

Combinatorial Agency with Audits

Raphael Eidenbenz
Computer Engineering and
Networks Laboratory (TIK)
ETH Zurich, Switzerland
eidenbenz@tik.ee.ethz.ch

Stefan Schmid
Chair for Efficient Algorithms
Computer Science Department
TU Munich, Germany
schmiste@in.tum.de

Abstract—This paper studies the question of how to overcome inefficiencies due to hidden actions in a rational milieu, such as a grid computing system with open clientele. We consider the so-called principal-agent model known from economic theory, where the members (or agents) of a distributed system collaborate in complex ways. We adopt the perspective of the principal and investigate auditing mechanisms that incentivize participants to contribute more to a common project. As conducting audits might be costly, the principal must balance the tradeoff between low auditing costs and the level of incentives offered to the participants to exert high effort. We present optimal solutions for this optimization problem in scenarios, where the project success either depends on all, on any or on the majority of the participants succeeding in their subtask. In the first case, we further find that with an increasing principal valuation, there is exactly one transition point where the optimal choices for achieving the maximal principal utility switch. Compared to a combinatorial agency without the leverage of audits, this transition occurs earlier.

I. INTRODUCTION

The Internet heralded a new computing paradigm that allows machines all over the planet to collaborate. It soon became the platform for many success stories such as the World Wide Web or the emailing system which had a deep impact on our society in general. Distributed computer networks today are decentralized. Different parts are operated by different stakeholders and the participating machines are no longer under a single administrative domain but under the control of individual users. This has led researchers to use game-theoretic models in order to gain deeper insights into the functioning of these systems.

Distributed computing solutions such as grid computing or peer-to-peer technology rely on resource contributions by the participants. However, in the absence of a corresponding incentive mechanism, participants may not be willing to share some of their bandwidth or CPU cycles, but prefer to free-ride. Unfortunately, it is often hard to uncover and penalize selfish participants, because their actions are *hidden*. For instance, in a computational grid network organized by a central server that has an open clientele, the client machines may compute different subtasks individually and return their results to contribute to the *overall* project. Some of these computations may be redundant, others not. The server seeks to provide incentives to the participants as to have them perform their calculations properly. However the server typically only evaluates the overall outcome rather than the quality of the individual computations. This situation

constitutes a so-called *principal-agent problem* where the principal seeks to influence the effort made by the agents indirectly by promising certain benefits or payments if the project as a whole succeeds. This problem is well-studied in economic theory and has many applications, e.g., in organization theory.

In this paper, we seek to complement the existing literature by focusing on the combinatorial structure of dependencies between the agents' actions. In particular, we study the question of whether and how auditing one or more agents can improve the efficiency of an agency. Of course, audits can be costly. In grid computing, for instance, in order to verify the result submitted by a client, the server may have to recompute the subtask. Or in the context of a software company, the Java classes programmed by an employee may be checked during a code review by the entire team. When auditing all agents is too expensive, it can make sense to audit only a subset of agents *at random*. This paper shows that such a sampling approach can indeed reduce the principal's cost to implement a desired outcome.

A. Related Work

Since the influential talk by Papadimitriou at STOC 2001 [25], many economic aspects (e.g., of the Internet) have been investigated from a computational (or algorithmic) point of view. In addition, measurement studies (e.g., [19]) have confirmed the presence of selfish participants in today's distributed systems, for example in peer-to-peer networks. Most often, selfish behavior has a negative impact on a system's performance. Researchers in the field of *algorithmic mechanism design* [23] seek to come up with implementations that improve the efficiency or social welfare in selfish milieus. The presence of selfish participants or cheaters in *grid computing* [15] is also well-known [7], [29]. Solutions based on so-called *ringers* [16] or *Merkle trees* [7] have been investigated. The work in [21] exploits the asymmetry between the time-consuming computation of subtasks from scratch and the more efficient result verification. In [6], the Lottery Tree mechanism encourages contributions (and the solicitation of new participants) in asymmetric network-effect systems such as BOINC [1]. One difficulty in providing cooperation incentives to the agents or players is *information asymmetry*. While algorithmic mechanism design studies how to extract *private information* from the participants, the principal-agent theory [9], [17], [18] deals with *hidden actions* that influence

the final outcome. It has its origins in the theory of the firms [4], [5], [20], [24]. Researchers have identified several practical means to improve an agency, such as investments into corporate culture, information systems (e.g., definition of milestones), reputation and trust solutions, incentive systems or bureaucratic control (cf also [13], [14]). Our work studies the possibilities of probabilistic auditing.

This paper builds upon the seminal work by Babaioff, Feldman and Nisan [2] who coined the notion of *combinatorial agency*. The authors initiated the study of handling *combinations* of agents rather than single agents, and put an emphasis on dealing with complex dependencies between the agents' actions. Babaioff et al. introduced the binary agency model as a simplification and derived very interesting results for different classes of action dependencies—which they call technologies. Their paper has opened a wide range of exciting questions. In [3], the same authors study combinatorial agency where the agents mix strategies, and it is shown that the principal can benefit compared to situations where merely pure Nash equilibria are considered. In [10], Emek and Feldman provide a complexity analysis of the *OR* and *series-parallel* technologies, and present a FPTAS for the *OR* technology. In contrast to this body of work, we investigate *the effect of audits*. We allow the principal to audit a number of agents and verify their work. While audits come at a certain cost, we show that the cooperation among the agents can be improved.

There are interesting results on the effect of control in *social psychology* [11]. It has been estimated that in the United States, more than 20 million workers were subject to electronic monitoring in 1993, and that companies spent more than USD 1 billion on monitoring software in 1996. For example, Longhaul trucking firms used the Global Positioning System (GPS) to track the truck driver's speed, fuel use, and route location [28]. Pitman et al. [26] find that a physically present surveillant can reduce the motivation of individuals; [27] comes to a similar conclusion in scenarios with video cameras. Falk et al. [12] evaluate the cost of distrust experimentally and find that the principal earns 31 percent less if agents are control-averse. In contrast to this work, we assume that the agents are rational in the strict sense that they aim at maximizing their individual revenue. As we will see, in this context, the possibility of conducting audits does not reduce the principal's utility. Moreover, we will shed light on the question of how much auditing is optimal for the principal.

B. Model

Our model is presented in three stages. We first explain the general principal-agent setting. Subsequently, we introduce the binary principal-agent problem as a simplification. Finally, we describe how a principal audits agents. Throughout this paper, we will refer to the principal as a *she* and to an agent as a *he*.

The Principal-Agent Setting: We consider a set of n (risk-neutral) *agents* $N = \{1, \dots, n\}$ which can be contracted and

audited by a *principal*. Each agent i has a set of possible *actions* A_i , where each action $a_i \in A_i$ entails a certain *cost* $c_i(a_i) \geq 0$ to the agent. The combination of the actions chosen by the agents yields a certain probabilistic *outcome* $o \in O$ according to a success function $t : A_1 \times \dots \times A_n \rightarrow \Delta(O)$, where Δ is the set of probability distributions on outcome O . This outcome is considered a monetary value (or some other form of benefit) which accrues at the principal. In the following, let $a_{-i} \in A_{-i}$ denote the $(n-1)$ -dimensional vector of the agents' actions excluding agent i , i.e., $a_{-i} = (a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n)$.

A crucial assumption underlying the principal-agent model is that the principal cannot see the individual agents' actions. The principal motivates the agents by offering each agent i a contract that specifies a *payment* p_i for each possible outcome *of the whole project*. Note that $p_i \geq 0$, i.e., agents cannot be fined (*limited liability constraint*). Given these contracts, the agents will optimize their personal utilities u_i , maximizing the expected payments p_i they receive minus the cost of the action $c_i(a_i)$. Since the agents' actions depend on each other, we can model the situation as a game and assume that the agents reach a Nash equilibrium.

In this paper, we adopt the point of view of the principal. The principal seeks to design an optimal set of contracts that maximizes her expected utility $u(o)$, i.e., the expected benefits minus the expected total payments $\sum_i p_i$. The principal thus creates a game for the agents, where the desired behavior is a Nash equilibrium. Note that although there might exist other Nash equilibria in the induced game, we assume that the agents reach the equilibrium desired by the principal. This is legitimate since any Nash equilibrium can be implemented for free, e.g., by providing insurances to the agents just in case the other agents would deviate from the specific Nash equilibrium (cf [8], [22]).

The Binary Agency: In this paper, we will investigate a simplified principal-agent model where each agent i only has two possible actions: *exert high effort* ($a_i = 1$) and *exert low effort* (or *shirk*) ($a_i = 0$). An agent's cost are 0 for $a_i = 0$ and c_i for $a_i = 1$. Moreover, the whole project can only either *succeed* or *fail*, that is, the set of possible outcomes O has cardinality two, and $O = \{\text{success}, \text{failure}\}$. As has been pointed out by Babaioff et al. [2], this model already contains the main interesting ingredients. How the agents are organized and how the project success builds on the combination of their contributions depends on the nature of the project. For instance, in an assembly line setting, each agent may be perfectly specialized and contribute exactly one unique part to the project. In this scenario, there is no redundancy and the project succeeds if and only if all agents successfully provide their part. We will refer to such a project as an *AND technology*. In contrast, in the *OR technology*, the project succeeds if and only if at least one of the agents successfully contributes its (sub-)task.

A *technology* is characterized by its success function $t(a)$. In this paper, we will focus on *anonymous technologies*: a technology is called anonymous if the success function

$t(a_1, \dots, a_n)$ depends only on the number—and not on a particular subset—of agents who exert high effort. We then write the success function as t_m , where m is the number of agents exerting high effort. Furthermore, we demand that also the costs c_i of all agents are identical, i.e., $c = c_i$. Let x_i denote the probabilistic event of agent i succeeding in his subtask. In particular, we will study the following technology functions:

- 1) *AND technology*: The project is successful if *all* agents succeed in their subtask, i.e., $f(x) = \bigwedge_i x_i$.
- 2) *OR technology*: The project is successful if *at least one* agent succeeds in his subtask, i.e., $f(x) = \bigvee_i x_i$.
- 3) *MAJORITY technology*: The project is successful if *a majority* of the agents succeeds in their subtasks, i.e., $f(x)$ is 1 if at least $n/2$ agents succeed and 0 otherwise.

Audits: We are now ready to introduce the final component to our model. We assume that in addition to observing the overall outcome, the principal can *audit a specific agent's action*. Auditing agent i will reveal a_i , but comes at a cost κ .

With the possibility of audits, the principal may offer slightly different contracts to the agents. She still offers agent i a payment p_i based on the outcome. If agent i is not audited, he is paid p_i if the project as a whole succeeds and 0 otherwise, independent of his action. However, if he is audited and it turns out that he has not exerted high effort, i.e., $a_i = 0$, the principal does not have to pay him, independent of the project outcome. We say that the principal *contracts with agent i* if $p_i > 0$. The principal informs all agents about all payments p_i . We assume that the principal then chooses k out of the $l \leq n$ contracted agents uniformly at random to audit their contribution to the project. This yields a probability $\pi = k/l$ that agent i is audited. This audit probability is communicated to the agents. While such a uniform audit strategy is reasonable with anonymous technologies, it might be better with non-anonymous technologies to inflict a more sophisticated probability distribution with individual audit probabilities π_i .¹

Let μ_i be 1 if agent i is audited and 0 otherwise. For each agent, the principal stipulates a payment depending on the project outcome, the agent's action and μ_i . The payment is hence a function $p_i : O \times A_i \times \{0, 1\} \rightarrow \mathbb{R}^+$. Agent i 's utility depends on the expected payment p_i he receives², which in turn is a function of the success probability under the given technology $t(a)$, his action a_i and the probability that he is audited, and the cost $c_i(a_i)$ of the action: $u_i(a) = p_i(t(a), a_i, \mu_i) - c_i(a_i)$. We assume that the success function $t(a)$ is *strictly monotone* in the sense that $\forall i \in N, \forall a_{-i} \in A_{-i} : t(1, a_{-i}) > t(0, a_{-i})$. The principal's valuation of a successful project is given by a scalar $v > 0$. The valuation

¹If, for instance, it is vital for the project that agent i succeeds whereas agent j 's contribution is less important for the project success, the principal might be best off by auditing agent i ($\pi_i = 1$) for sure and agent j with a low probability ($\pi_j \ll 1$). However, we do not explore these issues in this work.

²Note that in the following, we will simply write p_i instead of $\mathbb{E}[p_i]$.

of a project failure is 0. The principal's utility can hence be expressed by

$$u(a) = v(t(a)) - \sum_i p_i(t(a), a_i, \mu_i) - k \cdot \kappa.$$

The principal's aim is to maximize u under the assumption that the agents are rational in the sense that they act so as to maximize their expected utility in turn.

Finally, let us introduce some notation. In this paper, we will analyze two shirker models, the *shirking without contribution model* \mathcal{M}^0 and the *shirking with contribution model* \mathcal{M}^1 . We will indicate the studied technology as a subscript (*AND* for *AND* technologies, *OR* for *OR* technologies, and *MAJ* for *MAJORITY* technologies). For example, \mathcal{M}_{MAJ}^0 refers to the no contribution model under a *MAJORITY* technology.

II. GENERAL OBSERVATIONS

We start our analysis with some general observations. First of all, since the principal can always choose to audit no agents at all, we have the following trivial fact.

Fact 2.1: The possibility of auditing can never be detrimental to the principal's utility.

We now want to study the impact of audits on the prices that need to be paid to the agents.

Theorem 2.2: If the principal decides to contract agent i , she will offer him a payment

$$p_i = \frac{c_i}{t(1, a_{-i}) - t(0, a_{-i})(1 - \pi)}.$$

The principal's utility is

$$u(a, v) = t(a) \cdot (v - \sum_i p_i) - k \cdot \kappa.$$

Agent i 's utility is

$$u_i(1, a_{-1}) = c_i \left(\frac{t(1, a_{-i})}{t(1, a_{-i}) - t(0, a_{-i})(1 - \pi_i)} - 1 \right)$$

if he is contracted and 0 otherwise.

Proof: Given the other agents' actions a_{-i} , agent i will exert high effort if and only if the contracted payment p_i implies that $u_i(1, a_{-i}) \geq u_i(0, a_{-i})$, i.e., $p_i(t(1, a_{-i}), \mu_i, 1) - c_i(1) \geq p_i(t(0, a_{-i}), \mu_i, 0) - c_i(0)$. We have $p_i \cdot t(1, a_{-i}) - c_i \geq p_i \cdot t(0, a_{-i})(1 - \pi_i) \Leftrightarrow p_i \geq c_i / (t(1, a_{-i}) - t(0, a_{-i})(1 - \pi))$. Thus, the principal will only offer contracts with $p_i = c_i / (t(1, a_{-i}) - t(0, a_{-i})(1 - \pi_i))$. Offering a smaller or larger amount of money is wasteful as it has no influence on the agent's decision. The principal and agent utilities then follow from our definitions. ■

Note that in anonymous technologies, an optimal payment is $p_i = c / (t_l - t_l(1 - \pi))$ where l is the number of contracted agents. This yields a principal utility $u(v) = t_l \cdot (v - lp_i) - k\kappa$ and an agent utility $u_i(a_i = 1) = t_l c / (t_l - t_{l-1}(1 - \pi)) - c$ if contracted and 0 otherwise. Since the payments are identical for all contracted agents, we usually omit the index i , i.e., $p = p_i$. Additionally, Theorem 2.2 reveals the fact that auditing only affects the optimal payment offered to an agent whom the principal wants to contract with. Thus,

for the principal, auditing a non-contracted agent would not decrease the payments needed to have this agent exert high effort. On the contrary, it would even increase the costs for her, because the extra audit would cost her κ .

Fact 2.3: It makes no sense to audit non-contracted agents.

This justifies the design decision already incorporated in our model that the principal only conducts audits among the contracted agents. The principal may increase an agent's chance π of being audited by carrying out more random audits. The higher π is, the less she has to pay agent i in order to contract with him. However, increasing π comes at a cost: the principal must pay the extra audits. Thus, it is often not optimal for the principal to simply audit all contracted agents, but rather a fraction or even no agent at all. This depends on the magnitude of the audit cost κ .

III. SHIRKING WITHOUT CONTRIBUTION

This section studies the case where a shirker—an agent who does not exert high effort—does not contribute anything at all to the project. Hence, we presume that the probability of an agent's subtask being completed successfully is zero if he shirks. If he exerts high effort, his subtask succeeds with a certain probability $\delta_i < 1$. We call this case as the *no contribution case* and refer to it as \mathcal{M}^0 . Recall that under optimal contracts, a rational agent exerts high effort ($a_i = 1$) if he is contracted and he shirks ($a_i = 0$) otherwise. In the following, we further assume that all agents have the same success probability, i.e., $\delta_i = \delta \forall i \in N$. We proceed by examining the no contribution case under different technologies.

A. AND Technology

A first observation is that in the *AND* technology, the probability of project success is always zero if less than all n agents exert high effort. Hence, it is only worthwhile for the principal to either not contract with any agent and therewith not to make any payments to the agents, or to contract with all n agents. Applying Theorem 2.2, we get that the payment which needs to be offered to an agent i if the principal wants to contract with all agents is $p \geq c/\delta^n$, since $t(0, \cdot) = 0$. Observe that this payment is independent of the number of audits, k , and the audit costs κ . As conducting audits would only cost the principal and in no way increase her utility, auditing does not yield any benefits in the *AND* technology. Moreover, note that the payments grow exponentially in the number of agents, as the project's success relies on each agent contributing. We can state the following.

Theorem 3.1: The optimal utility of the principal in \mathcal{M}_{AND}^0 is given by

$$u(v) = \max\{\delta^n \cdot v - n \cdot c, 0\}.$$

If $v > n \cdot c/\delta^n$, it is worthwhile to contract all agents, otherwise no agent is contracted. It is always best to carry out no audits ($k = 0$).

Proof: Since auditing can only decrease the payments needed to contract with an agent, and since $p \geq c/\delta^n$ is

independent of k , it is never worthwhile to invest any money in audits. Thus, $k = 0$ and hence also the total auditing cost $\kappa \cdot k = 0$. We only need to consider two cases: either all agents are contracted or none. This yields a principal utility $u(v)$ of $\delta^n \cdot v - n \cdot c$ if all agents are contracted and 0 if no agent is contracted. The principal always chooses the better of the two options, and hence, the claim follows. ■

B. OR Technology

In the *OR* technology, the project already succeeds with a probability δ if only one agent exerts effort. Each additional contracted agent diminishes the failure probability by a factor of $(1 - \delta)$. It turns out that the analysis of the *OR* technology is more difficult.

Theorem 3.2: In \mathcal{M}_{OR}^0 , if the principal contracts l agents and conducts $k \leq l$ audits, the lowest payment for which a contracted agent exerts effort is $p = c/((1 - (1 - \delta)^l) - (1 - (1 - \delta)^{l-1})(1 - k/l))$. The resulting principal utility is $u(v) = (1 - (1 - \delta)^l)(v - l \cdot p) - k \cdot \kappa$.

Proof: The project's success probability in the *OR* technology is given by $t_l = (1 - (1 - \delta)^l)$ where l is the number of contracted agents. Applying this success function t_l to Theorem 2.2 yields p and u . ■

Theorem 3.2 allows the principal to compute the optimal number of contracted agents l and the optimal number of audits $k \leq l$. If done without any optimizations, this requires $O(n^2)$ function evaluations. It is interesting to study how l and k relate to each other. Assume for a moment that l is fixed. Auditing an additional agent ($k := k + 1$) will reduce the payment p needed to contract an agent, but at the same time, it costs κ . For a given l , the optimal choice of k follows from the equality of this additional gain and the additional loss, i.e., when $\frac{l \cdot c}{(1 - (1 - \delta)^l) - (1 - (1 - \delta)^{l-1})(1 - (k+1)/l)} - \frac{l \cdot c}{(1 - (1 - \delta)^l) - (1 - (1 - \delta)^{l-1})(1 - k/l)} = \kappa$. We omit an analytical characterization of the optimal k but present some numerical results. Figure 1 (left) plots the effect of different audit costs κ on the principal's utility u . The figure illustrates that, as expected, for larger κ , the principal's utility declines. Moreover, for larger κ , in order to maximize the utility, fewer audits are conducted. Figure 1 (right) studies the effect of different agent costs c . For larger c , the utility is smaller, and it is worthwhile to audit more agents. Furthermore, we see that for many configurations of κ and c , the optimal number of audits is not simply 0 or l , but a value in between.

C. MAJORITY Technology

To conclude our analysis of the no contribution case, we consider the *MAJORITY* technology where *at least half* of the agents must be successful.

Theorem 3.3: In \mathcal{M}_{MAJ}^0 , if the principal contracts l agents and conducts $k \leq l$ audits, the lowest payment for which a contracted agent exerts effort is

$$p = \begin{cases} \frac{c}{\Upsilon_{n/2}^l - \Upsilon_{n/2}^{l-1}(1-k/l)} & \text{if } l > n/2 \\ \frac{c}{\Upsilon_{n/2}^l} & \text{if } l = n/2, \\ \infty & \text{otherwise} \end{cases}$$




Fig. 1. *Left*: Principal's utility (as a function of the number of audited agents k) for different κ s: $\kappa = .5$, $\kappa = 1$, $\kappa = 2$, $\kappa = 4$, $\kappa = 6$, $\kappa = 8$. For this experiment, we used $l = 10$, $v = 100$, $c = 1$, and $\delta = .9$. *Right*: Principal's utility for different agent cost c with $l = 10$, $v = 100$, $\kappa = 8$ and $\delta = .9$.

where

$$\Upsilon_{\alpha}^{\beta} = \sum_{j=\alpha}^{\beta} \binom{j}{n/2} \cdot \delta^j (1-\delta)^{\beta-j}.$$

This yields a principal utility of

$$u(v) = \Upsilon_{n/2}^l \cdot (v - l \cdot p) - k \cdot \kappa,$$

where the payments p and $\Upsilon_{\alpha}^{\beta}$ are as given above.

Proof: The probability that the project succeeds with exactly $j \in \{n/2, \dots, l\}$ successful agents out of the l contracted agents is binomially distributed. The *MAJORITY* technology's success function t_l can be calculated by summing over all the possibilities of success. Hence t_l is $\Upsilon_{n/2}^l$ if $l \geq n/2$ and 0 otherwise. Payment p and utility u then follow from Theorem 2.2. ■

IV. SHIRKING WITH CONTRIBUTION

In the second part of this paper, we consider a setting where an agent exerting low effort still contributes to the project. We call this setting *shirking with contribution case* and denote it by \mathcal{M}^1 . Agents playing $a_i = 1$ succeed with a certain probability δ and an agent playing $a_i = 0$ succeeds with a probability $\gamma < \delta$. In the following, we assume that the principal willing to implement the project makes a small inalienable payment p_{ϵ} to each agent in order to make him participate in the project. This payment is independent of the agent's decision whether to exert high or low effort. An agent's utility is given by $u_i(a) = p_i(t(a), a_i, \mu_i) - c_i(a_i) + p_{\epsilon}$ and the principal utility is $u(a) = v(t(a)) - \sum_i p_i(t(a), a_i, \mu_i) - k\kappa - np_{\epsilon}$. The minimal payment offered by the principal to make an agent exert high effort can be computed similarly to Theorem 2.2. We still assume that the agents seek to maximize their expected utility. However, in contrast to the no contribution case, the inalienable payments p_{ϵ} are subtracted from the principal's utility and added to the agents' utilities.

Corollary 4.1: If the principal contracts with agent i , she offers him an optimal payment of p_i , which adds to p_{ϵ} :

$$p_i = \frac{c_i}{t(1, a_{-i}) - t(0, a_{-i})(1 - \pi)}.$$

The principal utility u and the agent utility u_i are

$$u(a, v) = t(a) \cdot (v - \sum_i p_i) - k\kappa - np_{\epsilon} \quad \text{and}$$

$$u_i(1, a_{-i}) = \left(\frac{t(1, a_{-i})}{t(1, a_{-i}) - t(0, a_{-i})(1 - \pi_i)} - 1 \right) c_i + p_{\epsilon}$$

if contracted and 0 otherwise.

Note that, as agents exerting low effort can also succeed in their subtask, the success function $t(a)$ is different in the shirking with contribution case. For instance, an *AND* technology project may be completed successfully even if no agent exerts high effort. Moreover, observe that also in case of inalienable payments p_{ϵ} , it is never worthwhile to audit non-contracted agents. This is due to the fact that not contracted agents receive only the minimal payment p_{ϵ} —independently of their action. Being audited does not decrease these payments and therefore the additional audits only entail costs.

A. Case Study: Agency With 2 Agents

In order to initiate the study of the shirking with contribution case, we consider a very small project with two agents only. This setting already raises interesting questions. We first study the *AND* technology. If no agent exerts effort, the project success probability is $t_0 = \gamma^2$. It is $t_1 = \gamma \cdot \delta$ if one agent exerts effort and $t_2 = \delta^2$ if both agents exert effort. The principal has three possibilities. She can either contract none ($l = 0$), one ($l = 1$), or both ($l = 2$) agents. The optimal mechanism can be found by considering these cases in turn.

Case $l = 0$: The principal does not contract *any agent*, and hence all agents exert low effort. The principal's utility becomes $u^{k,0}(v) = t_0 v - k\kappa - 2p_{\epsilon}$. (Henceforth, we will use the notation $u^{k,l}(v)$ to denote the principal utility when k agents are audited and l agents contracted.) Since non-contracted agents are never subject to auditing, $k = 0$ and thus

$$u^{0,0}(v) = t_0 \cdot v - 2p_{\epsilon}.$$

Case $l = 1$: The principal contracts *one agent*. Both agents are paid p_ϵ and from Corollary 4.1, we know that one agent is paid an additional $p = c/(t_1 - (1-\pi)t_0)$ in case of success. The principal's utility is $u^{k,1}(v) = t_1(v-p) - k\kappa - 2p_\epsilon$. The principal can either audit none ($k = 0$) or one ($k = 1$) agent. This yields an auditing probability for the contracted agent of $\pi = 1$ or $\pi = 0$ respectively, and hence,

$$\begin{aligned} u^{0,1}(v) &= t_1 \cdot \left(v - \frac{c}{t_1 - t_0} \right) - 2p_\epsilon \\ u^{1,1}(v) &= t_1 \cdot v - c - \kappa - 2p_\epsilon. \end{aligned}$$

Case $l = 2$: The principal contracts *two agents*. Both agents are paid an additional $p = c/(t_2 - (1-\pi)t_1)$ on success. The principal's utility is $u^{k,2}(v) = t_2(v-2p) - k\kappa - 2p_\epsilon$. Either, we can audit none ($k = 0, \pi = 0$), one ($k = 1, \pi = 1/2$), or both agents ($k = 2, \pi = 1$). We have

$$\begin{aligned} u^{0,2}(v) &= t_2 \cdot \left(v - 2 \frac{c}{t_2 - t_1} \right) - 2p_\epsilon \\ u^{1,2}(v) &= t_2 \cdot \left(v - 2 \frac{c}{t_2 - (1/2)t_1} \right) - \kappa - 2p_\epsilon \\ u^{2,2}(v) &= t_2 \cdot v - 2c - 2\kappa - 2p_\epsilon. \end{aligned}$$

Observe that the principal's utility is linear in v . Furthermore, the gradient is t_l and hence does not depend on c , k or κ . For given costs c and κ , there is an optimal choice of k independent of the principal's valuation v . For an optimal k and given l , we can compute the relationship between κ and c . For instance, when contracting one agent, there is one transition point at $\kappa = (t_1/(t_1 - t_0) - 1) \cdot c$, i.e., auditing the contracted agent yields a higher utility than without any audits when the costs κ do not exceed $(\delta/(\delta - \gamma) - 1) \cdot c$. Otherwise the principal is better off by not auditing. When contracting two agents, there are two transition points at $\kappa = (t_2/(t_2 - t_1) - t_2/(t_2 - (1/2)t_1)) \cdot 2c$ and $\kappa = (t_2/(t_2 - (1/2)t_1) - 1) \cdot 2c$, where for audit costs smaller than $(\delta/(\delta - \gamma/2) - 1) \cdot 2c$, auditing two agents, for costs over $(\delta/(\delta - \gamma) - \delta/(\delta - \gamma/2)) \cdot 2c$, no agent, and in-between, one agent yields the highest principal utility.

As an example, let $c = 1$, $\gamma = 1/4$, $\delta = 3/4$, and $p_\epsilon = 0$. Then we have a success probability of $t_0 = 1/16$ if no agent exerts effort, $t_1 = 3/16$ if one agent exerts effort and $t_2 = 9/16$ if both agents exert effort. Refer to Figure 2 for an overview of the three different cases that occur in a two-player *AND* technology. It illustrates the fact that the optimal number of audits can be either one, two or none at all.

Let us briefly consider the two-player example in an *OR* technology. All formulas given for the utility functions $u^{k,l}$ also hold with an *OR* technology. The only difference concerns the success function t_l . If the principal contracts with no agent, the project succeeds with probability $t_0 = 1 - (1-\gamma)^2$. For $l = 1$ and $l = 2$, we have $t_1 = 1 - (1-\gamma) \cdot (1-\delta)$ and $t_2 = 1 - (1-\delta)^2$.

If κ is slightly smaller than the cost c , there are two transition points v_1^* and v_2^* , where the principal is best off by contracting with no agent for a valuation $v \leq v_1^*$, by contracting with one agent and conducting *one* audit for $v_1^* \leq v \leq v_2^*$ and by contracting with both agents and

auditing *both* agents for $v \geq v_2^*$. If κ is twice as large as c , there are also two transition points v_1^* and v_2^* . The principal is best off by contracting with no agent for a valuation $v \leq v_1^*$, by contracting with one agent and auditing *none* for $v_1^* \leq v \leq v_2^*$ and by contracting with both agents and conducting *one* audit for $v \geq v_2^*$. If κ is about a factor twenty larger than c , there are again two transition points v_1^* and v_2^* . This time, the principal is best off by contracting with no agent for a valuation $v \leq v_1^*$, by contracting with one agent and monitoring *none* for $v_1^* \leq v \leq v_2^*$ and by contracting with both agents and auditing *none* for $v \geq v_2^*$. This suggests that in an *OR* technology, there are always n transition points $v_1^* < \dots < v_n^*$ where it is optimal for a $v \in (v_l^*, v_{l+1}^*)$ to contract with l agents and audit a certain fraction k of the l contracted agents (cf Conjecture 4.10).

The two-player case already discloses some insights about how audits can be useful to a principal. It shows that audits facilitate a higher principal benefit in many cases. In certain configurations, it even enables the principal to carry out a project that otherwise would not prospect profit. We say a project is *profitable* if the principal's expected utility is positive.

Theorem 4.2: In the *AND* and *OR* technology, there are situations where a project is only profitable if the principal employs audits.

Proof: In a \mathcal{M}_{AND}^1 technology with two players, let $p_\epsilon = 4$, $\gamma = 1/4$, $\delta = 3/4$, $c = 1$, $p_\epsilon = 0.1$ and $\kappa = 1$. Let the principal have a valuation $v = 35$. Then the optimal number of contracts without audits is $l = 2$. However, this yields a negative utility of $u^{0,2}(35) = -0.3125$, whereas auditing both agents yields a positive utility $u^{2,2}(35) = 1.6875$. Creating a respective situation for *OR* technologies is straightforward. ■

B. Anonymous Technologies

We will now generalize our results to anonymous technologies with n agents and provide the corresponding results for two anonymous technologies, namely for the *AND* technology as well as for the *OR* technology.

Recall that in *anonymous* technologies, the success function only depends on the number of agents exerting high effort and not on the particular set of agents exerting high effort. For the *AND* technology for instance, we have $t(a) = t_m = \delta^m \gamma^{n-m}$ where m is the number of agents exerting high effort. Let l be the number of agents contracted by the principal. If the agents are rational decision makers, all contracted agents will exert high effort and thus $m = l$. Our aim is to give advice to the principal on how to choose the number l of contracts and the number k of audits given her valuation v and the parameters c , κ and n .

In an anonymous technology where the principal contracts with l out of n agents and conducts k audits among the l contracted agents uniformly at random, the minimal payment needed to convince a rational agent not to cheat is

$$p = \frac{c}{t_l - t_{l-1} \cdot (1 - \frac{k}{l})} \quad (1)$$

v	$c = 1, \kappa = 1$				$c = 2, \kappa = 1$				$c = 4, \kappa = 1$			
	k, l	$u^{k,l}(v)$	benefit	in %	k, l	$u^{k,l}(v)$	benefit	in %	k, l	$u^{k,l}(v)$	benefit	in %
1	0, 0:	-0.1375	0	0%	0, 0:	-0.1375	0	0%	0, 0 :	-0.1375	0	0%
5	0, 0:	0.1125	0	0%	0, 0:	0.1125	0	0%	0, 0 :	0.1125	0	0%
10	0, 2:	2.425	0	0%	0, 0:	0.425	0	0%	0, 0 :	0.425	0	0%
15	0, 2:	5.2375	0	0%	1, 2:	2.4375	0.2	8.9%	0, 0 :	0.7375	0	0%
20	0, 2:	8.05	0	0%	1, 2:	5.25	0.2	4.0%	0, 0 or 2, 2:	1.05	0	0%
25	0, 2:	10.8625	0	0%	1, 2:	8.0625	0.2	2.5%	2, 2 :	3.8625	2	107.4%
50	0, 2:	24.925	0	0%	1, 2:	22.125	0.2	0.9%	2, 2 :	17.925	2	12.6%

Fig. 2. Example of the benefits of auditing for the principal in the AND technology with two players, depending on her project valuation v where $\kappa = 1$, $\gamma = 1/4$, $\delta = 3/4$, and $p_\epsilon = 0.1$. In the first case ($c = 1, k = 1$), there is one transition between $v = 5$ and $v = 10$ where the principal's optimal choice switches from $k = 0, l = 0$ to $k = 0, l = 2$. However, for this case, auditing never brings any benefits. In the second case ($c = 2, k = 1$), there is also one transition between $v = 10$ and $v = 15$ where the principal's optimal choice switches from $k = 0, l = 0$ to $k = 1, l = 2$. Here, auditing one of the two contracted agents results in a benefit of 0.2 for a valuation of 15 or larger. In the third case ($c = 4, k = 1$), there is one transition at $v = 20$. If the principal's valuation of the project is lower than 20, her optimal choice is $k = 0, l = 0$, if it is higher, $k = 2, l = 2$ is best. In this case, auditing both contracted agents brings a benefit of 2 for a valuation of 15 or larger. As the absolute value of the audit benefit is constant for a certain k and l , the gain from auditing decreases in terms of percentage with v . Notice also that for a low valuation $v = 1$ the principal's utility is negative. This is because the basic payments p_ϵ made to all agents exceed the benefits from the project and hence the principal will not carry out the project.

and the principal's utility is

$$u^{k,l}(v) = t_l \cdot (v - lp) - k\kappa - np_\epsilon. \quad (2)$$

A first thing to notice is that $u^{k,l}$ is a linear function of v . Secondly, the gradient of $u^{k,l}(v)$ is independent of k . It depends only on n and the number of contracted agents l .

Theorem 4.3: In any anonymous technology under \mathcal{M}^1 , if the principal contracts with l agents, it is optimal for her to audit k^* agents, where

$$k^*(l) = \begin{cases} l & \text{for } \kappa \leq \frac{t_{l-1}c}{t_l} \\ 0 & \text{for } \kappa \geq \frac{t_l t_{l-1}}{(t_l - t_{l-1})^2} c \\ (1 + \sqrt{\frac{t_l}{t_{l-1}} \frac{c}{\kappa} - \frac{t_l}{t_{l-1}}}) l & \text{otherwise.} \end{cases}$$

Proof: The first and second derivative of $u^{k,l}(v)$ with respect to k are

$$\frac{d}{dx} u^{k,l}(v) = \frac{t_l t_{l-1} c l^2}{(t_l l - t_{l-1} l + t_{l-1} k)^2} - \kappa$$

and

$$\frac{d^2}{dx^2} u^{k,l}(v) = -\frac{t_l t_{l-1}^2 c l^2}{(t_l l - t_{l-1} l + t_{l-1} k)^3} < 0.$$

Since the second derivative is always negative, $u^{k,l}(v)$ has exactly one extremum, namely a maximum at the root of $\frac{d}{dk} u^{k,l}(v)$. Hence computing the root allows us to express the value of k which maximizes the principal's utility as a function k^* of l , in particular, $k^*(l) = (1 + \sqrt{\frac{t_l}{t_{l-1}} \frac{c}{\kappa} - \frac{t_l}{t_{l-1}}}) l$. As k can only take values in $(0, l)$, we must restrict the term inside parentheses to values in $(0, 1)$. If the term is larger than 1, the optimal feasible value for k is l . If it is negative, the optimal feasible value is 0. This gives the first two cases in the definition of k^* . ■

Theorem 4.3 immediately implies the following corollary.

Corollary 4.4: Audits never pay off in anonymous technologies if the audit costs κ are greater or equal to $(t_l t_{l-1}) / ((t_l - t_{l-1})^2)$ times the cost c that incur when an individual agent exerts high effort.

C. AND Technology

In the AND technology, from Equation (1) it follows that the principal pays $p_i = c / (\delta^l \gamma^{n-l} - \delta^{l-1} \gamma^{n-l+1} \cdot (1 - k/l))$ in order to contract with an agent i . The resulting expected principal utility is $u^{k,l}(v) = \delta^l \gamma^{n-l} v - l^2 c / (l\delta - l\gamma + k\gamma) - k\kappa - np_\epsilon$.

We observe that the gradient of $u^{k,l}$ is independent of k and, since $\gamma < \delta$, the gradient is maximal for $l = n$, namely δ^n , and minimal for $l = 0$, γ^n . Hence, we know that for a valuation v large enough, the principal will contract with all agents. We also see that $u^{0,0}(v) = \gamma^n v - np_\epsilon$ is the best choice for a v small enough, as the costs of np_ϵ are inevitable for any choice of l and k . Of course, if $\gamma^n v < np_\epsilon$, the principal's utility is negative. In that case, she would not carry out the entire project at all. Interestingly, it never pays off to contract with only a fraction of the agents. The principal utility is maximized if she either contracts with nobody and obviously conducts no audits or if she contracts with all agents and chooses an optimal number of audits.

Theorem 4.5: In \mathcal{M}_{AND}^1 , there exists a value v^* such that for any value $v \leq v^*$ it is optimal for the principal to contract with no agent (and thus audit no agent) and for any value $v \geq v^*$ it is optimal to contract with all n agents and audit k^* agents at random, where k^* is given by

$$k^* = \begin{cases} n & \text{for } \kappa \leq \frac{\gamma}{\delta} c \\ 0 & \text{for } \kappa \geq \frac{\delta\gamma}{(\delta-\gamma)^2} c \\ (1 + \sqrt{\frac{\delta c}{\gamma\kappa} - \frac{\delta}{\gamma}}) n & \text{otherwise} \end{cases}$$

and v^* is $v^* = \frac{1}{\delta^n - \gamma^n} \left(\frac{\delta c n^2}{\delta n - \gamma n + \gamma k^*(n)} + \kappa \cdot k^*(n) \right)$.

In order to prove Theorem 4.5, we will first apply the insight gained in Theorem 4.3 to the AND technology, namely that for any number of contracted agents l , there is an optimal audit number $k^*(l)$ in Lemma 4.6. Second, we show that at the intersection point of $u^{0,0}$ and $u^{k^*(n),n}$, no other choice of l yields a higher principal utility even with the corresponding optimal $k^*(l)$ in Lemma 4.7. Finally, we show that this is

sufficient. The following lemma is an immediate implication of Theorem 4.3.

Lemma 4.6: If the principal contracts with l agents in \mathcal{M}_{AND}^1 , it is optimal for her to conduct k^* audits given by

$$k^*(l) = \begin{cases} l & \text{for } \kappa \leq \frac{\gamma}{\delta}c \\ 0 & \text{for } \kappa \geq \frac{\delta\gamma}{(\delta-\gamma)^2}c \\ (1 + \sqrt{\frac{\delta c}{\gamma\kappa} - \frac{\delta}{\gamma}})l & \text{otherwise.} \end{cases}$$

Lemma 4.7: At the intersection point \bar{v} of $u^{0,0}(v)$ and $u^{k^*(n),n}(v)$ in \mathcal{M}_{AND}^1 , no other choice of l and k yields a higher principal utility, i.e., if $u^{0,0}(\bar{v}) = u^{k^*(n),n}(\bar{v})$ then $\forall l, k : u^{0,0}(\bar{v}) = u^{k^*(n),n}(\bar{v}) \geq u^{k,l}(\bar{v})$.

Proof: Computing the intersection point \bar{v} of $u^{0,0}(v)$ and $u^{k^*(n),n}(v)$ gives $\bar{v} = \frac{1}{\delta^n - \gamma^n} \cdot \left(\frac{\delta c n^2}{\delta n - \gamma n + \gamma k^*(n)} + \kappa \cdot k^*(n) \right)$. As shown in Lemma 4.6, for each choice of l , there is an optimal choice of k , namely $k^*(l)$. Thus it suffices to show that $u^{0,0}(\bar{v}) \geq u^{k^*(l),l}$ holds for any feasible l , since $\forall k, l, v : u^{k^*(l),l}(v) \geq u^{k,l}(v)$. $u^{0,0}(\bar{v}) \geq u^{k^*(l),l}$
 $\Leftrightarrow \frac{1}{\gamma^{n-l}\delta^l - \gamma^n} \cdot \left(\frac{\delta c l^2}{\delta l - \gamma l + \gamma k^*(l)} + \kappa \cdot k^*(l) \right) \geq \frac{1}{\delta^n - \gamma^n} \cdot \left(\frac{\delta c n^2}{\delta n - \gamma n + \gamma k^*(n)} + \kappa \cdot k^*(n) \right)$. Let $\alpha = (1 + \sqrt{\frac{\delta c}{\gamma\kappa} - \frac{\delta}{\gamma}})$. Then we have $k^*(l) = \alpha l$.

$$\begin{aligned} \frac{\frac{\delta c l}{\delta - \gamma + \gamma\alpha} + \kappa\alpha l}{\frac{\delta c n}{\delta - \gamma + \gamma\alpha} + \kappa\alpha n} &\geq \frac{\gamma^{n-l}\delta^l - \gamma^n}{\delta^n - \gamma^n} \\ \frac{l}{n} &\geq \frac{\gamma^{n-l}\delta^l - \gamma^n}{\delta^n - \gamma^n} \\ l \left(\frac{\delta}{\gamma} \right)^n + n &\geq n \left(\frac{\delta}{\gamma} \right)^l + l \end{aligned} \quad (3)$$

In order to see that inequality (3) holds, let us look at the difference function

$$f(l) = l \left(\frac{\delta}{\gamma} \right)^n + n - n \left(\frac{\delta}{\gamma} \right)^l - l.$$

Observe that f has two roots at $l = 0$ and $l = n$. The second derivative of f is

$$\frac{d^2}{dx^2} f(l) = -n \left(\frac{\delta}{\gamma} \right)^l \left(\ln \frac{\delta}{\gamma} \right)^2.$$

Since $\delta > \gamma$ and $0 \leq l \leq n$, the second derivative of f is strictly negative for $n > 0$. Thus, the first derivative is strictly descending and consequently, the gradient of f is strictly greater at $v = 0$ than at $v = n$. As $f(0) = 0$ and $f(n) = 0$, it must hold that $f(l) > 0$ for any l in $(0, n)$, and the claim follows. ■

Equipped with Lemma 4.7, we are finally ready to conclude the proof of Theorem 4.5.

Proof: Note that the transition point indicated in Theorem 4.5, v^* , is exactly the intersection point of $u^{0,0}(v)$ and $u^{k^*(n),n}(v)$. We conclude by showing that no choice of k and l yields a higher utility than $u^{0,0}(v)$ for a $v < v^*$ and $u^{k^*(n),n}(v)$ for a $v > v^*$.

Recall that $u^{k,l}(v)$ is linear and that for a sufficiently small v , choosing $l = k = 0$ is optimal, i.e., $u^{0,0}(v) = \max_{l,k} u^{k,l}(v)$, and for a sufficiently large v , choosing $l = n$ and $k = k^*(n)$ is optimal, i.e., $u^{k^*(n),n}(v) = \max_{l,k} u^{k,l}(v)$. This is because the gradient of $u^{k,l}(v)$ is maximal for $l = n$ and minimal for $l = 0$. We now show that no choice of k and l yields a higher utility than $u^{0,0}(v)$ for a $v < v^*$ and $u^{k^*(n),n}(v)$ for a $v > v^*$. Assume for the sake of contradiction that there is a k' , a $l' > 0$ and a $v' < v^*$ such that $u^{k',l'}(v') \geq u^{0,0}(v')$. As $u^{k',l'}(v)$ is linear and its gradient is larger than the gradient of $u^{0,0}(v)$, it holds that $u^{k',l'}(v^*) > u^{0,0}(v^*)$. This is a contradiction to Lemma 4.7. Assume for the sake of contradiction that there is a k' , a $l' < n$ and a $v' > v^*$ such that $u^{k',l'}(v') \geq u^{k^*(n),n}(v')$. As the gradient of $u^{k',l'}(v)$ is lower than the gradient of $u^{k^*(n),n}(v)$, it holds that $u^{k',l'}(v^*) > u^{k^*(n),n}(v^*)$. Contradiction. ■

Knowing the number of agents to contract and the number of audits, we are finally interested in the question of how much the principal can actually benefit from employing audits.

Theorem 4.8: With an *AND* technology, the benefit of auditing is $u^{k^*(n),n}(v) - \max\{u^{0,0}(v), u^{0,n}(v)\}$ for $v > v^*$ and 0 otherwise where v^* is the intersection point of $u^{k^*(n),n}(v)$ and $u^{0,0}(v)$.

Proof: From Theorem 4.5, we know that for a $v \leq v^*$, the optimal k and l is $k = l = 0$. In this case, the principal does not conduct any audits and hence does not profit from the possibility of employing audits. Babaioff et al. have shown in [2] that without audits, it is optimal for the principal to either contract with no agent or with all agents (in an *AND* technology). With audits, choosing $k = k^*(n)$ and $l = n$ is optimal for a $v > v^*$ (Theorem 4.5). By subtracting the optimal gain without audits, $\max\{u^{0,0}(v), u^{0,n}(v)\}$, from the utility achieved with this optimal choice, we get the auditing surplus. ■

Refer to Figure 3 for a plot of the benefit from auditing compared to the absolute revenue in a technology with two agents. Once v grows larger than the transition point, the benefit increases and reaches a maximum level at the transition point without audits. Obviously, the benefit surplus compared to the absolute revenue in % then decreases with growing valuation.

D. OR Technology

The success function of the anonymous *OR* technology is $t_m = 1 - (1 - \gamma)^{n-m}(1 - \delta)^m$, where m is the number of agents exerting high effort. Applying Theorem 4.3 to the *OR* technology gives the following result.

Lemma 4.9: In \mathcal{M}_{OR}^1 with n agents, it holds that if the principal contracts with l agents, it is optimal for her to audit k^* agents given by

$$k^*(l) = \begin{cases} l & \text{for } \kappa \leq \phi_l^{-1}c \\ 0 & \text{for } \kappa \geq \frac{\phi_l}{(\phi_l - 1)^2}c \\ (1 + \sqrt{\frac{\phi_l c}{\kappa} - \phi_l})l & \text{otherwise} \end{cases}$$

where $\phi_l = \frac{1 - (1 - \delta)^l (1 - \gamma)^{n-l}}{1 - (1 - \delta)^{l-1} (1 - \gamma)^{n-l+1}}$.




Fig. 3. Benefit from audits with growing v in \mathcal{M}_{AND}^1 (left) and in \mathcal{M}_{OR}^1 (right) with two agents where $c = 1$, $\kappa = .8$, $p_\epsilon = .1$, $\gamma = 1/4$, $\delta = 3/4$.

Proof: Since the *OR* technology is anonymous, the formula for k^* follows directly from Theorem 4.3 by plugging in the *OR* technology success function $t_l = 1 - (1 - \gamma)^{n-l} (1 - \delta)^l$. ■

We presume that the *OR* technology has n transition points on the value of v , where it is optimal for the principal to increase the number of contracts by one with increasing v and to audit k^* agents. Unfortunately, we did not succeed in proving this property completely. However, we can provide a proof under the assumption that the following conjecture holds.

Conjecture 4.10: Let v_l^* be the transition point where choosing $k = k^*(l-1)$, $l = l-1$ yields the same principal utility as $k = k^*(l)$, $l = l$. With *OR* technologies, it holds for all $l \in \{1, \dots, n-1\}$ that v_l^* is smaller than v_{l+1}^* .

A proof may be similar to the one of Lemma 4.7: Compute v_l^* and v_{l+1}^* by solving $u^{k^*(l),l}(v_l^*) = u^{k^*(l-1),l-1}(v_l^*)$ for v_l^* and $u^{k^*(l+1),l+1}(v_{l+1}^*) = u^{k^*(l),l}(v_{l+1}^*)$ for v_{l+1}^* respectively. Then define the difference function $f(l) = v_{l+1}^* - v_l^*$ and show that it is always positive for all $l \in \{1, \dots, n-1\}$.

We did not succeed in the last step of the proof analytically, because the maths are not as advantageous as in the *AND* technology. The reason for this is mainly that in the *AND* technology, the term t_l/t_{l-1} dissolves to δ/γ and thus it is independent of l and n . In the *OR* technology, however, this term still depends on l and n . We conducted simulations of the difference function $f(l)$ for a variety of configurations and could not find a counter-example to our conjecture. Another fact that supports our conjecture is that an anonymous *OR* technology has n transitions in the purely-hidden-actions case as well as in the observable-actions case [2]. Since the case with audits could be interpreted as a combination of these two cases, it seems likely that there are also n transitions in this case.

Lemma 4.11: Let $k^*(l)$ be the function defined in Lemma 4.9. If Conjecture 4.10 holds then for any *OR* technology with n agents, there exist n transition points $v_1^*, v_2^*, \dots, v_n^*$ where for a valuation $v \in (v_{l-1}^*, v_l^*)$, it is optimal for the principal to contract with l agents and audit $k^*(l)$ agents out of the l contracted agents. For a valuation $v \leq v_1^*$, choosing $k = l = 0$ and for a $v \geq v_n^*$, choosing $k = k^*(n)$, $l = n$ is optimal.

Proof: If Conjecture 4.10 holds, the intersection point of a utility function $u^{k^*(l),l}$ with $u^{k^*(l-1),l-1}$ is reached with a smaller valuation v than the intersection point with

$u^{k^*(l+1),l+1}$. Moreover, $v_1^* < v_2^* < \dots < v_n^*$. Since the utility functions are linear in v and the gradient of any $u^{k,l+1}(v)$ is larger than the gradient of any $u^{k,l}(v)$, the utility $u^{k^*(l),l}(v)$ is higher than $u^{k^*(l+1),l+1}(v)$ in the interval $(v_l^*, v_l^* + 1)$. Similarly, we see that $u^{k^*(l-1),l-1}(v)$ is lower than $u^{k^*(l),l}(v)$ in $(v_l^*, v_l^* + 1)$. Applying this argument recursively for all $l \in \{l+2, l+3, \dots, n\}$ and all $l \in \{1, \dots, l-2\}$, we get that the largest possible utility lies in the interval $(v_l^*, v_l^* + 1)$, $u^{k^*(l),l}(v)$. Furthermore, for a sufficiently small valuation, choosing $k = 0, l = 0$ yields the highest utility since the principal pays only the minimal payments p_ϵ and still has a chance of success. (If the minimal payments p_ϵ exceed the principal's valuation, her utility is negative, but still as large as possible.) Thus, if the principal's valuation is in the interval $(0, v_1^*)$, choosing $k = l = 0$ yields the best utility possible, in any interval $(v_l^*, v_l^* + 1)$ where $l \in \{1, \dots, n-1\}$, choosing $k = k^*(l)$, $l = l$ is best and in (v_n^*, ∞) , $k = k^*(n)$, $l = n$ is best. ■

We have provided guidelines for the principal on how to maximize her revenue in an *OR* technology. For a particular project, where the parameters n , δ , γ , κ , c and p_ϵ are known, she could e.g., compute $l^* = \arg \max_l u^{k^*(l),l}$ and then contract with l^* agents and audit $k^*(l^*)$ out of the contracted agents. Additionally, we would like to know how much she can actually benefit from employing audits.

Corollary 4.12: Let v_l^* denote the intersection point of $u^{k^*(l),l}(v)$ and $u^{k^*(l-1),l-1}(v)$. If Conjecture 4.10 holds, the benefit of audits in an *OR* technology, is 0 if $v \leq v_1^*$, it is $u^{k^*(l),l}(v) - \max_{l \in \{0, \dots, n\}} u^{0,l}(v)$ if $v \in (v_l^*, v_{l+1}^*)$ and $u^{k^*(n),n}(v) - \max_{l \in \{0, \dots, n\}} u^{0,l}(v)$ if $v \geq v_n^*$.

Figure 3 shows the benefit from auditing compared to the absolute revenue with two agents. One can nicely observe the two transition points where with growing valuation, the benefit first increases and then decreases again until the next transition point is reached.

E. MAJORITY Technology

In the contribution case, the *MAJORITY* technology is quite complex as there is also a chance of success if less than half of the agents exert high effort. To complete our analysis on the effects of allowing the principal to employ audits, we provide the principal with the optimal payments and her resulting utility given that she contracts with l agents.

Theorem 4.13: In \mathcal{M}_{MAJ}^1 , if the principal contracts l agents and conducts $k \leq l$ audits,

the optimal payment to contract with an agent is $p = 1/(\sum_{j=n/2}^n \sum_{\max\{0, l-n+j\} \leq r \leq l} \Psi_{r, j-r} - (1 - \frac{k}{l}) \cdot \sum_{j=n/2}^n \sum_{\max\{0, l-1-n+j\} \leq r \leq l-1} \Psi_{r, j-r}) \cdot c$ where

$$\Psi_{r,s} = \binom{l}{r} \cdot \delta^r (1-\delta)^{l-r} \cdot \binom{n-l}{s} \cdot \gamma^s (1-\gamma)^{n-l-s}.$$

This yields a principal utility of

$$u(v) = \left(\sum_{j=n/2}^n \sum_{\max\{0, l-n+j\} \leq r \leq l} \Psi_{r, j-r} \right) \cdot (v-l \cdot p) - k\kappa,$$

where the payments p and $\Psi_{r,s}$ are as given above.

Proof: $\Psi_{r,s}$ represents the probability that exactly r contracted agents and exactly s non-contracted agents succeed. The *MAJORITY* technology's success function t_l can be calculated by summing over all the possibilities of success, i.e., the probabilities of the project succeeding with exactly $j \in \{n/2, \dots, l\}$ successful agents, which in turn are composed of all feasible combinations of r and s such that $r + s = j$. Payment p and utility u then follow from Theorem 2.2. ■

V. CONCLUSION

Many open distributed systems have to deal with the problem of hidden actions, i.e., actions or contributions of (often anonymous) participants that are not observed directly. If one can, however, reveal certain participants' actions by conducting audits, the coordinator may improve the cooperation and therewith the efficiency of such systems. The cost of auditing and the efficiency gain constitute an optimization problem. In both a model where shirkers contribute and a model where they do not contribute, there are general solutions on how to optimally choose the number of audits, the number of contracts and the payments to the participants in a combinatorial principal-agent setting with *AND*, *OR*, or *MAJORITY* technologies. It turns out that with increasing principal valuation, a system with the additional leverage of audits often exhibits similar transitions as in settings without audits in terms of the optimal choices for achieving the maximal principal utility. However, with audits, these transitions occur earlier than without audits. Sometimes, employing audits makes a project profitable which could otherwise not be carried out.

It would be interesting to test our results "in the wild"—our paper has studied anonymous technologies in theory only and using a simplified model. Another natural extension would be to study audits with non-anonymous forms of interaction. In such technologies, the question of which *set* of agents to audit arises. Furthermore, we ask: Can we devise general algorithms to solve the principal's optimization problem for arbitrary technologies? And what is the computational complexity?

ACKNOWLEDGMENTS

The authors would like to thank Markus Kern from the Max-Planck-Institute in Heidelberg and Alban Fischer from the Chair of Strategic Management and Innovation at ETH Zurich for interesting discussions.

REFERENCES

- [1] D. P. Anderson. BOINC: A System for Public-Resource Computing and Storage. In *Proc. 5th IEEE/ACM International Workshop on Grid Computing.*, pages 4–10, 2004.
- [2] M. Babaioff, M. Feldman, and N. Nisan. Combinatorial Agency. In *Proc. 7th ACM Conference on Electronic Commerce (EC)*, pages 18–28, 2006.
- [3] M. Babaioff, M. Feldman, and N. Nisan. Mixed Strategies in Combinatorial Agency. In *Proc. 2nd International Workshop on Internet and Network Economics (WINE)*, pages 353–364, 2006.
- [4] H. R. Coase. The Nature of the Firm. *Economica*, 4:386–405, 1937.
- [5] M. A. Crew. *Theory of the Firm*. Longman, 1975.
- [6] J. R. Douceur and T. Moscibroda. Lottery Trees: Motivational Deployment of Networked Systems. In *Proc. 2007 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM)*, pages 121–132, 2007.
- [7] W. Du, J. Jia, M. Mangal, and M. Murugesan. Uncheatable Grid Computing. In *Proc. 24th International Conference on Distributed Computing Systems (ICDCS)*, 2004.
- [8] R. Eidenbenz, Y. A. Oswald, S. Schmid, and R. Wattenhofer. Manipulation in Games. In *Proc. 18th International Symposium on Algorithms and Computation (ISAAC)*, Sendai, Japan, Springer LNCS 4835, December 2007.
- [9] K. Eisenhardt. Agency Theory: An Assessment and Review. *Academy of Management Review*, 14(1):57–74, 1989.
- [10] Y. Emek and M. Feldman. Computing Optimal Contracts in Series-Parallel Heterogeneous Combinatorial Agencies. In *Working Paper*, 2008.
- [11] M. E. Enzle and S. C. Anderson. Surveillant Intentions and Intrinsic Motivation. *J. Personality and Social Psychology*, 64:257–266, 1993.
- [12] A. Falk and M. Kosfeld. Distrust – The Hidden Cost of Control. In *IZA Discussion Paper, IZA DP No. 1203*, 2004.
- [13] E. Fehr, A. Klein, and K. M. Schmidt. Fairness and Contract Design. *Econometrica*, 75(1):121–154, 2007.
- [14] E. Fehr and K. M. Schmidt. Fairness and Incentives in a Multi-task Principal-Agent Model. *Scand J. of Economics*, 106(3):453–474, 2004.
- [15] I. Foster and C. Kesselman. *The Grid: Blueprint for a New Computing Infrastructure*. Elsevier, 2004.
- [16] P. Golle and I. Mironov. Uncheatable Distributed Computations. In *Proc. 2001 Conference on Topics in Cryptology (CT-RSA)*, pages 425–440, 2001.
- [17] S. J. Grossman and O. D. Hart. An Analysis of the Principle Agent Problem. *Econometrica*, 51(1):7–46, 1983.
- [18] B. Holmström. Moral Hazard and Observability. *The Bell Journal of Economics*, 10(1):74–91, 1979.
- [19] D. Hughes, G. Coulson, and J. Walkerdine. Free Riding on Gnutella Revisited: The Bell Tolls? *IEEE Distributed Systems Online*, 6(6), 2005.
- [20] M. Jensen and W. Meckling. Theory of the Firm: Managerial Behavior, Agency Costs, and Ownership Structure. *J. Financial Economics*, 3(4):305–360, 1976.
- [21] M. Kuhn, S. Schmid, and R. Wattenhofer. Distributed Asymmetric Verification in Computational Grids. In *Proc. 22nd IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, 2008.
- [22] D. Monderer and M. Tennenholtz. k-Implementation. In *Proc 4th ACM Conference on Electronic Commerce (EC)*, pages 19–28, 2003.
- [23] N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [24] C. H. Papadimitriou. Computational Aspects of Organization Theory. In *Proc. 4th Annual European Symposium on Algorithms (ESA)*, 1996.
- [25] C. H. Papadimitriou. Algorithms, Games, and the Internet. In *Proc. ACM Symposium on Theory of Computing (STOC)*, pages 749–753, 2001.
- [26] T. S. Pittman, M. E. Davey, K. A. Alafat, K. V. Wetherill, and N. A. Kramer. Informational Versus Controlling Verbal Rewards. *J. Personality and Social Psychology Bulletin*, 6:228–233, 1980.
- [27] R. Plant and R. M. Ryan. Intrinsic Motivation and the Effects of Self-Consciousness, Self-Awareness, and Ego-Involvement: An Investigation of Internally Controlling Styles. *J. Personality*, 53:435–449, 1985.
- [28] P. W. Schmitz. Workplace Surveillance, Privacy Protection, and Efficiency Wages. *J. Labour Economics*, 12:727–738, 2005.
- [29] D. Szajda, B. Lawson, and J. Owen. Hardening Functions for Large Scale Distributed Computations. In *Proc. 2003 IEEE Symposium on Security and Privacy*, 2003.