

Mixed Strategies in Combinatorial Agency

[Extended Abstract]*

Moshe Babaioff **, Michal Feldman ***, and Noam Nisan†

Abstract. We study a setting where a principal needs to motivate a team of agents whose combination of hidden efforts stochastically determines an outcome. In a companion paper we devise and study a basic “combinatorial agency” model for this setting, where the principal is restricted to inducing a pure Nash equilibrium. Here, we show that the principal may possibly gain from inducing a mixed equilibrium, but this gain can be bounded for various families of technologies (in particular if a technology has symmetric combinatorial structure). In addition, we present a sufficient condition under which mixed strategies yield no gain to the principal.

1 Introduction

1.1 Background: Combinatorial Agency

The well studied principal-agent problem deals with how a “principal” can motivate a rational “agent” to exert costly effort towards the welfare of the principal. The difficulty in this model is that the agent’s action (i.e. whether he exerts effort or not) is invisible to the principal and only the final outcome, which is probabilistic and also influenced by other factors, is visible. “Invisible” here is meant in a wide sense that includes “not precisely measurable”, “costly to determine”, or “non-contractible” (meaning that it can not be upheld in “a court of law”). This problem is well studied in many contexts in classical economic theory and we refer the readers to introductory texts on economic theory such as [8] Chapter 14. The solution is based on the observation that a properly designed contract, in which the payments are contingent upon the final outcome, can influence a rational agent to exert the required effort.

In [2] we initiated a general study of handling *combinations* of agents rather than a single agent. While much work was previously done on motivating teams of agents [6, 10, 7, 4], our emphasis is on dealing with the complex combinatorial structure of dependencies between agents’ actions. In the general case, each

* A full version of this paper is available at the authors’ web sites. All proofs and some additional results appear in the full version.

** UC Berkeley School of Information. *E-mail*: moshe@sims.berkeley.edu.

*** School of Computer Science, Hebrew University of Jerusalem. *E-mail*: mfeldman@cs.huji.ac.il.

† School of Computer Science, Hebrew University of Jerusalem. *E-mail*: noam@cs.huji.ac.il.

combination of efforts exerted by the n different agents may result in a different expected gain for the principal. The general question asks, given an exact specification of the expected utility of the principal for each combination of agents' actions, which conditional payments should the principal offer to which agents as to maximize his net utility? We view this problem of hidden actions in computational settings as a complementary problem to the problem of hidden information that is the heart of the field of Algorithmic Mechanism Design [9]. An example that was discussed in [5] is Quality of Service routing in a network: every intermediate link or router may exert a different amount of "effort" (priority, bandwidth, ...) when attempting to forward a packet of information. While the final outcome of whether a packet reached its destination is clearly visible, it is rarely feasible to monitor the exact amount of effort exerted by each intermediate link – how can we ensure that they really do exert the appropriate amount of effort? For example, in Internet routing, IP routers may delay or drop packets, and in mobile ad hoc networks, devices may strategically drop packets to conserve their constrained energy resources.

In the general model presented in [2], each of n agents has a set of possible *actions*, the combination of actions by the players results in some *outcome*, where this happens probabilistically. The main part of the specification of a problem in this model is a function ("the technology") that specifies this distribution for each n -tuple of agents' actions. Additionally, the problem specifies the principal's utility for each possible outcome, and for each agent, the agent's cost for each possible action. The principal motivates the agents by offering to each of them a *contract* that specifies a payment for each possible outcome of the whole project. Key here is that the actions of the players are non-observable ("hidden-actions") and thus the contract cannot make the payments directly contingent on the actions of the players, but rather only on the outcome of the whole project.

Given a set of contracts, each agent optimizes his own utility; i.e., chooses the action that maximizes his expected payment minus the cost of the action. Since the outcome depends on the actions of all players together, the agents are put in a game here and are assumed to reach a Nash Equilibrium (NE). The principal's problem is that of designing the *optimal contract*: i.e. the vector of contracts to the different agents that induce an equilibrium that will optimize his expected utility from the outcome minus his expected total payment. The main difficulty is that of determining the required Nash equilibrium point.

Our interest in this paper (as in [2]), is focused on the binary case: each agent has only two possible actions "exert effort" and "shirk" and there are only two possible outcomes "success" and "failure". Our motivating examples comes from the following more restricted and concrete "structured" subclass of problem instances: Every agent i performs a subtask which succeeds with a low probability γ_i if the agent does not exert effort and with a higher probability $\delta_i > \gamma_i$, if the agent does exert effort. The whole project succeeds as a deterministic Boolean function of the success of the subtasks.

1.2 This Paper: Mixed Equilibria

In [2] we studied the notion of Nash-equilibrium in pure strategies: we did not allow the principal to attempt inducing an equilibrium where agents have mixed strategies over their actions. In the observable-actions case (where the principal can condition the payments on the agents' individual actions) the restriction to pure strategies is without loss of generality: mixed actions can never help since they simply provide a convex combination of what would be obtained by pure actions. Yet, surprisingly, we show this is not the case for the hidden-actions case which we are studying: in some cases, a Mixed-Nash equilibrium can provide better expected utility to the principal than what he can obtain by equilibrium in pure strategies. In particular, this already happens for the "OR" function with two players, with a certain (quite restricted) range of parameters (see Section 3).

Our main goal is to quantify the principal's gain from inducing mixed equilibrium, rather than pure. To do that, we analyze the worst ratio (over all principal's values) between the principal's optimal utility with mixed equilibrium, and his optimal utility with pure equilibrium. We term this ratio "the price of purity" (POP) of the instance under study. We prove that for a class of instances, those with "increasing returns to scale", which contains in particular the *AND* function, the price of purity is trivial (i.e., $POP = 1$). Moreover, we show that for any other Boolean function, there is an assignment of the parameters (agents' individual success probabilities) for which the obtained structured technology has non trivial POP (i.e., $POP > 1$). (Section 4).

While the price of purity may be strictly greater than 1, we obtain quite a large number of results bounding this ratio (Section 5). These bounds range from very weak ones (e.g., $POP \leq n$ for any anonymous or DRS technology) to better ones for restricted cases (e.g., $POP \leq 1.154\dots$ for a family of anonymous OR technologies, and $POP \leq 2$ for any technology with 2 agents). We conjecture that there exists a universal constant C that bounds the POP for any technology, thus shrinking the large gaps between our conjecture and the obtained bounds is the main open problem of this paper.

Conjecture 1. There exists a constant $C > 1$, such that for any technology t , $POP(t) \leq C$.

A more extreme form of the conjecture states that a non-anonymous OR technology with 2 agents is the most extreme case and yields the highest possible POP for any structured technology.

Additionally, we study some other properties of mixed equilibrium. We show that mixed Nash equilibria are more delicate than pure ones. In particular, we show that unlike the pure case, in which the optimal contract is also a "strong equilibrium" [1] (i.e., resilient to deviations by coalitions), an optimal mixed contract (in which at least two agents truly mix) never satisfies the requirements of a strong equilibrium (Section 6).

Finally, we study the computational hardness of the optimal mixed Nash equilibrium, and show that the hardness results from the pure case hold for the mixed case as well (Section 7).

2 Model and Preliminaries

We focus on the simple “binary action, binary outcome” scenario where each agent has two possible actions (“exert effort” or “shirk”) and there are two possible outcomes (“failure”, “success”). We begin by presenting the model with pure actions (which is a generalization of [11]), and then move to the mixed case. A principal employs a set of agents N of size n . Each agent $i \in N$ has a set of two possible actions $A_i = \{0, 1\}$ (binary action), the low effort action (0) has a cost of 0 ($c_i(0) = 0$), while the high effort action (1) has a cost of $c_i > 0$ ($c_i(1) = c_i$). The played profile of actions determine, in a probabilistic way, a “contractible” outcome, $o \in \{0, 1\}$, where the outcomes 0 and 1 denote project failure and success, respectively (binary-outcome). The outcome is determined according to a success function $t : A_1 \times \dots \times A_n \rightarrow [0, 1]$, where $t(a_1, \dots, a_n)$ denotes the probability of project success where players play with the action profile $a = (a_1, \dots, a_n) \in A_1 \times \dots \times A_n = A$. We use the notation (t, \mathbf{c}) to denote a technology (a success function and a vector of costs, one for each agent).

The principal’s value of a successful project is given by a scalar $v > 0$, where he gains no value from a project failure. In this hidden-actions model the actions of the players are invisible, but the final outcome is visible to him and to others, and he may design enforceable contracts based on this outcome. We assume that the principal can pay the agents but not fine them (known as the *limited liability* constraint). The contract to agent i is thus given by a scalar value $p_i \geq 0$ that denotes the payment that i gets in case of project success. If the project fails, the agent gets no money (this is in contrast to the “observable-actions” model in which payment to an agent can be contingent on his action).

Given this setting, the agents have been put in a game, where the utility of agent i under the profile of actions $a = (a_1, \dots, a_n) \in A$ is given by $u_i(a) = p_i \cdot t(a) - c_i(a_i)$. As usual, we denote by $a_{-i} \in A_{-i}$ the $(n-1)$ -dimensional vector of the actions of all agents excluding agent i . i.e., $a_{-i} = (a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n)$. The agents will be assumed to reach Nash equilibrium, if such an equilibrium exists. The principal’s problem (which is our problem in this paper) is how to design the contracts p_i as to maximize his own expected utility $u(a, v) = t(a) \cdot (v - \sum_{i \in N} p_i)$, where the actions a_1, \dots, a_n are at Nash-equilibrium. In the case of multiple Nash equilibria, in our model we let the principal choose the desired one, and “suggest” it to the agents, thus focusing on the “best” Nash equilibrium.

As we wish to concentrate on motivating agents, rather than on the coordination between agents, we assume that more effort by an agent always leads to a better probability of success. Formally, $\forall i \in N, \forall a_{-i} \in A_{-i}$ we have that $t(1, a_{-i}) > t(0, a_{-i})$. We also assume that $t(a) > 0$ for any $a \in A$.

We next consider the extended game in which an agent can mix between exerting effort and shirking (randomize over the two possible pure actions). Let q_i denote the probability that agent i exerts effort, and let q_{-i} denote the $(n-1)$ -dimensional vector of investment probabilities of all agents except for agent i . We can extend the definition of the success function t to the range of mixed strategies, by taking the expectation.

$$t(q_1, \dots, q_n) = \sum_{a \in \{0,1\}^n} \left(\prod_{i=1}^n q_i^{a_i} \cdot (1 - q_i)^{(1-a_i)} \right) t(a_1, \dots, a_n)$$

Note that for any agent i and any (q_i, q_{-i}) it holds that $t(q_i, q_{-i}) = q_i \cdot t(1, q_{-i}) + (1 - q_i) \cdot t(0, q_{-i})$. A mixed equilibrium profile in which at least one agent mixes with probability $p_i \in [0, 1]$ is called a *non-degenerate* mixed equilibrium.

In pure strategies, the marginal contribution of agent i , given $a_{-i} \in A_{-i}$, is defined to be: $\Delta_i(a_{-i}) = t(1, a_{-i}) - t(0, a_{-i})$. For the mixed case we define the marginal contribution of agent i , given q_{-i} to be: $\Delta_i(q_{-i}) = t(1, q_{-i}) - t(0, q_{-i})$. Since t is monotone, Δ_i is a positive function.

We next characterize what payment can result in an agent mixing between exerting effort and shirking.

Claim. Agent i 's best response is to mix between exerting effort and shirking with probability $q_i \in (0, 1)$ only if he is indifferent between $a_i = 1$ and $a_i = 0$. Thus, given a profile of strategies q_{-i} , agent i mixes only if:

$$p_i = \frac{c_i}{\Delta_i(q_{-i})} = \frac{c_i}{t(1, q_{-i}) - t(0, q_{-i})}$$

which is the payment that makes him indifferent between exerting effort and shirking. The expected utility of agent i , who exerts effort with probability q_i is: $u_i(q) = c_i \cdot \left(\frac{t(q)}{\Delta_i(q_{-i})} - q_i \right)$.

A profile of mixed strategies $q = (q_1, \dots, q_n)$ is a *Mixed Nash equilibrium* if for any agent i , q_i is agent i 's best response, given q_{-i} .

The principal's expected utility under the mixed Nash profile q is given by $u(q, v) = (v - P) \cdot t(q)$, where P is the total payment in case of success, given by $P = \sum_{i|q_i > 0} \frac{c_i}{\Delta_i(q_{-i})}$. An *optimal mixed contract* for the principal is an equilibrium mixed strategy profile $q^*(v)$ that maximizes the principal's utility at the value v . In [2] we show a similar characterization of optimal pure contract $a \in A$. An agent that exerts effort is paid $\frac{c_i}{\Delta_i(a_{-i})}$, and the utilities are the same as the above, when given the pure profile. In the pure Nash case, given a value v , an *optimal pure contract* for the principal is a set of agents $S^*(v)$ that exert effort in equilibrium, and this set maximizes the principal's utility at the value v .

A simple but crucial observation, generalizing a similar one in [2] for the pure Nash case, shows that the optimal mixed contract exhibits some monotonicity properties in the value.

Lemma 1. (Monotonicity lemma): *For any technology (t, c) the expected utility of the principal at the optimal mixed contract, the success probability of the optimal mixed contract, and the expected payment of the optimal mixed contract, are all monotonically non-decreasing with the value.*

The proof shows that the same monotonicity holds in the observable-actions case as well. Additionally, the lemma holds in more general settings, where each

agent has an arbitrary action set (not restricted to the binary-actions model considered here).

We wish to quantify the gain by inducing mixed Nash equilibrium, over inducing pure Nash. We define the *price of purity* as the worse ratio (over v) between the maximum utilities that are obtained in mixed and pure strategies.

Definition 1. *The price of purity $POP(t, \mathbf{c})$ of a technology (t, \mathbf{c}) is defined as the worse ratio, over v , between the principal's optimal utility in the mixed case and his optimal utility in the pure case. Formally,*

$$POP(t, \mathbf{c}) = \text{Sup}_{v>0} \frac{t(q^*(v)) \left(v - \sum_{i|q_i^*(v)>0} \frac{c_i}{\Delta_i(q_i^*(v))} \right)}{t(S^*(v)) \left(v - \sum_{i \in S^*(v)} \frac{c_i}{\Delta_i(a_{-i})} \right)}$$

where $S^*(v)$ denotes an optimal pure contract and $q^*(v)$ denotes an optimal mixed contract, for the value v .

The price of purity is at least 1, and may be greater than 1, as we later show. Additionally, it is obtained at some value that is a transition point of the pure case (a point in which the principal is indifferent between two optimal pure contracts).

Lemma 2. *For any technology (t, \mathbf{c}) , the price of purity is obtained at a finite v that is a transition point between two optimal pure contracts.*

2.1 Structured Technology Functions

In order to be more concrete, we next present technology functions whose structure can be described easily as being derived from independent agent tasks – we call these *structured technology functions*. This subclass gives us some natural examples of technology functions, and also provides a succinct and natural way to represent technology success functions.

In a structured technology function, each individual succeeds or fails in his own “task” independently. The project’s success or failure deterministically depends, maybe in a complex way, on the set of successful sub-tasks. Thus we will assume a monotone Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ which indicates whether the project succeeds as a function of the success of the n agents’ tasks.

A structured technology function t is defined by $t(a_1, \dots, a_n)$ being the probability that $f(x_1, \dots, x_n) = 1$ where the bits x_1, \dots, x_n are chosen according to the following distribution: if $a_i = 0$ then $x_i = 1$ with probability $\gamma_i \in [0, 1)$ (and $x_i = 0$ with probability $1 - \gamma_i$); otherwise, i.e. if $a_i = 1$, then $x_i = 1$ with probability $\delta_i > \gamma_i$ (and $x_i = 0$ with probability $1 - \delta_i$). Thus, a structured technology is defined by n , f and the parameters $\{\delta_i, \gamma_i\}_{i \in N}$.

Let us consider two simple structured technology functions, “AND” and “OR”. First consider the “AND” technology: $f(x_1, \dots, x_n)$ is the logical conjunction of x_i ($f(x) = \bigwedge_{i \in N} x_i$). Thus the project succeeds only if all agents

succeed in their tasks. For this technology, the probability of success is the product of the individual success probabilities. Agent i succeeds with probability $\delta_i^{a_i} \cdot \gamma_i^{1-a_i}$, thus $t(a) = \prod_{i \in N} \delta_i^{a_i} \cdot \gamma_i^{1-a_i}$.

Next, consider the "OR" technology: $f(x_1, \dots, x_n)$ is the logical disjunction of x_i ($f(x) = \bigvee_{i \in N} x_i$). Thus the project succeeds if at least one of the agents succeed in their tasks. For this technology, the probability of success is 1 minus the probability that all of them fail. Agent i fails with probability $(1 - \delta_i)^{a_i} \cdot (1 - \gamma_i)^{1-a_i}$, thus $t(a) = 1 - \prod_{i \in N} (1 - \delta_i)^{a_i} \cdot (1 - \gamma_i)^{1-a_i}$.

These are just two simple examples. One can consider other more interesting examples as the Majority function (the project succeed if the majority of the agents are successful), or the OR-Of-ANDs technology, which is a disjunction over conjunctions (several teams, the project succeed if all the agents in any one of the teams are successful). For additional examples see [2].

A success function t is called *anonymous* if it is symmetric with respect to the players. I.e. $t(a_1, \dots, a_n)$ depends only on $\sum_i a_i$. For example, in an anonymous OR technology there are parameters $1 > \delta > \gamma > 0$ such that each agent i succeed with probability γ with no effort, and with probability $\delta > \gamma$ with effort. If m agents exert effort, the success probability is $1 - (1 - \delta)^m \cdot (1 - \gamma)^{n-m}$.

A technology has *identical costs* if there exists a c such that for any agent i , $c_i = c$. As in the case of identical costs the POP is independent of c , we use $POP(t)$ to denote the POP for technology t with identical costs. We abuse notation and denote a technology with identical costs by its success function t . Throughout the paper, unless explicitly stated otherwise, we assume identical costs. A technology t with identical costs is *anonymous* if t is anonymous.

3 Example: Mixed Nash Outperforms Pure Nash!

If the actions are observable (henceforth, the observable-actions case), then an agent that exerts effort is paid exactly his cost, and the principal's utility equals the social welfare. In this case, the social welfare in mixed strategies is a convex combination of the social welfare in pure strategies; thus, it is clear that the optimal utility is always obtained in pure strategies. However, surprisingly enough, in the hidden-actions case, the principal might gain higher utility when mixed strategies are allowed. This is demonstrated in the following example:

Example 1. Consider an anonymous OR technology with two agents, where $c = 1$, $\gamma = \gamma_1 = \gamma_2 = 1 - \delta_1 = 1 - \delta_2 = 0.09$ and $v = 348$. The mixed strategies $q_1 = q_2 = 0.92$ achieve a utility of 324.27, while the optimal contract with pure strategies is obtained when both agents exert effort and achieves a utility of 318.3. This implies that by moving from pure strategies to mix strategies, one gains at least $324.27/318.3 > 1.0187$ factor improvement (which is approximately 1.8%).

A worse ratio exists for the more general case (in which it does not necessarily hold that $\delta = 1 - \gamma$) of $\gamma = 0.0001$, $\delta = 0.9$ and $v = 233$. For this case we get that the optimal pure contract is with one agent, gives utility of 208.7, while the

mixed contract $q_1 = q_2 = 0.92$ gives utility of 213.569, and the ratio is at least 1.0233 (approximately 2.3%).

To complete the example, Diagram 1 presents the optimal contract for *OR* of 2 agents, as a function of γ (when $\delta = 1 - \gamma$) and v . It shows that for some parameters of γ and v , the optimal contract is obtained when both agents exert effort with equal probabilities.




Fig. 1. Optimal mixed contracts in *OR* technologies with 2 agents. The white area corresponds to both agents exert effort with the same non-trivial probability, q . For any fixed γ , q increases in v .

The following lemma shows that optimal mixed contracts in any anonymous *OR* technology have this specific structure. That is, all agents that do not shirk, mix with exactly the same probability.

Lemma 3. *For any anonymous OR technology (any $\delta > \gamma, c, n$) and value v , either the optimal mixed contract is a pure contract, or, in the optimal mixed contract $k \in \{2, \dots, n\}$ agents exert effort with equal probabilities $q_1 = \dots = q_k \in (0, 1)$, and the rest of the agents exert no effort.*

4 When is Pure Nash Good Enough?

Next, we identify a class of technologies for which the price of purity is 1; that is, the principal cannot improve his utility by moving from pure Nash equilibrium to mixed Nash equilibrium. These are technologies for which the marginal contribution of any agent is non-decreasing in the effort of the other agents. Formally, for two pure action profiles $a, b \in A$ we denote $b \succeq a$ if for all j , $b_j \succeq_j a_j$ (effort b_j is at least as high as the effort a_j).

Definition 2. *A technology success function t exhibits (weakly) increasing returns to scale (IRS) if for every i , and every pure profiles $b \succeq a$*

$$t(b_i, b_{-i}) - t(a_i, b_{-i}) \geq t(b_i, a_{-i}) - t(a_i, a_{-i})$$

Any AND technology exhibits IRS [11, 2]. For IRS technologies we show that $POP=1$.

Theorem 1. *Assume that t exhibits increasing returns to scale (IRS). For any cost vector \mathbf{c} , $POP(t, \mathbf{c}) = 1$. Moreover, a non-degenerate mixed contract is never optimal.*

We show that *AND* (on some subset of bits) is the only function such that any structured technology based on this function exhibits IRS, that is, this is the only function such that for any choices of parameters (any n and any $\{\delta_i, \gamma_i\}_{i \in N}$), the structured technology exhibits IRS. For any other Boolean function, there is an assignment for the parameters such that the created structured technology is essentially OR over 2 inputs (see lemma in full version), thus it has non-trivial POP (recall Example 1).

Theorem 2. *Let f be any monotone Boolean function with $n \geq 2$ inputs, that is not constant and not a conjunction of some subset of the input bits. Then there exist parameters $\{\gamma_i, \delta_i\}_{i=1}^n$ such that the POP of the structured technology with the above parameters (and identical cost $c = 1$) is greater than 1.0233.*

Thus, our goal now is to give upper bounds on the POP for various technologies.

5 Quantifying the Gain by Mixing

5.1 POP for General Technologies

We first show that the POP can be bounded by the principal's price of unaccountability [3]. The principal's price of unaccountability ($POU_P(t)$) is the worse ratio (over all v), of the utility of the principal in the observable-actions case, and the utility of the principal in the hidden-actions case (see formal definition of POU_P in the full version).

Theorem 3. *For any technology t it holds that $POU_P(t) \geq POP(t)$.*

However, this bound is rather weak. To best see this, note that the principal's price of unaccountability for AND might be unbounded (see [3]). Yet, as shown in section 1, $POP(AND) = 1$.

In this section we provide better bounds on technologies with identical costs. We begin by characterizing the payments for a mixed contract. We show that under a mixed profile, each agent in the support of the contract is paid at least the minimal payment to a single agent under a pure profile with the same support, and at most the maximal payment.

For a mixed profile $q = (q_1, q_2, \dots, q_n)$, let $S(q)$ be the support of q , that is, $i \in S(q)$ if and only if $q_i > 0$. Similarly, for a pure profile $a = (a_1, a_2, \dots, a_n)$ let $S(a)$ be the support a . Under the mixed profile q , agent $i \in S(q)$ is being paid $p_i(q_{-i}) = \frac{c_i}{t(1, q_{-i}) - t(0, q_{-i})}$. Similarly, under the pure profile a , agent $i \in S(a)$ is being paid $p_i(S(a) \setminus \{i\}) = p_i(a_{-i}) = \frac{c_i}{t(S(a)) - t(S(a) \setminus \{i\})}$, where $t(T)$ is the success probability when $a_j = 1$ for $j \in T$, and $a_j = 0$ for $j \notin T$.

Lemma 4. For a mixed profile $q = (q_1, q_2, \dots, q_n)$, and for any agent $i \in S(q)$ let $S_{-i} = S(q) \setminus \{i\}$ be the support of q excluding i . It holds that

$$\max_{T \subseteq S_{-i}} p_i(T) \geq p_i(q_{-i}) \geq \min_{T \subseteq S_{-i}} p_i(T)$$

In what follows, we consider two general families of technologies with n agents: anonymous technologies and technologies that exhibit decreasing returns to scale (DRS). DRS technologies are technologies with decreasing marginal contribution (more effort by others decrease the contribution of an agent). For both families we present a bound of n on the POP.

We begin with a formal definition of DRS technologies.

Definition 3. A technology success function t exhibits (weakly) **decreasing returns to scale (DRS)** if for every i , and every $b \succeq a$

$$t(b_i, b_{-i}) - t(a_i, b_{-i}) \leq t(b_i, a_{-i}) - t(a_i, a_{-i})$$

Theorem 4. For any anonymous technology or a (non-anonymous) technology that exhibits DRS, it holds that $POP(t) \leq n$.

We also prove a bound on the POP for any technology with 2 agents (even not anonymous), and an improved bound for the anonymous case.

Theorem 5. For any technology t (even non-anonymous) with 2 agents, it holds that $POP(t) \leq 2$. If t is anonymous then $POP(t) \leq 3/2$.

We do not provide bounds for any non-anonymous technology, this is left as an open problem for future research.

Open Problem 1 Provide an upper bound on the POP for general technologies.

As mentioned in the introduction, we believe that the obtained bounds are very weak. In particular, we conjecture that there exists a constant C that bounds the POP for any technology. Moreover, we believe that a non-anonymous OR technology with 2 agents yields the highest possible POP. This motivates us to explore the POP for the OR technology in more detail.

5.2 POP for the OR Technology

As any OR technology (even non-anonymous) exhibits DRS (see claim in the full version), this implies a bound of n on the POP of the OR technology. Yet, for anonymous OR technology we present improved bounds on the POP. In particular, if $\gamma = 1 - \delta < 1/2$ we can bound the POP by 1.154...

Theorem 6. For any anonymous OR technology with n agents:

1. If $1 > \delta > \gamma > 0$:
 - (a) $POP \leq \frac{1 - (1 - \delta)^n}{\delta} \leq n - (n - 1)\delta$.
 - (b) POP goes to 1 as n goes to ∞ (for any fixed δ) or when δ goes to 1 (for any fixed $n \geq 2$).

2. If $\frac{1}{2} > \gamma = 1 - \delta > 0$:

(a) $POP \leq \frac{2(3-2\sqrt{3})}{3(\sqrt{3}-2)} (= 1.154..)$.

(b) POP goes to 1 as γ goes to 0 or as γ goes to $\frac{1}{2}$ (for any fixed $n \geq 2$).

While the bounds for anonymous OR technologies for the case in which $\delta = 1 - \gamma$ are much better than the general bounds, they are still not tight. The highest POP we were able to obtain by simulations was of 1.0233 for $\delta > \gamma$, and 1.0187 for $\delta = 1 - \gamma$ (see Section 3), but deriving the exact bound analytically is left as an open problem.

Open Problem 2 *What is the POP for an anonymous OR technology? what is it for a non-anonymous OR technology?*

6 The Robustness of Mixed Nash Equilibria

In order to induce an agent i to truly mix between exerting effort and shirking, p_i must be equal exactly to $c_i/\Delta_i(q_{-i})$ (see claim 2). Even under an increase of ϵ in p_i , agent i is no longer indifferent between $a_i = 0$ and $a_i = 1$, and the equilibrium falls apart. This is in contrast to the pure case, in which any $p_i \geq \frac{c_i}{\Delta_i(a_{-i})}$ will maintain the required equilibrium. This delicacy exhibits itself through the robustness of the obtained equilibrium to deviations in *coalitions* (as opposed to the unilateral deviations as in Nash). A "strong equilibrium" [1] requires that no subgroup of players (henceforth *coalition*) can coordinate a joint deviation such that every member of the coalition strictly improves his utility.

Definition 4. *A mixed strategy profile $q \in [0, 1]^n$ is a strong equilibrium (SE) if there does not exist any coalition $\Gamma \subseteq N$ and a strategy profile $q'_\Gamma \in \times_{i \in \Gamma} [0, 1]$ such that for any $i \in \Gamma$, $u_i(q'_\Gamma, q_\Gamma) > u_i(q)$.*

In [3] we show that under the payments that induce the pure strategy profile S^* as the best pure Nash equilibrium (i.e., the pure Nash equilibrium that maximizes the principal's utility), S^* is also a *strong equilibrium*. In contrast to the pure case, we next show that any non-degenerate mixed Nash equilibrium q in which there exist at least two agents that truly mix (i.e., $\exists i \neq j$ s.t. $q_i, q_j \in (0, 1)$), can never be a strong equilibrium. This is because if the coalition $\Gamma = \{i | q_i \in (0, 1)\}$ deviate to q'_Γ in which each $i \in \Gamma$ exerts effort with probability 1, each agent $i \in \Gamma$ strictly improves his utility.

Theorem 7. *If the mixed optimal contract q includes at least two agents that truly mix ($\exists i \neq j$ s.t. $q_i, q_j \in (0, 1)$), then q is not a strong equilibrium.*

In any OR technology, for example, it holds that in any non-degenerate mixed equilibrium at least two agents truly mix (see lemma 3). Therefore, no non-degenerate contract in the OR technology can be a strong equilibrium.

As generically a mixed Nash contract is not a strong equilibrium while a pure Nash contract always is, if the principal wishes to induce a strong Nash equilibrium (e.g., when the agents can coordinate their moves), he can restrict himself to inducing a pure Nash equilibrium, and his loss from doing so is bounded by the POP (see section 5).

7 Algorithmic Aspects

The computational hardness of finding the optimal mixed contract depends on the representation of the technology and how it is being accessed. For a black-box access and for the special case of read-once networks, we generalize our hardness results of the pure case [3] to the mixed case. The main open question is whether it is possible to find the optimal mixed contract in polynomial time, given a table representation of the technology (the optimal pure contract can be found in polynomial time in this case). Our generalization theorems follow.

Theorem 8. *Given as input a black box for a success function t (when the costs are identical), and a value v , the number of queries that is needed, in the worst case, to find the optimal mixed contract is exponential in n .*

Even if the technology is a structured technology and further restricted to be the source-pair reliability of a read-once network (see [3]), computing the optimal mixed contract is hard.

Theorem 9. *The optimal mixed contract problem for read once networks is $\#P$ -hard (under Turing reductions).*

References

1. R. Aumann. Acceptable Points in General Cooperative n -Person Games. In *Contributions to the Theory of Games*, volume 4, 1959.
2. M. Babaioff, M. Feldman, and N. Nisan. Combinatorial agency. In *ACM EC'06*, pages 18–28, 2006.
3. M. Babaioff, M. Feldman, and N. Nisan. Combinatorial agency. In *Full version. Under submission.*, 2006.
4. Y. K. Che and S. W. Yoo. Optimal Incentives in Teams. *American Economic Review*, 91:525–541, 2001.
5. M. Feldman, J. Chuang, I. Stoica, and S. Shenker. Hidden-Action in Multi-Hop Routing. In *EC'05*.
6. B. Holmstrom. Moral Hazard in Teams. *Bell Journal of Economics*, 13:324–340, 1982.
7. H. Itoh. Incentives to Help Multi-Agent Situations. *Econometrica*, 59:611–636, 1991.
8. A. Mass-Colell, M. Whinston, and J. Green. *Microeconomic Theory*. Oxford University Press, 1995.
9. N. Nisan and A. Ronen. Algorithmic mechanism design. *Games and Economic Behaviour*, 35:166 – 196, 2001. A preliminary version appeared in STOC 1999.
10. R. Strausz. Moral hazard in sequential teams. *Departmental Working Paper. Free University of Berlin*, 1996.
11. E. Winter. Incentives and Discrimination. *American Economic Review*, 94:764–773, 2004.