Game Theory	Spring Semester, 2005/2006

Lecture 9: May 2006

Lecturer: Michal Feldman *Scribe: Sergey Pisotskiy, Ilya Even-Han, Liza Potikha* [1]

9.1 Introduction

9.1.1 Background

One of the most striking characteristics of modern computer networks – in particular the Internet – is that different parts of it are owned and operated by different individuals, firms, and organizations. The analysis and design of protocols for this environment thus naturally needs to take into account the different "selfish" economic interests of the different participants. A significant part of the difficulty stems from underlying asymmetries of information: one participant may not know everything that is known or done by another.

In this lecture we will deal with the complementary lack of knowledge, that of *hidden actions*. In many cases the actual behaviors – actions – of the different participants are "hidden" from others and only influence the final outcome indirectly. In this lecture we will study hidden actions in multi-agents computational settings - which means that any agent chooses it's behavior, which is unknown to others, and only final result is public. It's completely understandable that each agent has it's influence on the final result, but we have no way to determinate what actions have been chosen by the agents.

A good example for the problem 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 (and there is no problem if the packet have reached there more than one time) 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?

Our approach to this problem is based on the well studied principal-agent problem in economic theory: How can a "principal" motivate a rational "agent" to exert costly effort towards the welfare of the principal? The crux of the model is that the agent's action (i.e.

[1] Based on "Combinatorial Agency" paper by Moshe Babaioff, Michal Feldman, Noam Nisan.

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. 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.

9.1.2 Our Model

In this model each of n agents has a set of possible *actions* - in our simple case there will be two simple actions (in other words the agent will have binary actions) : effort ($a_i = 1$ and $c_i(1) = c$ for some $c > 0$) or no effort ($a_i = 0$ and $c_i(0) = 0$); the combination of actions by the players results in some *outcome* (project succeeds, and principal gets value v, or project fails,and principal gets value 0), where the outcome is determined probabilistically. The main part of the specification of a problem in this model is a function t that specifies this distribution for each n-tuple of agents' actions: $t : \{0,1\}^n \to [0,1]$, where $t(a_1,\dots,a_n)$ denotes the probability of project success, given $(a_1,\dots,a_n)$ vector of agents actions. We assume a very rational property of this function, which is monotonicity:

$$\forall i \in N,\ \forall a_{-i} \in A_{-i} \quad t(1, a_{-i}) > t(0, a_{-i})$$

In other words, we want that the distribution function will have the next property: more effort by an agent leads to a better probability of success. 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 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. In other words, the general type of the contract is : project succeeds $\to$ agent i receives p_i, otherwise he gets 0.

Players' utilities, under action profile $a = (a_1,\dots,a_n)$ and value v: Given a set of contracts, the agents will each optimize his own utility: i.e. will choose the action that maximizes his expected payment minus the cost of his action. Since the outcome depends on the actions of all players together, the agents are put in a game and are assumed to reach a Nash equilibrium. The principal's problem is of designing an optimal set of contracts: i.e. contracts that maximize his expected utility from the outcome, minus his expected total payment.

We then consider a more concrete model which concerns a subclass of problem instances where this exponential size table is succinctly represented. This subclass will provide many natural types of problem instances. In this subclass every agent performs a subtask which succeeds with a low probability γ if the agent does not exert effort and with a higher probability $\delta > \gamma$, if the agent does exert effort. The whole project succeeds as a deterministic Boolean function of the success of the subtasks. This Boolean function can now be represented in various ways. Two basic examples are the "AND" function in which the project succeeds only if all subtasks succeed, and the "OR" function which succeeds if any of the subtasks succeeds.

9.2 Model and Preliminaries

9.2.1 The General Setting

A principal employs a set of agents N of size n. Each agent $i \in N$ has a possible set of actions $A_i \in \{0,1\}$, and a cost (effort) $c_i(a_i) \geq 0$ for each possible action $a_i \in A_i$ ($c_i : A_i \rightarrow \Re_+$). The actions of all players determine, in a probabilistic way, a "contractible" outcome $o \in O$, according to a success function $t : A_1 \times, \ldots \times A_n \rightarrow \Delta(O)$ (where $\Delta(O)$ denotes the set of probability distributions on O). A technology is a pair, (t, c), of a success function, t, and cost functions, $c = (c_1, c_2, \ldots, c_n)$. The principal has a certain value for each possible outcome, given by the function $v : O \rightarrow \Re$. Actions of the players are invisible, but the final outcome o is visible to him and to others (in particular the court), and he may design enforceable contracts based on the final outcome.

Thus the contract for agent i is a function (payment) $p_i : O \rightarrow \Re$; again, we will also view p_i as a function on $\Delta(O)$.

In this notation the success function $t : \{0,1\}^n \rightarrow [0,1]$, where $t(a_1, \ldots, a_n)$ denotes the probability of project success where players with $a_i = 0$ do not exert effort and incur no cost, and players with $a_i = 1$ do exert effort and incur a cost of c_i.

Additionally, we assume that $t(a) > 0$ for any $a \in A$ (or equivalently, $t(0, 0, \ldots, 0) > 0$).

The monotonicity property for this function is as follows. Denote by $a_{-i} \in A_{-i}$ the $(n-1)$-dimensional vector of the actions of all agents excluding agent i. i.e., $a_{-i} = (a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n)$. Then a success function t must satisfy:

$$\forall i \in N,\ \forall a_{-i} \in A_{-i} \quad t(1, a_{-i}) > t(0, a_{-i})$$

Given this setting, the agents have been put in a game, where the utility of agent i under the vector of actions $a = (a_1, \ldots, a_n)$ is given by $u_i(a) = p_i(t(a)) - c_i(a_i)$, where p_i is the payment function of player i.

The agents will be assumed to reach Nash equilibrium, if such 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 $u(a) = v(t(a)) - \sum_i p_i(t(a))$, where the actions $a_1, \ldots, a_n$ are at Nash-equilibrium.

In the case of multiple Nash equilibria we let the principal choose the equilibrium, thus focusing on the "best" Nash equilibrium. A variant, which is similar in spirit to "strong implementation" in mechanism design would be to take the worst Nash equilibrium, or even, stronger yet, to require that only a single equilibrium exists.

Finally, the *social welfare* for $a \in A$ is $u(a) + \sum_{i \in N} u_i(a) = v(t(a)) - \sum_{i \in N} c_i(a_i)$.

9.2.2 Nash equilibrium in our model and POU

Definition 1 *The marginal contribution of agent i, denoted by Δ_i, is the difference between the probability of success when i exerts effort and when he shirks.*

$$\Delta_i(a_{-i}) = t(1, a_{-i}) - t(0, a_{-i})$$

Note that since t is monotone, Δ_i is a strictly positive function. At this point we can already make some simple observations. The best action, $a_i \in A_i$, of agent i can now be easily determined as a function of what the others do, $a_{-i} \in A_{-i}$, and his contract p_i.

Claim 9.1 *Given a profile of actions a_{-i}, agent i's best strategy is $a_i = 1$ if $p_i \geq \frac{c_i}{\Delta_i(a_{-i})}$, and is $a_i = 0$ if $p_i \leq \frac{c_i}{\Delta_i(a_{-i})}$.(In the case of equality the agent is indifferent between the two alternatives.)*

Since $p_i \geq \frac{c_i}{\Delta_i(a_{-i})}$ if and only if $u_i(1, a_{-i}) = p_i \cdot t(1, a_{-i}) - c_i \geq p_i \cdot t(0, a_{-i}) = u_i(0, a_{-i})$, i's best strategy is to choose $a_i = 1$ in this case.

This allows us to specify the contracts that are the principal's optimal, for inducing a given equilibrium.

Observation 1 *The best contracts (for the principal) that induce $a \in A$ as an equilibrium are $p_i = 0$ for agent i who exerts no effort ($a_i = 0$), and $p_i = \frac{c_i}{\Delta_i(a_{-i})}$ for agent i who exerts effort ($a_i = 1$).*

In this case, the expected utility of agent i who exerts effort is $c_i \cdot \left(\frac{t(1,a_{-i})}{\Delta_i(a_{-i})} - 1\right)$, and 0 for an agent who shirk. The principal's expected utility is given by $u(a, v) = (v - P) \cdot t(a)$, where P is the total payment in case of success, given by $P = \sum_{i|a_i=1} \frac{c_i}{\Delta_i(a_{-i})}$.

We say that the principal *contracts with* agent i if $p_i > 0$ (and $a_i = 1$ in the equilibrium $a \in A$). The principal's goal is to maximize his utility given his value v, i.e. to determine the profile of actions $a^* \in A$, which gives the highest value of $u(a, v)$ in equilibrium.

Choosing $a \in A$ corresponds to choosing a set S of agents that exert effort ($S = \{i|a_i = 1\}$). We call the set of agents S^* that the principal contracts with in a^* ($S^* = \{i|a_i^* = 1\}$) an *optimal contract* for the principal at value v. We sometimes abuse notation and denote $t(S)$ instead of $t(a)$, when S is exactly the set of agents that exert effort in $a \in A$.

A natural yardstick by which to measure this decision is the non-strategic case, i.e. when the agents need not be motivated but are rather controlled directly by the principal (who also bears their costs). In this case the principal will simply choose the profile $a \in A$ that optimizes the social welfare (global efficiency), $t(a) \cdot v - \sum_{i|a_i=1} c_i$. The worst ratio between the social welfare in this non-strategic case and the social welfare for the profile $a \in A$ chosen by the principal in the agency case, may be termed the *price of unaccountability*.

Given a technology $(t, \vec{c})$, let $S^*(v)$ denote the optimal contract in the agency case and let $S^*_{ns}(v)$ denote an optimal contract in the non-strategic case, when the principal's value is v. The social welfare for value v when the set S of agents is contracted is $t(S) \cdot v - \sum_{i \in S} c_i$ (in both the agency and non-strategic cases).

Definition 2 *The* price of unaccountability $POU(t, \vec{c})$ *of a technology* $(t, \vec{c})$ *is defined as the worst ratio (over v) between the total social welfare in the non-strategic case and the agency case:*

$$POU(t, \vec{c}) = Sup_{v>0} \frac{t(S^*_{ns}(v)) \cdot v - \sum_{i \in S^*_{ns}(v)} c_i}{t(S^*(v)) \cdot v - \sum_{i \in S^*(v)} c_i}$$

In cases where several sets are optimal in the agency case, we take the worst set (i.e., the set that yields the lowest social welfare).

When the technology $(t, \vec{c})$ is clear in the context we will use POU to denote the price of unaccountability for technology $(t, \vec{c})$. Note that the POU is at least 1 for any technology.

9.2.3 Structured Technology Functions

In order to be more concrete, we will especially focus on technology functions whose structure can be described easily as being derived from independent agent tasks – we call these *structured technology functions*. This subclass will first give us some natural examples of technology function, and will also provide a succinct and natural way to represent the technology functions.

In a structured technology function, each individual succeeds or fails in his own "task" independently. The project's success or failure depends, possibly in a complex way, on the set of successful sub-tasks. Thus we will assume a monotone Boolean function $f : \{0, 1\}^n \to \{0, 1\}$ which denotes whether the project succeeds as a function of the success of the n agents' tasks (and is not determined by any set of $n - 1$ agents). Additionally there are constants $0 < \gamma_i < \delta_i < 1$, where γ_i denotes the probability of success for agent i if he does not exert effort, and δ_i $(> \gamma_i)$ denotes the probability of success if he does exert effort. In order to reduce the number of parameters, we will restrict our attention to the case where $\gamma_1 = \ldots = \gamma_n = \gamma$ and $\delta_1 = \ldots = \delta_n = 1 - \gamma$ thus leaving ourselves with a single parameter γ s.t. $0 < \gamma < \frac{1}{2}$.

Under this structure, the technology function t is defined by $t(a_1, \ldots, a_n)$ being the probability that $f(x_1, \ldots, x_n) = 1$ where the bits $x_1, \ldots, x_n$ are chosen according to the following distribution: if $a_i = 0$ then $x_i = 1$ with probability γ and $x_i = 0$ with probability $1-\gamma$; otherwise, i.e. if $a_i = 1$, then $x_i = 1$ with probability $1-\gamma$ and $x_i = 0$ with probability γ. We denote $x = (x_1, \ldots, x_n)$.

A few simple examples should be in order here:

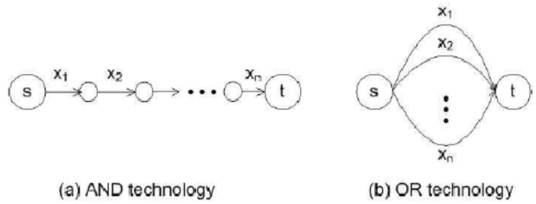

Figure 9.1: (a) AND and (b) OR technologies.

1. The "AND" technology: $f(x_1, \ldots, x_n)$ 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. This is shown graphically as a read-once network in Figure 9.1(a). If m agents exert effort ($\sum_i a_i = m$), then $t(a) = t_m = \gamma^{n-m}(1-\gamma)^m$. E.g. for two players, the technology function $t(a_1 a_2) = t_{a_1+a_2}$ is given by $t_0 = t(00) = \gamma^2$, $t_1 = t(01) = t(10) = \gamma(1-\gamma)$, and $t_2 = t(11) = (1-\gamma)^2$.
2. The "OR" technology: $f(x_1, \ldots, x_n)$ 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. This is shown graphically as a read-once network in Figure 9.1(b). If m agents exert effort, then $t_m = 1 - \gamma^m(1-\gamma)^{n-m}$. E.g. for two players, the technology function is given by $t(00) = 1 - (1-\gamma)^2$, $t(01) = t(10) = 1 - \gamma(1-\gamma)$, and $t(11) = 1 - \gamma^2$.
3. The "Or-of-Ands" (OOA) technology: $f(x)$ is the logical disjunction of conjunctions. In the simplest case of equal-length clauses (denote by n_c the number of clauses and by n_l their length), $f(x) = \bigvee_{j=1}^{n_c}(\bigwedge_{k=1}^{n_l} x_k^j)$. Thus the project succeeds if in at least one clause all agents succeed in their tasks. This is shown graphically as a read-once network in Figure 9.2(a). If m_i agents on path i exert effort, then $t(m_1, ..., m_{n_c}) = 1 - \prod_i(1-\gamma^{n_l - m_i}(1-\gamma)^{m_i})$. E.g. for four players, the technology function $t(a_1^1\, a_2^1, a_1^2\, a_2^2)$ is given by $t(00,00) = 1 - (1-\gamma^2)^2$, $t(01,00) = t(10,00) = t(00,01) = t(00,10) = 1 - (1-\gamma(1-\gamma))(1-\gamma^2)$, and so on.
4. The "And-of-Ors" (AOO) technology: $f(x)$ is the logical conjunction of disjunctions. In the simplest case of equal-length clauses (denote by n_l the number of clauses and by n_c their length), $f(x) = \bigwedge_{j=1}^{n_l}(\bigvee_{k=1}^{n_c} x_k^j)$. Thus the project succeeds if at least one agent from each disjunctive-form-clause succeeds in his tasks. This is shown graphically as a read-once network in Figure 9.2(b). If m_i agents on clause i exert effort, then $t(m_1, ..., m_{n_c}) = \prod_i(1-\gamma^{m_i}(1-\gamma)^{n_c - m_i})$. E.g. for four players, the technology function $t(a_1^1\, a_2^1, a_1^2\, a_2^2)$ is given by $t(00,00) = (1-(1-\gamma)^2)^2$, $t(01,00) = t(10,00) = t(00,01) = t(00,10) = (1-\gamma(1-\gamma))(1-(1-\gamma)^2)$, and so on.

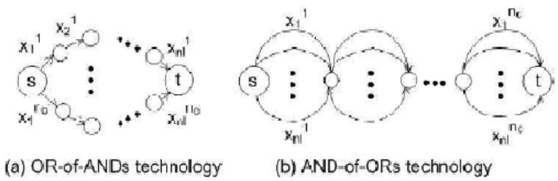

Figure 9.2: Graphical representations of (a) OOA and (b) AOO technologies.

9.3 Analysis of Some Anonymous Technologies

A success function t is called *anonymous* if it is symmetric with respect to the players. I.e. $t(a_1, \ldots, a_n)$ depends only on $\sum_{i \in N} a_i$ (the number of agents that exert effort). A technology (t, c) is *anonymous* if t is anonymous and the cost c is identical to all agents. Of the examples presented above, the AND and OR technologies were anonymous (but not AOO and OOA). As for an anonymous t only the number of agents that exert effort is important, we can shorten the notations and denote $t_m = t(1^m, 0^{n-m})$, $\Delta_m = t_{m+1} - t_m$, $p_m = \frac{c}{\Delta_{m-1}}$,meaning the payment principal has to give to each one of m agents to motivate them, and his utility in this case is $u_m = t_m \cdot (v - m \cdot p_m)$, for the case of identical cost c,

9.3.1 AND and OR Technologies

Let us start with a direct and full analysis of the AND and OR technologies for two players for the case $\gamma = 1/4$ and $c = 1$.

Example 1 AND **technology with two agents,** $c = 1$**,** $\gamma = 1/4$**:** *we have* $t_0 = \gamma^2 = 1/16$*,* $t_1 = \gamma(1-\gamma) = 3/16$*, and* $t_2 = (1-\gamma)^2 = 9/16$ *thus* $\Delta_0 = 1/8$ *and* $\Delta_1 = 3/8$*. The principal has 3 possibilities: contracting with 0, 1, or 2 agents. Let us write down the expressions for his utility in these 3 cases:*

- **0 Agents:** *No agent is paid thus and the principal's utility is* $u_0 = t_0 \cdot v = v/16$*.*
- **1 Agent:** *This agent is paid* $p_1 = c/\Delta_0 = 8$ *on success and the principal's utility is* $u_1 = t_1(v - p_1) = 3v/16 - 3/2$*.*
- **2 Agents:** *each agent is paid* $p_2 = c/\Delta_1 = 8/3$ *on success, and the principal's utility is* $u_2 = t_2(v - 2p_2) = 9v/16 - 3$*.*

Notice that the option of contracting with one agent is always inferior to either contracting with both or with none, and will never be taken by the principal. The principal will contract

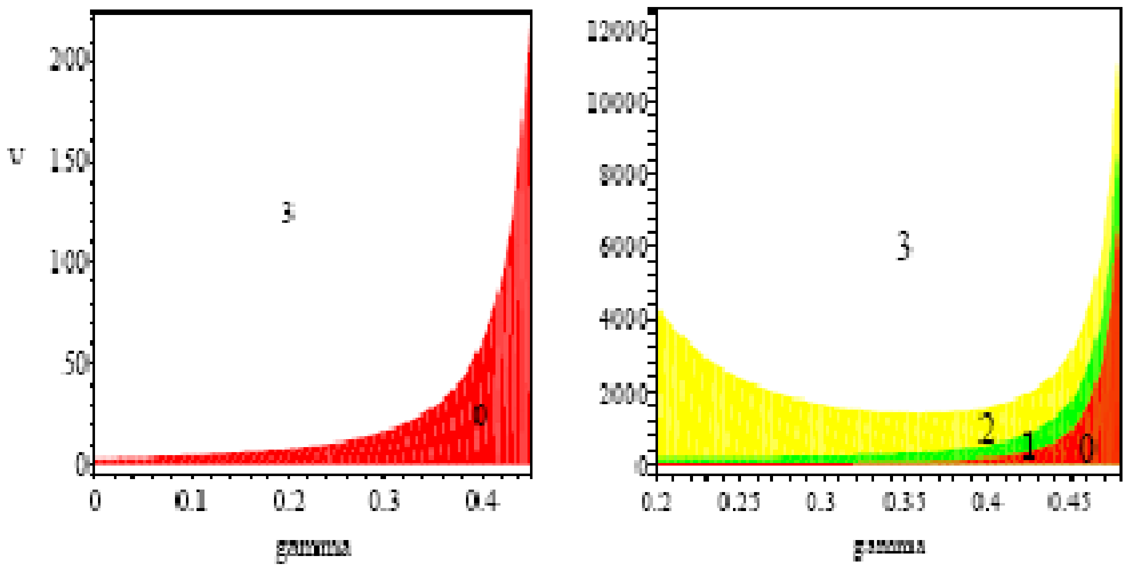

Figure 9.3: Number of agents in the optimal contract of the AND (left) and OR (right) technologies with 3 players, as a function of γ and v. AND technology: either 0 or 3 agents are contracted, and the transition value is monotonic in γ. OR technology: for any γ we can see all transitions.

with no agent when $v < 6$, with both agents whenever $v > 6$, and with either non or both for $v = 6$.

This should be contrasted with the non-strategic case in which the principal completely controls the agents (and bears their costs) and thus simply optimizes globally. In this case the principal will make both agents exert effort whenever $v \geq 4$. Thus for example, for $v = 6$ the globally optimal decision (non-strategic case) would give a global utility of $6 \cdot 9/16 - 2 = 11/8$ while the principal's decision (in the agency case) would give a global utility of $3/8$, giving a ratio of $11/3$.

It turns out that this is the worst price of unaccountability in this example, and it is obtained exactly at the transition point of the agency case, as we show below.

Example 2 **OR technology with two agents, $c = 1$, $\gamma = 1/4$:** *we have $t_0 = 1-(1-\gamma)^2 = 7/16$, $t_1 = 1 - \gamma(1-\gamma) = 13/16$, and $t_2 = 1 - \gamma^2 = 15/16$ thus $\Delta_0 = 3/8$ and $\Delta_1 = 1/8$. Let us write down the expressions for the principal's utility in these three cases:*

- **0 Agents:** *No agent is paid and the principal's utility is $u_0 = t_0 \cdot v = 7v/16$.*
- **1 Agent:** *This agent is paid $p_1 = c/\Delta_0 = 8/3$ on success and the principal's utility is $u_1 = t_1(v - p_1) = 13v/16 - 13/6$.*

- **2 Agents:** *each agent is paid* $p_2 = c/\Delta_1 = 8$ *on success, and the principal's utility is* $u_2 = t_2(v - 2p_2) = 15v/16 - 15/2$.

Now contracting with one agent is better than contracting with none whenever $v > 52/9$ *(and is equivalent for* $v = 52/9$*), and contracting with both agents is better than contracting with one agent whenever* $v > 128/3$ *(and is equivalent for* $v = 128/3$*), thus the principal will contract with no agent for* $0 \le v \le 52/9$*, with one agent for* $52/9 \le v \le 128/3$*, and with both agents for* $v \ge 128/3$*.*

In the non-strategic case, in comparison, the principal will make a single agent exert effort for $v > 8/3$*, and the second one exert effort as well when* $v > 8$*.*

It turns out that the price of unaccountability here is $19/13$, and is achieved at $v = 52/9$, which is exactly the transition point from 0 to 1 contracted agents in the agency case. This is not a coincidence that in both the AND and OR technologies the POU is obtained for v that is a transition point .

Lemma 9.2 *For any given technology* $(t, \vec{c})$ *the price of unaccountability* $POU(t, \vec{c})$ *is obtained at some value* v *which is a transition point, of either the agency or the non-strategic cases.*

Proof sketch: We look at all transition points in both cases. For any value lower than the first transition point, 0 agents are contracted in both cases, and the social welfare ratio is 1. Similarly, for any value higher than the last transition point, n agents are contracted in both cases, and the social welfare ratio is 1. Thus, we can focus on the interval between the first and last transition points. Between any pair of consecutive points, the social welfare ratio is between two linear functions of v (the optimal contracts are fixed on such a segment). We then show that for each segment, the suprimum ratio is obtained at an end point of the segment (a transition point). As there are finitely many such points, the global suprimum is obtained at the transition point with the maximal social welfare ratio. □

We already see a qualitative difference between the AND and OR technologies (even with 2 agents): in the first case either all agents are contracted or none, while in the second case, for some intermediate range of values v, exactly one agent is contracted. Figure 9.3 shows the same phenomena for AND and OR technologies with 3 players.

Theorem 9.3 *For any anonymous* AND *technology*[1]*:*

- *there exists a value*[2] $v_* < \infty$ *such that for any* $v < v_*$ *it is optimal to contract with no agent, for* $v > v_*$ *it is optimal to contract with all* n *agents, and for* $v = v_*$*, both contracts* $(0, n)$ *are optimal.*

[1] AND technology with any number of agents n and any γ, and any identical cost c.

[2] v_* is a function of n, γ, c.

- *the price of unaccountability is obtained at the transition point of the agency case, and is*

$$POU = \max\{(\frac{1}{\gamma} - 1)^{n-1}, (1 - \frac{\gamma}{1-\gamma})\}$$

Proof sketch: For any fixed number of contracted agents, k, the principal's utility is a linear function in v, where the slope equals the success probability under k contracted agents. Thus, the optimal contract corresponds to the maximum over a set of linear functions. Let v_* denote the point at which the principal is indifferent between contracting with 0 or n agents. At v_*, the principal's utility from contracting with 0 (or n) agents is higher than his utility when contracting with any number of agents $k \in \{1, \ldots, n-1\}$. As the number of contracted agents is monotonic non-decreasing in the value (due to Lemma 9.5), for any $v < v_*$, contracting with 0 agents is optimal, and for any $v > v_*$, contracting with n agents is optimal. This is true for both the agency and the non-strategic cases. □

The property of a single transition occurs in both the agency and the non-strategic cases, where the transition occurs at a smaller value of v in the non-strategic case. Notice that the POU is not bounded across the AND family of technologies (for various n, γ) as $POU \to \infty$ either if $\gamma \to 0$ (for any given $n \geq 2$) or $n \to \infty$ (for any fixed $\gamma \in (0, \frac{1}{2})$).

Next we consider the OR technology and show that it exhibits all n transitions.

Theorem 9.4 *For any anonymous OR technology, there exist finite positive values $v_1 < v_2 < \ldots < v_n$ such that for any v s.t. $v_k < v < v_{k+1}$, contracting with exactly k agents is optimal (for $v < v_1$, no agent is contracted, and for $v > v_n$, all n agents are contracted). For $v = v_k$, the principal is indifferent between contracting with $k-1$ or k agents.*

Proof sketch: To prove the claim we define v_k to be the value for which the principal is indifferent between contracting with $k-1$ agents, and contracting with k agents. We then show that for any k, $v_k < v_{k+1}$. As the number of contracted agents is monotonic non-decreasing in the value (due to Lemma 9.5), $v_1 < v_2 < \ldots < v_n$ is a sufficient condition for the theorem to hold. □

Observation 2 *While in the AND technology the POU for $n = 2$ is not bounded from above (for $\gamma \to 0$), the highest POU in OR technology with two agents is 2 (for $\gamma \to 0$).*

9.4 Non-Anonymous Technologies

In non-anonymous technologies (even with identical costs), we need to talk about the contracted *set* of agents and not only about the number of contracted agents. In this section, we identify the sets of agents that can be obtained as the optimal contract for some v. These sets construct the *orbit* of a technology.

Definition 3 *For a technology t, a set of agents S is in the* orbit of t *if for some value v, the optimal contract is exactly with the set S of agents (where ties between different S's are broken according to a lexicographic order*[3]*). The* k-orbit *of t is the collection of sets of size exactly k in the orbit.*

in the non-strategic case the k-orbit of any technology with identical cost c is of size at most 1 (as all sets of size k has the same cost, only the one with the maximal probability can be on the orbit). Thus, the orbit of any such technology in the non-strategic case is of size at most $n+1$. We show that the picture in the agency case is very different.

A basic observation is that the orbit of a technology is actually an ordered list of sets of agents, where the order is determined by the following lemma.

Lemma 9.5 *(***Monotonicity lemma***) For any technology $(t, \vec{c})$, in both the agency and the non-strategic cases, the expected utility of the principal at the optimal contracts, the success probability of the optimal contracts, and the expected payment of the optimal contract, are all monotonically non-decreasing with the value.*

Proof: Suppose the sets of agents S_1 and S_2 are optimal in v_1 and $v_2 < v_1$, respectively. Let $Q(S)$ denote the expected total payment to all agents in S in the case that the principal contracts with the set S and the project succeeds (for the agency case, $Q(S) = t(S) \cdot \sum_{i \in S} \frac{c_i}{t(S) - t(S \setminus i)}$, while for the non-strategic case $Q(S) = \sum_{i \in S} c_i$). The principal's utility is a linear function of the value, $u(S, v) = t(S) \cdot v - Q(S)$. As S_1 is optimal at v_1, $u(S_1, v_1) \geq u(S_2, v_1)$, and as $t(S_2) \geq 0$ and $v_1 > v_2$, $u(S_2, v_1) \geq u(S_2, v_2)$. We conclude that $u(S_1, v_1) \geq u(S_2, v_2)$, thus the utility is monotonic non-decreasing in the value.

Next we show that the success probability is monotonic non-decreasing in the value. S_1 is optimal at v_1, thus:

$$t(S_1) \cdot v_1 - Q(S_1) \geq t(S_2) \cdot v_1 - Q(S_2)$$

S_2 is optimal at v_2, thus:

$$t(S_2) \cdot v_2 - Q(S_2) \geq t(S_1) \cdot v_2 - Q(S_1)$$

Summing these two equations, we get that $(t(S_1) - t(S_2)) \cdot (v_1 - v_2) \geq 0$, which implies that if $v_1 > v_2$ than $t(S_1) \geq t(S_2)$.

Finally we show that the expected payment is monotonic non-decreasing in the value. As S_2 is optimal at v_2 and $t(S_1) \geq t(S_2)$, we observe that:

$$t(S_2) \cdot v_2 - Q(S_2) \geq t(S_1) \cdot v_2 - Q(S_1) \geq t(S_2) \cdot v_2 - Q(S_1)$$

or equivalently, $Q(S_2) \leq Q(S_1)$, which is what we wanted to show. □

[3]This implies that there are no two sets with the same success probability in the orbit.

9.4.1 AOO and OOA Technologies

We begin our discussion of non-anonymous technologies with two examples; the And-of-Ors (*AOO*) and Or-of-Ands (*OOA*) technologies.

The *AOO* technology (see figure 9.2) is composed of multiple *OR-components* that are "And"ed together.

Theorem 9.6 *Let h be an anonymous OR technology, and let $f = \bigwedge_{j=1}^{n_c} h$ be the AOO technology that is obtained by a conjunction of n_c of these OR-components on disjoint inputs. Then for any value v, an optimal contract contracts with the same number of agents in each OR-component. Thus, the orbit of f is of size at most $n_l + 1$, where n_l is the number of agents in h.*

Conjecture 1 *In an OOA technology which is a disjunction of the same anonymous paths (with the same number of agents, γ and c, but over disjoint inputs), for any value v the optimal contract is constructed from some number of fully-contracted paths. Moreover, there exist $v_1 < \ldots < v_{n_l}$ such that for any v, $v_i \leq v \leq v_{i+1}$, exactly i paths are contracted.*

9.4.2 Orbit Characterization

The *AOO* is an example of a technology whose orbit size is linear in its number of agents. If conjecture 1 is true, the same holds for the *OOA* technology. What can be said about the orbit size of a general non-anonymous technology?

In case of identical costs, it is impossible for **all** subsets of agents to be on the orbit. This holds by the observation that the 1-orbit (a single agent that exerts effort) is of size at most 1. Only the agent that gives the highest success probability (when only he exerts effort) can be on the orbit (as he also needs to be paid the least). Nevertheless, we next show that the orbit can have exponential size.

A collection of sets of k elements (out of n) is "admissible", if every two sets in the collection differ by at least 2 elements (e.g. for k=3, 123 and 234 can not be together in the collection, but 123 and 345 can be).

Theorem 9.7 *Every admissible collection can be obtained as the $k-orbit$ of some t.*

Proof sketch: The proof is constructive. Let $\mathcal{S}$ be some admissible collection of k-size sets. For each set $S \in \mathcal{S}$ in the collection we pick ϵ_S, such that for any two admissible sets $S_i \neq S_j$, $\epsilon_{S_i} \neq \epsilon_{S_j}$. We then define the technology function t as follows: for any $S \in \mathcal{S}$, $t(S) = 1/2 - \epsilon_S$ and $\forall i \in S$, $t(S \setminus i) = 1/2 - 2\epsilon_S$. Thus, the marginal contribution of every $i \in S$ is ϵ_S. Note that since $\mathcal{S}$ is admissible, t is well defined, as for any two sets $S, S' \in \mathcal{S}$ and any two agents i, j, $S \setminus i \neq S' \setminus j$. For any other set Z, we define $t(Z)$ in a way that ensures that the marginal contribution of each agent in Z is a very small ϵ (the technical details appear in the paper). This completes the definition of t.

We show that each admissible set $S \in \mathcal{S}$ is optimal at the value $v_S = \frac{ck}{2\epsilon_S^2}$. We first show that it is better than any other $S' \in \mathcal{S}$. At the value $v_S = \frac{ck}{2\epsilon_S^2}$, the set S that corresponds to ϵ_S maximizes the utility of the principal. This result is obtained by taking the derivative of $u(S, v)$. Therefore S yields a higher utility than any other $S' \in \mathcal{S}$. We also pick the range of ϵ_S to ensure that at v_S, S is better than any other set $S' \setminus i$ s.t. $S' \in \mathcal{S}$. Now we are left to show that at v_S, the set S yields a higher utility than any other set $Z \notin \mathcal{S}$. The construction of $t(Z)$ ensures this since the marginal contribution of each agent in Z is such a small ϵ, that the payment is too high for the set to be optimal. □

Lemma 9.8 *For any $n \geq k$, there exists an admissible collection of k-size sets of size $\Omega(\frac{1}{n} \cdot \binom{n}{k})$.*

Proof sketch: The proof is based on an error correcting code that corrects one bit. Such a code has a distance ≥ 3, thus admissible. It is known that there are such codes with $\Omega(2^n/n)$ code words. To ensure that an appropriate fraction of these code words have weight k, we construct a new code by XOR-ing each code word with a random word r. The properties of XOR ensure that the new code remains admissible. Each code word is now uniformly mapped to the whole cube, and thus its probability of having weight k is $\binom{n}{k}/2^n$. Thus the expected number of weight k words is $\Omega(\binom{n}{k}/n)$, and for some r this expectation is achieved or exceeded. □

For $k = n/2$ we can construct an exponential size admissible collection, which by Theorem 9.7 can be used to build a technology with exponential size orbit.

Corollary 9.9 *There exists a technology (t, c) with orbit of size $\Omega(\frac{2^n}{n\sqrt{n}})$.*