})$ for every $1 \leq i \leq k$ and $(v_i, r(J_i)(1+\hat{\sigma}))$ for $i = k+1$. By Corollary 3.8 and Proposition 3.9, it is sufficient to bound (from above) the ratio $\Phi = \sum_{i=1}^{k+1} |\mathcal{B}_i|/v^*$. This is carried out in two stages: (1) bounding the ratio $\Phi_1 = \sum_{i=1}^{k} |\mathcal{B}_i|/v_{k+1} > \sum_{i=1}^{k} |\mathcal{B}_i|/v^*$ (this inequality holds as $v_{k+1} = v^*(\mathcal{J}) = v^*(\bar{\mathcal{K}}) < v^*(\mathcal{K}) \leq v^*(\mathcal{L}) = v^*$), and (2) bounding the ratio $\Phi_2 = |\mathcal{B}_{k+1}|/v^*$. Eventually, we will show that $\Phi_i \leq \hat{\epsilon}/2$ for $i = 1, 2$, thus establishing Theorem 3.

Bounding Φ_2 is easy: we have $\Phi_2 < \frac{|\mathcal{B}_{k+1}|}{r(J_{k+1})} < \frac{r(J_{k+1})(1+\hat{\sigma}) - r(J_{k+1})}{r(J_{k+1})} = \hat{\sigma}$. By the choice of $\hat{\sigma}$, it follows that $\Phi_2 \leq \hat{\epsilon}/2$. The bound on Φ_1 is more involved and depends on the geometric insight established in Lemma 3.10. In the scope of this lemma, we ignore the combinatorial agency interpretation of $J_0, \ldots, J_{k+1}$ and consider them merely as lines with slopes

$\Delta \leq s(J_0) < \cdots < s(J_{k+1}) < 1$ and roots $0 = r(J_0) < \cdots < r(J_{k+1})$. We write in short $s(J_i) = s_i$ and $r(J_i) = r_i$ for every $0 \leq i \leq k+1$. Observe that if $k = 0$ or if the bad intervals $\mathcal{B}_i$ are empty for all $1 \leq i \leq k$, then the bound $\Phi_1 \leq \hat{\epsilon}/2$ holds by definition. Therefore in what follows we assume that $k > 0$ and that there exist some $1 \leq i \leq k$ such that $|\mathcal{B}_i| > 0$.

Lemma 3.10. *The lines $J_0, \ldots, J_{k+1}$ must satisfy $\Phi_1 \leq 1 - e^{-\hat{\sigma}(2 \log k \ln(1/\Delta)+3)}$.*

Proof. For the sake of the analysis, we modify the line collection $\{J_0, \ldots, J_{k+1}\}$ in a manner that can only increase Φ_1. First, if the line J_i, $1 \leq i \leq k$, exhibits an empty bad interval, i.e., if $v_i \geq r_i(1+\hat{\sigma})$, then we remove it from the collection $\{J_0, \ldots, J_{k+1}\}$. This is repeated until every remaining line J_i, $1 \leq i \leq k$, exhibits a non-empty bad interval. In attempt to avoid cumbersome notation, we assume that the remaining lines are renamed $J_0, \ldots, J_{k+1}$ from scratch with s_i, r_i, and v_i defined as before, after this step. (The parameter k may have decreased as a result of the above step, but this causes the required bound $1 - e^{-\hat{\sigma}(2\log k \ln(1/\Delta)+3)}$ to decrease, hence it is sufficient to prove the assertion for a smaller k.) By removing those lines, we cannot decrease the ratio Φ_1 since good reals may have turned bad, but not vice versa. Note that v_{k+1} may have decreased due to the removals, which causes the ratio Φ_1 to increase.

Next, we fix $s_0, \ldots, s_{k+1}$ and modify $r_1, \ldots, r_{k+1}$ so that eventually we have $v_{i+1} \leq r_i(1+\hat{\sigma})$ for every $1 \leq i \leq k$. While doing so, we will ensure (in a manner specified below) that the ratio Φ_1 does not decrease. First, if $v_{k+1} > r_k(1+\hat{\sigma})$, then we fix $r_0, \ldots, r_k$ and multiply r_{k+1} by a factor of $1-d$ for sufficiently small positive d. Consequently, the bad intervals $\mathcal{B}_1, \ldots, \mathcal{B}_k$ remain intact and v_{k+1} is multiplied by a (positive) factor no greater than $1-d$ (see Eq. (13)), hence the ratio Φ_1 can only increase. We choose d so that the newly obtained v_{k+1} coincides with $r_k(1+\hat{\sigma})$.

The following step is repeated for $i = k-1, \ldots, 1$. Assume by induction that $v_{j+1} \leq r_j(1+\hat{\sigma})$ for every $i < j \leq k$. If $v_{i+1} > r_i(1+\hat{\sigma})$, then we fix $r_0, \ldots, r_i$ and multiply all r_js, $i < j \leq k+1$, by a factor of $1-d$ for sufficiently small positive d. Consequently, we get (i) the bad intervals $\mathcal{B}_1, \ldots, \mathcal{B}_i$ remain intact, (ii) the size of the bad interval $\mathcal{B}_j$ is multiplied by a factor of $1-d$ for every $i < j \leq k$ (see Eq. (13)), (iii) v_{k+1} is multiplied by a factor of $1-d$, and (iv) the assumption $v_{j+1} \leq r_j(1+\hat{\sigma})$ for every $i < j \leq k$ is not violated. By (i) and (ii), we conclude that the numerator in the ratio Φ_1 is multiplied by a factor no smaller than $1-d$, thus, combined with (iii), the ratio Φ_1 can only increase. Once again, we choose d so that the newly obtained v_{i+1} coincides with $r_i(1+\hat{\sigma})$.

So, in what follows, we may assume without loss of generality that $r_i < v_i < v_{i+1} \leq r_i(1+\hat{\sigma})$ for every $1 \leq i \leq k$, which means that $\mathcal{B}_i = (v_i, v_{i+1})$ for every $1 \leq i \leq k$ and it remains to bound the ratio $\Phi_1 = (v_{k+1} - v_1)/v_{k+1}$. Instead, we will bound the larger ratio $(v_{k+1} - r_1)/v_{k+1}$.

By Eq. (13), the assumption $v_{i+1} \leq r_i(1+\hat{\sigma})$ implies that $r_i \geq r_{i+1}\frac{s_{i+1}}{s_{i+1}(1+\hat{\sigma})-\hat{\sigma}s_i}$. Consider some integer $0 \leq q \leq \log k$. If $s_i \geq (1-2^{-q})s_{i+1}$, then $r_i \geq r_{i+1}\frac{1}{1+\hat{\sigma}/2^q}$. How many indices $1 \leq i \leq k$ can satisfy the inequality $s_i < (1-2^{-q})s_{i+1}$? Since $s_1 > s_0 \geq \Delta$ and $s_{k+1} < 1$, it follows that if there exists m such indices i, then $(1-2^{-q})^m > \Delta$. Thus $e^{-m/2^q} > \Delta$ and $m < 2^q \ln(1/\Delta)$.

We shall partition the indices $1, \ldots, k$ into $\lfloor \log k \rfloor + 1$ *categories*, denoted by $C_1, \ldots, C_{\lfloor \log k \rfloor+1}$. For every $1 \leq q \leq \lfloor \log k \rfloor$, the category C_q consists of all indices $1 \leq i \leq k$ such that $(1-2^{1-q})s_{i+1} \leq s_i < (1-2^{-q})s_{i+1}$. Recall that $|C_q| < 2^q \ln(1/\Delta)$. The category $C_{\lfloor \log k \rfloor+1}$ consists of all indices $1 \leq i \leq k$ such that $(1-2^{-\lfloor \log k \rfloor})s_{i+1} \leq s_i$. Clearly, $|C_{\lfloor \log k \rfloor+1}| \leq k$. Therefore

$$\begin{aligned} r_1 &\geq \prod_{q=1}^{\lfloor \log k \rfloor+1} \left(\frac{1}{1+\hat{\sigma}/2^{q-1}}\right)^{|C_q|} \cdot r_{k+1} \\ &> \prod_{q=1}^{\lfloor \log k \rfloor} \left(\frac{1}{1+\hat{\sigma}/2^{q-1}}\right)^{2^q \ln(1/\Delta)} \cdot \left(\frac{1}{1+\hat{\sigma}/2^{\lfloor \log k \rfloor}}\right)^{k} \cdot r_{k+1} \\ &> \prod_{q=1}^{\lfloor \log k \rfloor} e^{-2\hat{\sigma}\ln(1/\Delta)} \cdot e^{-2\hat{\sigma}} \cdot r_{k+1} \\ &\geq e^{-2\hat{\sigma}(\log k \ln(1/\Delta)+1)} \cdot r_{k+1}. \end{aligned}$$

Since $v_{k+1} \leq r_k(1+\hat{\sigma}) < r_{k+1}(1+\hat{\sigma}) < e^{\hat{\sigma}} r_{k+1}$, we have $r_1 > e^{-\hat{\sigma}(2\log k \ln(1/\Delta)+3)} \cdot v_{k+1}$. Therefore $\frac{v_{k+1}-r_1}{v_{k+1}} = 1 - r_1/v_{k+1} < 1 - e^{-\hat{\sigma}(2\log k \ln(1/\Delta)+3)}$.

The proof is completed by showing that $1 - e^{-\hat{\sigma}(2\log k \ln(1/\Delta)+3)} \leq \hat{\epsilon}/2$. This holds by the choice of $\hat{\sigma}$ since

$$\begin{aligned} &1 - \hat{\epsilon}/2 \leq e^{-\hat{\sigma}(2\log k \ln(1/\Delta)+3)} \\ &\Longleftarrow e^{-\hat{\epsilon}/2} \leq e^{-\hat{\sigma}(2\log k \ln(1/\Delta)+3)} \\ &\Longleftrightarrow \hat{\sigma} \leq \hat{\epsilon}/(4\log k \ln(1/\Delta)+6) \end{aligned}$$

and since $\log k \leq n$. $\square$

3.4. A note on arbitrary technologies

We conclude the paper with a note on arbitrary technologies. The following theorem shows that every technology admits a collection of polynomial size that gives a $(1+\epsilon)$-approximation to the optimal contract.

Theorem 4. *For every technology $t = \langle N, \{\gamma_i\}_{i=1}^n, \{\delta_i\}_{i=1}^n, \{c_i\}_{i=1}^n, \varphi \rangle$ and for any $\epsilon > 0$, the orbit of t admits a $(1+\epsilon)$-approximation of size* $\mathrm{poly}(|t|, 1/\epsilon)$.

Proof. Consider an arbitrary technology $t = \langle N, \{\gamma_i\}_{i=1}^n, \{\delta_i\}_{i=1}^n, \{c_i\}_{i=1}^n, \varphi \rangle$ and some $\epsilon > 0$. The contract collection $\mathcal{C}$ is constructed in a single stage of Algorithm `Calibrate`: we first calibrate all contracts in 2^N into a scale of precision $1+\epsilon$ and then remove from each interval all contracts excluding the one with minimum payment (under t). More formally, the collection $\mathcal{C}$ contains at most one contract S that is calibrated to the interval $\mathcal{I}$, in this case $p(S) \leq p(S')$ for every contract $S' \subseteq N$ such that S' is calibrated to $\mathcal{I}$. Following the line of arguments presented earlier in this section, we show that $|\mathcal{C}| = O\left(\frac{1}{\epsilon} \log \frac{1}{\Delta}\right)$. Moreover, if an arbitrary contract $T \subseteq N$ is not in $\mathcal{C}$, then it was compensated by some contract $S \in \mathcal{C}$ such that S and T are calibrated to the same interval. Therefore $f(S) \geq \frac{f(T)}{1+\epsilon}$ and since $p(S) \leq p(T)$, it follows that $\frac{U_T(v)}{U_S(v)} \leq 1+\epsilon$ for every payoff $v > 0$. □

4. Conclusions

The hidden action problem lies at the heart of economic theory and has been recently studied from an algorithmic perspective. In this article, we continue the study initiated by Babaioff et al. [1] of the computational complexity of optimal team incentives under hidden actions. Our contribution focuses on the OR technology, whose computational complexity was raised as an open question in [1]. The importance of our results comes from the observation that OR technologies are very common in real life since agents' actions quite often serve as substitutions for each other. Indeed, the OR technology is one of the most fundamental and common interrelations between agents' tasks. We establish the NP-hardness of the problem of computing an optimal contract in an OR technology, and we also show that there exist OR technologies with exponentially large orbits (thus disproving a conjecture of [1]). On the positive side, we devise an FPTAS for OR technologies.

In addition, we consider the more general family of *series–parallel* (SP) technologies, which are constructed inductively from AND and OR technologies. For SP technologies, we establish a scheme that provides a $(1+\epsilon)$-approximation for all but an $\hat{\epsilon}$-fraction of the relevant instances in time polynomial in the size of the technology and in the reciprocals of ϵ and $\hat{\epsilon}$. It remains as an open problem whether there exists an approximation scheme for SP technologies. This article makes a significant step in understanding the computational complexity of the combinatorial agency model, which is an example of the important interaction between game theory, economic theory, and computer science.

Acknowledgments

The authors are indebted to Noam Nisan for his help on various parts of the paper. The second author's work was partially supported by the Israel Science Foundation (grant number 1219/09), by the Leon Recanati Fund of the Jerusalem school of business administration and by the Google Inter-university center for Electronic Markets and Auctions.

References

[1] Moshe Babaioff, Michal Feldman, Noam Nisan, Combinatorial agency, in: ACM EC'06, 2006.
[2] Moshe Babaioff, Michal Feldman, Noam Nisan, Mixed strategies in combinatorial agency, in: 2nd International Workshop on Internet and Network Economics, WINE, 2006.
[3] Moshe Babaioff, Michal Feldman, Noam Nisan, Free-riding and free-labor in combinatorial agency, in: Symposium on Algorithmic Game Theory, SAGT, 2009.
[4] Moshe Babaioff, Michal Feldman, Noam Nisan, Eyal Winter, Combinatorial agency, Journal of Economic Theory 147 (3) (2012) 999–1034.
[5] Felipe Balmaceda, Santiago R. Balseiro, Jose R. Correa, Nicolas E. Stier-Moses, Cost of moral hazard and limited liability in the principal-agent problem, in: Workshop on Internet and Network Economics, WINE, 2010.
[6] Raphael Eidenbenz, Stefan Schmid, Combinatorial agency with audits, in: IEEE International Conference on Game Theory for Networks, GameNets, 2009.
[7] Yuval Emek, Iftach Heitner, Combinatorial agency: the observable-action model. 2006. A manuscript.
[8] Joan Feigenbaum, Scott Shenker, Distributed algorithmic mechanism design: recent results and future directions, in: Proceedings of the International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications, 2002.
[9] Michal Feldman, John Chuang, Ion Stoica, Scott Shenker, Hidden-action in multi-hop routing, in: ACM Conference on Electronic Commerce, 2005, pp. 117–126.
[10] M.R. Garey, D.S. Johnson, Computers and Intractability: A Guide to the Theory of NP-Completeness, W. H. Freeman, New York, 1979.
[11] Bengt Holmstrom, Moral hazard in teams, Bell Journal of Economics 13 (1982) 324–340.
[12] Patrick Legeros, Steven Matthews, Efficient and nearly-efficient partnerships, Review of Economic Studies 60 (3) (1993) 599–611.
[13] Shaili Jain, David Parkes, Combinatorial agency of threshold functions, in: Symposium on Algorithmic Game Theory, SAGT, 2011.
[14] Dilip Mookherjee, Optimal incentive schemes with many agents, Review of Economic Studies 51 (3) (1984) 433–446.
[15] Noam Nisan, Algorithms for selfish agents, in: Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science, 1999, pp. 1–15.
[16] Noam Nisan, Amir Ronen, Algorithmic mechanism design, in: Proceedings of the 31st Symposium on Theory of Computing, 1999.
[17] Noam Nisan, Tim Roughgarden, Eva Tardos, Vijay V. Vazirani, Algorithmic Game Theory, Cambridge University Press, 2007.
[18] Christos Papadimitriou, Algorithms, games, and the internet, in: Proceedings of 33rd STOC, 2001, pp. 749–753.
[19] E. Rasmusen, Moral hazard in risk-averse teams, Rand Journal of Economics 18 (1987) 428–435.
[20] Ronald Strausz, Moral hazard in sequential teams, in: Departmental Working Paper. Free University of Berlin, 1996.
[21] Eyal Winter, Incentives and Discrimination, American Economic Review 94 (2004) 764–773.