Reassignment-Based Strategy-Proof Mechanism for Interdependent Task Allocation with Private Costs and Execution Failures

Ayman Ghoneim[1,2,*] **and Jussi Rintanen**[1]
The Australian National University[1] and NICTA[2†]
Canberra ACT, Australia

Abstract

In this study, we consider a task allocation model with interdependent tasks, where tasks are assigned based on what agents report about their privately known capabilities and costs. Since selfish agents may strategically misreport their private information in order to increase their payments, mechanism design is used to determine a payment schema that guarantees truthful reporting. Misreported information may cause execution failures, creating interdependencies between the agents' valuations. For this problem, efficient and strategy-proof mechanisms have not been proposed yet. In this study, we show that such mechanisms exist if the failing tasks are reassigned, in addition, individual rationality and center rationality are obtained. Then, we extend the model to consider agents who have limited resources, and show that the center rationality property is lost.

1 Introduction

Task allocation is an important and challenging problem that occurs in various real-life applications, ranging from construction, service providing, to computing and research projects. Adopting a general model, a center wants to assign some tasks to a number of self-interested agents, where each agent has its own private information (i.e., type) that describes its abilities and costs for executing tasks. Given that the center aims for an efficient assignment (i.e., one that maximizes the social welfare) and provides payments to the agents, agents may strategically misreport their types in order to increase their payments. Thus, mechanism design is used to determine the payments that guarantee truthful reporting.

In this study, we consider the interdependent task allocation (ITA) problem, where tasks may fail during the execution because of the agents' strategically misreported information (i.e., agents claim the ability to perform tasks that they cannot perform). This model of failures is suitable when assuming selfish agents, and for mimicking the one-shot interaction situations in which agents don't care much about future implications (e.g., reputation, future opportunities). Given the interdependencies between tasks, an agent may not be able to execute its assigned tasks if their predecessor tasks have failed. This implies that an agent's *actual value* of its assigned tasks may depend on other agents' *actual types*, and that agents in such settings have *interdependent valuations*. When valuations are interdependent, mechanisms that achieve the strongest and most preferable form of truthfulness in dominant strategy (i.e., strategy-proof) have *not* been proposed yet for *any* domain (see Section 5).

In this study, we *prove* that it is impossible for an efficient mechanism to achieve strategy-proofness using a single allocation round, even if agents have sufficient resources. Then, we *contribute* a novel efficient mechanism that achieves strategy-proofness by using multiple allocation rounds (i.e., reassign the failing tasks). Finally, we extend the ITA model to consider agents with limited resources, and *prove* that the center rationality property is lost. In the next section, we formulate the task allocation problem as a mechanism design problem. In Sections 3 and 4, we propose the reassignment mechanism and discuss limited resources. Section 5 discusses related work, and finally, we conclude the study and discuss future work in Section 6.

2 Task Allocation and Preliminary Concepts

Basic Model. Assume a center that has a set $T = \{t_1, \ldots, t_m\}$ of m tasks. There are *predefined* interdependencies (i.e., an ordering) between these tasks, where some tasks can't be executed unless their predecessor tasks were executed successfully. Thus, each task t may have a set of successor tasks $t_{\succ}$ and a set of predecessor tasks $t_{\prec}$. The center gains a reward $R(t)$ (e.g., a market value) for each successful task t. The center wants to allocate the tasks to a set α of n self-interested agents, where each agent has its own private information (i.e., type) and knows nothing about other agents' types. The type $\theta_i = \langle T_i; C_i(t), \forall t \in T_i \rangle$ of agent i consists of: *1.* the set of tasks $T_i \subseteq T$ that the agent can perform, and *2.* the cost $C_i(t)$ for which the agent can execute each task $t \in T_i$.

Outcome. The center wants to determine an assignment

---

*Author to whom correspondence should be addressed. Email: ayman.ghoneim@anu.edu.au

†NICTA is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy and the Australian Research Council through the ICT Centre of Excellence program.

(i.e., outcome) $o = \{(t_1, i), (t_2, j), \ldots\}$, where each pair indicates the agent who is assigned a certain task, e.g., (t_1, i) means that agent i is assigned t_1. Under an outcome o, agent i is assigned the tasks in $T_i(o) = \{t_k | (t_k, i) \in o\}$, and $T_A(o) = \bigcup_{i\in\alpha} T_i(o)$ is the set of assigned tasks. An assignment may *not* contain all the offered tasks by the center in T, i.e., a task t and its successor tasks $t_{\succ}$ will not be assigned if no agent reported the ability to perform task t.

Tentative Values and Efficiency. The tentative value of agent i of an outcome o is $v_i(o, \theta_i) = -\sum_{t\in T_i(o)} C_i(t)$. The center's tentative value of an outcome o is $V(o) = \sum_{t\in T_A(o)} R(t)$. Given an outcome o, its social welfare - considering the center and the agents - is $SW(o) = V(o) + \sum_{i\in\alpha} v_i(o, \theta_i)$. Alternatively, the social welfare $SW(o)$ can be viewed as the summation of the social welfare of each assigned task in o, i.e., $SW(o) = \sum_{t\in T_A(o)} SW(t)$, where $SW(t) = R(t) - C_i(t)$ is the social welfare from assigning task t to agent i. Based on the vector $\theta = (\theta_1, \ldots, \theta_n)$ of the agents' reported types, the center will determine an efficient outcome $o_d \equiv o_d(\theta)$ from the set O of all possible outcomes.

Definition 1. *The determined outcome o_d is* efficient *if o_d maximizes the social welfare, i.e., $o_d = argmax_{o\in O} SW(o)$, and $SW(o_d) \geq 0$.*

Under o_d, each task t is simply assigned to agent i who can perform it for the cheapest cost (i.e., highest $SW(t)$), given that the predecessor tasks $t_{\prec}$ of t are assigned.

Utilities and Mechanism Design. Given the determined efficient outcome o_d, the center pays each agent i a payment $p_i(o_d)$ for its contributions in o_d. Assuming quasi-linear utilities, the utility of agent i is $u_i(o_d, \theta_i) = v_i(o_d, \theta_i) + p_i(o_d)$, while the center's utility is $U(o_d, \theta) = V(o_d) - \sum_{i\in\alpha} p_i(o_d)$. To guarantee the efficiency of o_d, the center must propose a payment schema $p_i(o_d)$ for each agent i that guarantees that the agent reports its private type truthfully. Clearly, this is a mechanism design problem [Mas-Colell *et al.*, 1995]. We will focus our attention here on *direct revelation* (DR) mechanisms, where an agent reports *all* its private information to the center that determines o_d and organizes payments to the agents. The revelation principle states that the properties of any mechanism can be replicated by a DR mechanism, and thus, any obtained results here immediately generalize to other indirect mechanisms. The mechanism needs *primarily* to establish truthfulness under some solution concept (Definition 2), either in *dominant strategies* (i.e., *strategy-proof*) or in *ex-post incentive compatibility*. Dominant strategy implementation is the strongest and most preferable solution concept, as it ensures that an agent reports truthfully irrespective of other agents' behavior.

Definition 2. *Given a true type θ_i of agent i, a strategically misreported type $\theta_i^{'}$ of agent i, a vector of reported types $\theta_{-i} = (\theta_1, \ldots, \theta_{i-1}, \theta_{i+1}, \ldots, \theta_n)$ of other agents except agent i, an outcome o_d that is determined if agent i reports θ_i, and an outcome $o_d^{'}$ that is determined if agent i reports $\theta_i^{'}$, a DR mechanism achieves* truthfulness *in*

Dominant Strategy: *For any agent i, reporting truthfully is always an optimal strategy regardless of whether other agents are reporting truthfully or not, i.e., $\forall i \in \alpha$, $u_i(o_d, \theta_i) \geq u_i(o_d^{'}, \theta_i)$ for any reported θ_{-i}.*

Ex-Post Incentive Compatibility: *For any agent i, reporting truthfully is always an optimal strategy given that other agents are reporting truthfully, i.e., $\forall i \in \alpha$, $u_i(o_d, \theta_i) \geq u_i(o_d^{'}, \theta_i)$ given that θ_{-i} holds the true types of other agents.*

DR mechanisms are preferred to possess other properties such as *individual rationality* and *center rationality*.

Definition 3. *A DR mechanism is* individually rational *if for every truthful agent i, its participation guarantees it a non-negative utility (i.e., $u_i(o_d, \theta_i) \geq 0$) given any $o_d \in O$.*

Definition 4. *A DR mechanism is* center rational *if in the truth-telling equilibrium, the center has a non-negative utility (i.e., $U(o_d, \theta) \geq 0$) given any outcome $o_d \in O$.*

Strategic Misreporting. Recalling the type $\theta_i = \langle T_i; C_i(t), \forall t \in T_i\rangle$ of agent i, agent i may increase its utility by strategically misreporting its type to the center in the following three ways: *1.* over-report its ability to perform more tasks than its actual ability (i.e., over-report $T_i^{'} \supset T_i$), this implies a larger set of outcomes $O^{'} \supset O$ from which the center will determine the problem's outcome; *2.* under-report its ability to perform tasks than its can actually perform (i.e., under-report $T_i^{'} \subset T_i$), this implies a smaller set of outcomes $O^{'} \subset O$; and *3.* misreport different costs for performing tasks than the actual costs (i.e., misreport cost $C_i^{'}(t) \neq C_i(t)$ for any task $t \in T$), this implies that the same assignments in $O^{'}$ and O may correspond to different social welfare.

Failures, Executed Outcome and Actual Values. Given the possibility that agents may over-report, we define a *failure point* as a task that wasn't executed successfully. Given any possible failing task, all its successor tasks will *not* be executed. We denote o_e as the part of the determined outcome o_d that was successfully executed, $T(o_e)$ as the set of successful tasks, and $T_i(o_e)$ as the set of successful tasks executed by agent i. Given the possibility that the unexecuted tasks may include tasks that belong to agent i, the *actual* value of agent i is $v_i(o_e, \theta_i) = -\sum_{t\in T_i(o_e)} C_i(t)$, which may differ from its tentative value $v_i(o_d, \theta_i) = -\sum_{t\in T_i(o_d)} C_i(t)$. We define $T_f(o_d)$ as the set of tasks that weren't executed successfully (i.e., $T_f(o_d) = T_A(o_d) - T(o_e)$), which includes all failure points and their successor tasks, and we define $T_f^i(o_d) \subseteq T_f(o_d)$ as the set of tasks that were assigned to agent i and weren't executed because of preceding failures.

Interdependent Valuations and Center rationality. Our problem differs from a classical mechanism design problem in two main aspects. *First, interdependent valuations.* Classical mechanism design normally assumes that the value $v_i(o_d, \theta_i)$ of an agent i of o_d depends *only* on its type θ_i (i.e., independent valuations). But here valuations are interdependent, as the actual value $v_i(o_e, \theta_i)$ of agent i clearly depends on its type, and the *actual* types of other agents who may cause execution failures (i.e., $o_d \neq o_e$). *Second, center rationality.* Classical mechanism design usually assumes that the central authority that determines the problem's outcome is an unbiased party that has no self-interests, as it solves a *social choice* problem that involves *only* the agents. And thus, it is not preferred that this authority ends up with any left-over from the agents' payments (i.e., happens if a weakly budget balanced mechanism is used), and redistributing the left-over using redistribution mechanisms is required. Here, we assume a commercial ITA model, where the center has its value of the determined outcome, and if center rationality holds, any left-over contributes toward the center's utility (i.e., profit). Thus, we follow Porter et al. [2008] in denoting budget balance as center rationality to point out this issue.

Investment Example. An investment company wants to improve the suitability of a piece of land for construction in order to sell it for a higher price. Possible interdependent tasks for the land improvement are site clearing, removal of trees, general excavation, installation of sewer lines, etc. Assume that the company decided on seven tasks that have the interdependencies $t_1 \prec t_2 \prec \ldots \prec t_7$. The company gets a reward of 10 from each completed task (i.e., the land's price increases by 10 after each task), and wants to assign the tasks to two contactors i and j.

Table 1: Investment Example

	t_1	t_2	t_3	t_4	t_5	t_6	t_7
θ_i	3	6	6	7	4	∞	6
θ_j	∞	4	∞	5	8	∞	3
$\theta_j^{'}$	5	4	5	5	8	∞	3
$\theta_j^{''}$	15	4	5	5	8	7	3

Table 1 includes the contractors' true types θ_i and θ_j, and two misreported types $\theta_j^{'}$ and $\theta_j^{''}$ for contractor j. We use ∞ to denote the inability to perform a task. If $\theta_j^{'}$ was reported, then $o_d = \{(t_1, i), (t_2, j), (t_3, j), (t_4, j), (t_5, i)\}$, where t_3 is assigned to contractor j instead of contractor i if θ_j was reported. Based on o_d, $T_A(o_d) = \{t_1, \ldots, t_5\}$. o_d will fail at task t_3 (because agent j can't execute it), and thus, $T(o_e) = \{t_1, t_2\}$, $T_i(o_e) = \{t_1\}$, $T_f(o_d) = \{t_3, t_4, t_5\}$, and $T_f^i(o_d) = \{t_5\}$.

Non-Negative ITA Model. In this study, we consider a non-negative ITA (NN-ITA) model, where each assigned task must incur a non-negative social welfare (Assumption II). We define our assumptions as follows.

Assumption I. Failure-Detection: *If any task t failed, this failure is detected and the responsible agent is identified.*

This assumption provides a task-by-task monitoring, and is very reasonable when the outcome of a problem is executable. This assumption was used by all similar studies (discussed in Section 5) that deal with outcome failures.

Assumption II. Non-Negative $SW(t)$: *The center will assign a task $t \in T$ only if it incurs a non-negative social welfare, i.e., $SW(t) \geq 0$.*

This assumption narrows down the situations where an efficient outcome is determined, but it is crucial for maintaining the center rationality property. In the general case, the center should assign a task t for a negative social welfare if this will allow assigning its successor tasks, and these successor tasks have a positive social welfare that compensates the negative social welfare of t. This is exactly the same as assuming that the center has combinatorial rewards for the tasks (e.g., gets a single reward of 10 from both t_1 and t_2), and it is proved that achieving center rationality is impossible for combinatorial rewards, even if there are no interdependencies between tasks [Porter *et al.*, 2008, Theorem 4.2].

3 Execution Failures and Sufficient Resources

In this section, we deal only with execution failures assuming that agents have sufficient resources. In other words, given the set of tasks T_i that agent i can perform, the agent has sufficient resources to execute all the tasks assigned for it from T_i. We present this impossibility result.

Theorem 1. *There is no efficient mechanism that achieves strategy-proofness for NN-ITA by using a single allocation round, even if agents have sufficient resources.*

Proof outline. If agent i reported its true type θ_i, then the outcome o_d may: *A.* have a successful execution, or *B.* fail by another agent $j \neq i$. If agent i reported $\theta_i^{'} \neq \theta_i$, then the outcome $o_d^{'}$ may: *1.* have a successful execution, *2.* fail by agent i, or *3.* fail by another agent $j \neq i$. To prove strategy-proofness (Definition 2), we need to show that $u_i(o_d, \theta_i) \geq u_i(o_d^{'}, \theta_i)$ holds in the six possible combinations of A and B respectively with 1, 2 and 3: *Case A1.* Both o_d and $o_d^{'}$ are successful, *Case A2.* o_d is successful and $o_d^{'}$ fails by agent i, *Case A3.* o_d is successful and $o_d^{'}$ fails by another agent $j \neq i$, *Case B1.* o_d fails by another agent $j \neq i$ and $o_d^{'}$ is successful, *Case B2.* o_d fails by another agent $j \neq i$ and $o_d^{'}$ fails by agent i, and *Case B3.* Both o_d and $o_d^{'}$ fail by another agent $j \neq i$. To prove Theorem 1, we prove that there is no payment schema that can cover cases A3 and B1 simultaneously. Let $o_e^{'}$ and $o_f^{'}$ be the executed and unexecuted parts of $o_d^{'}$, respectively.

Proof. Given that the actual value $v_i(o_e, \theta_i) = -\sum_{t \in T_i(o_e)} C_i(t)$ of agent i, the agent's payment $p_i(o_e)$ must increase with each task executed by agent i to compensate the decrease in the agent's value. Any payment schema *either* pays agent i based on only the executed tasks (i.e., $p_i(o_e)$), *or* will also include payments for the unexecuted tasks $T(o_f)$ (i.e., $p_i(o_e, o_f)$). For $p_i(o_e)$, $u_i(o_d, \theta_i) \geq u_i(o_d^{'}, \theta_i)$ will not hold for case B1. This is because agent i may incur some extra costs and prevent the failure [1], which increases the number of executed tasks (i.e., $o_e \subset o_e^{'}$), and thus, its payment. Agent i has incentive to do so if its utility with payment $p_i(o_e^{'})$ will be greater than its utility with payment $p_i(o_e)$. For $p_i(o_e, o_f)$, we want to stress that agent i can by strategic misreporting: 1. make tasks from o_f under o_d belong to $o_e^{'}$ under $o_d^{'}$ (e.g., as in case B1). The agent will do this if the increase in its utility from executed tasks is more than from unexecuted tasks; or 2. make tasks from o_e under o_d belong to $o_f^{'}$ under $o_d^{'}$ (e.g., as [2] in case A3). The agent will do this if the increase in its utility from unexecuted tasks is more than from executed tasks. For $u_i(o_d, \theta_i) \geq u_i(o_d^{'}, \theta_i)$

[1] In the investment example, θ_i and $\theta_j^{'}$ were reported, and o_d will fail by contractor j at t_3. Contractor i can claim t_3 under $o_d^{'}$ by reporting $\theta_i^{'}$ that misreports the cost of t_3 to be 4. Here, $o_d^{'}$ will not fail at t_3, because contractor i can perform t_3, however, for a higher cost than reported.

[2] In the investment example, θ_i and $\theta_j^{'}$ were reported. o_d will fail by contractor j at t_3. Contractor i can report $\theta_i^{'}$ that misreports the cost of t_1 to be 6, and makes t_1 assigned to contractor j under $o_d^{'}$, and $o_d^{'}$ will fail at t_1 because contractor j can't perform it.

to hold for both cases A3 and B1, the increase in the utility of agent i from executed tasks or unexecuted tasks must be the same. To achieve this, $p_i(o_e, o_f)$ must depend directly on the agent's privately known costs for the unexecuted tasks, which can be misreported. □

Reassignment Mechanism. One way to overcome this impossibility result is to design mechanisms that *reassign* failing tasks, i.e., if task t failed, then the center will reassign *only* task t to the agent who reported the second cheapest cost, and then, the execution can start again. The reassignment may happen several time for the same task (e.g., task t failed due to agent i, then reassigned to agent j and failed, then reassigned to agent k and succeeded), and may happen to more than one task. The reassignment will end if all the tasks in o_d were executed successfully, or if there is a *permanent failure* (i.e., a task that failed and can't be reassigned). We define a *temporary failure* as a task that failed and then was executed successfully after reassignment. Using reassignment is very reasonable and common in real-life applications, where the center needs the tasks to be executed. We will now propose a reassignment NN-ITA mechanism, and prove its properties. We denote α_{-i} as the set of agents without agent i, and we denote o_{re} as the executed outcome after the reassignment process. Given the executed outcome o_{re}, we define $SW_{-i}(o_{re})$ as the social welfare of o_{re} without the social welfare of the executed tasks by agent i, i.e., $SW_{-i}(o_{re}) = \sum_{j \in \alpha_{-i}} \sum_{t \in T_j(o_{re})} SW(t)$. As well, we define $SW(o^{-i}(o_{re}))$ as the social welfare of a *virtual* outcome $o^{-i}(o_{re})$, where $o^{-i}(o_{re})$ is the assignment that maximizes the social welfare given the types of other agents $j \neq i$ from the successfully executed tasks in o_{re} (i.e., $T(o_{re})$), while considering Assumption II, neglecting the reported information by agent $j \neq i$ regarding a certain task t if the agent caused its failure, and neglecting the dependencies between the tasks in $T(o_{re})$. *For instance, if θ_i and $\theta_j^{''}$ are reported in the investment example, $o_d = \{(t_1, i), (t_2, j), (t_3, j), (t_4, j), (t_5, i), (t_6, j), (t_7, j)\}$. o_d will fail at t_3, which will be reassigned to contractor i. Then, o_d will fail again at t_6 which is a permanent failure because it can't be reassigned to contractor i. Thus, $o_{re} = \{(t_1, i), (t_2, j), (t_3, i), (t_4, j), (t_5, i)\}$, and $SW_{-i}(o_{re}) = SW(t_2) + SW(t_4) = 6 + 5 = 11$. $SW(o^{-i}(o_{re})) = SW(t_2) + SW(t_4) + SW(t_5) = 6 + 5 + 2 = 13$, because $T(o_{re}) = \{t_1, t_2, t_3, t_4, t_5\}$ and when assigning them to contractor j, t_1 is not assigned because of its negative social welfare, t_3 is not assigned because it failed due to contractor j, and t_2, t_4, t_5 are assigned because we neglected their dependency on t_1 and t_3.*

Definition 5. *A* reassignment NN-ITA mechanism *is defined as follows.*

1. *The center announces the set of the offered tasks T. Then, agents report their types $\theta = (\theta_1, \ldots, \theta_n)$ to the center that will determine an efficient outcome o_d (Definition 1 under Assumption II).*
2. *The outcome o_d then will be executed resulting in o_{re} after reassignments. Each agent i will be paid as follows.*
   a. If agent i caused any temporary or permanent failure, then agent i will get no payment, i.e., $p_i(o_{re}) = 0$.
   b. If the outcome was executed successfully (possibly after reassignment) or permanently failed because of another agent $j \neq i$, then agent i will be paid $p_i(o_{re}) = \sum_{t \in T_i(o_{re})} R(t) + SW_{-i}(o_{re}) - SW(o^{-i}(o_{re}))$.

Theorem 2. *The reassignment NN-ITA mechanism is individually rational for every truthful agent.*

Proof. If agent i caused temporary or permanent failure, then its utility will be

$$u_i(o_d, \theta_i) = -\sum_{t \in T_i(o_{re})} C_i(t), \quad (1)$$

which is negative or 0 if agent i didn't execute any tasks (i.e., $T_i(o_{re}) = \emptyset$). If the execution was successful (possibly after reassignment) or permanently failed due to another agent $j \neq i$, then the utility of agent i will be

$$u_i(o_d, \theta_i) = \sum_{t \in T_i(o_{re})} R(t) - \sum_{t \in T_i(o_{re})} C_i(t) \quad (2)$$
$$+SW_{-i}(o_{re}) - SW(o^{-i}(o_{re})).$$

For every truthful agent i, its utility is Eq. 2, which can be re-written as $u_i(o_d, \theta_i) = \sum_{t \in T_i(o_{re})} SW(t) + SW_{-i}(o_{re}) - SW(o^{-i}(o_{re})) = SW(o_{re}) - SW(o^{-i}(o_{re}))$. Given that $o_{-i}(o_{re})$ is determined by assigning the executed tasks $T(o_{re})$, then $SW(o_{re}) \geq SW(o^{-i}(o_{re}))$ holds. This is because agent i executes its tasks in o_{re} for the cheapest possible cost (i.e., highest social welfare), but these tasks are assigned in $o^{-i}(o_{re})$ to other agents for higher costs. □

Theorem 3. *The reassignment NN-ITA mechanism is strategy-proof and efficient.*

Proof outline. Considering reassignment, we re-write the six cases in the proof outline of Theorem 1 as follows: *Case A1.* Both o_d and $o_d^{'}$ are successful (possibly after reassignment), *Case A2.* o_d is successful (possibly after reassignment) and any task in $o_d^{'}$ fails temporary or permanently by agent i, *Case A3.* o_d is successful (possibly after reassignment) and $o_d^{'}$ fails permanently by another agent $j \neq i$, *Case B1.* o_d fails permanently by another agent $j \neq i$ and $o_d^{'}$ is successful (possibly after reassignment), *Case B2.* o_d fails permanently by another agent $j \neq i$ and any task in $o_d^{'}$ fails temporary or permanently by agent i, and *Case B3.* Both o_d and $o_d^{'}$ fail permanently by another agent $j \neq i$. To prove strategy-proofness based on Definition 2, we will prove that $u_i(o_d, \theta_i) \geq u_i(o_d^{'}, \theta_i)$ holds in these six cases for any θ_{-i}, given that agent i may practise each type of strategic misreporting (i.e., over-reporting, under-reporting and misreporting costs) separately. By showing that practicing each lying type separately decreases the agent's utility under $o_d^{'}$, then we will have shown any combined strategic misreporting that involves more than one lying type may further decrease the agent's utility under $o_d^{'}$. We *stress* that the payment applies for all the agents who reported their information, and we don't assume that each agent is necessarily assigned tasks under o_d. Once strategy-proofness is established, efficiency follows from step 1 in Definition 5.

Proof. **Cases A2 and B2.** Under the outcome $o_d^{'}$, the utility of agent i will be negative or 0 (expressed by Eq. 1). However, under the outcome o_d, the agent has a non-negative utility expressed by Eq. 2 (established in Theorem 2). And thus, $u_i(o_d, \theta_i, \theta_{-i}) \geq u_i(o_d^{'}, \theta_i, \theta_{-i})$ holds. **Cases A1, A3, B1 and B3.** In all the four cases, the utility of agent i under o_d or $o_d^{'}$ is expressed by Eq. 2, and we want to prove that $u_i(o_d, \theta_i) = \sum_{t \in T_i(o_{re})} R(t) - \sum_{t \in T_i(o_{re})} C_i(t) + SW_{-i}(o_{re}) - SW(o^{-i}(o_{re})) \geq u_i(o_d^{'}, \theta_i) = \sum_{t \in T_i(o_{re}^{'})} R(t) - \sum_{t \in T_i(o_{re}^{'})} C_i(t) + SW_{-i}(o_{re}^{'}) - SW(o^{-i}(o_{re}^{'}))$ holds. *Over-reporting:* Given that $o_d^{'}$ is successful (possibly after reassignment) in cases A1 and B1, any over-reported tasks in $\theta_i^{'}$ weren't assigned to agent i. Given that $o_d^{'}$ fails permanently by another agent $j \neq i$ in cases A3 and B3, any over-reported tasks in $\theta_i^{'}$ before the permanent failure point weren't assigned to agent i. Given the previous and that Eq. 2 has no terms related to unexecuted tasks, over-reporting has no effect on the agent's utility. *Under-reporting:* If agent i was the only one capable of performing the task t that it under-reported or report a cost that is higher than the task's reward, then t will not be assigned (no agent can perform it or because of Assumption II) and its successor tasks will not be assigned under $o_d^{'}$. This may decrease the payment that agent i pays the center (i.e., $SW(o^{-i}(o_{re}^{'})) < SW(o^{-i}(o_{re}))$) if the unassigned tasks under $o_d^{'}$ contain tasks that were assigned to other agents $j \in \alpha_{-i}$ under o_d. However, this decrease corresponds to an equal decrease in the agent's received payment from the center (i.e., $SW_{-i}(o_{re}) < SW_{-i}(o_{re}^{'})$). As well, $\sum_{t \in T_i(o_{re})} R(t) - \sum_{t \in T_i(o_{re})} C_i(t) > \sum_{t \in T_i(o_{re}^{'})} R(t) - \sum_{t \in T_i(o_{re}^{'})} C_i(t)$ may hold if the unassigned tasks under $o_d^{'}$ contain tasks that were assigned to agent i under o_d, as any executed task by agent i corresponds to non-negative increase in its utility under Assumption II. *Misreporting costs:* By using reassignment, we **stress** that agent i doesn't need to misreport costs to prevent failures (as in footnote 1), as any failing tasks will be reassigned to agent i or any other agent $j \neq i$ who can execute them successfully. And thus, we can assume that misreporting costs doesn't affect the execution horizon (i.e., $T(o_{re}) = T(o_{re}^{'})$), which implies $SW(o^{-i}(o_{re}^{'})) = SW(o^{-i}(o_{re}))$. Given that $u_i(o_d, \theta_i) = \sum_{t \in T_i(o_{re})} R(t) - \sum_{t \in T_i(o_{re})} C_i(t) + SW_{-i}(o_{re}) = SW(o_{re})$, and $u_i(o_d^{'}, \theta_i) = \sum_{t \in T_i(o_{re}^{'})} R(t) - \sum_{t \in T_i(o_{re}^{'})} C_i(t) + SW_{-i}(o_{re}^{'}) = SW(o_{re}^{'})$, $SW(o_{re}) \geq SW(o_{re}^{'})$ holds because the center initially determines an efficient outcome that maximizes the social welfare, and reassigning failing tasks happens in a manner that maximizes the social welfare (i.e., reassign to the agent who reported the second cheapest cost). □

Theorem 4. *The reassignment NN-ITA mechanism is center rational, and provides profit for the center.*

Proof. In the truth-telling equilibrium, the center pays $p_i(o_{re}) = \sum_{t \in T_i(o_{re})} R(t) + SW_{-i}(o_{re}) - SW(o^{-i}(o_{re}))$ for each agent i. The center's utility of the executed outcome is $U(o_{re}, \theta) = V(o_{re}) - \sum_{i \in \alpha} p_i(o_{re}) = \sum_{t \in T(o_{re})} R(t) - \sum_{i \in \alpha} p_i(o_{re})$, and we need to show that $U(o_{re}, \theta) \geq 0$ holds. The term $\sum_{i \in \alpha} \sum_{t \in T_i(o_{re})} R(t)$ offsets the first term $\sum_{t \in T_i(o_{re})} R(t)$ of each payment $p_i(o_{re})$. Thus, we can represent the center's utility by the remaining terms of each $p_i(o_e)$, i.e., $U(o_{re}, \theta) = \sum_{i \in \alpha} SW(o^{-i}(o_{re})) - SW_{-i}(o_{re})$, and we need to prove that $SW(o^{-i}(o_{re})) \geq SW_{-i}(o_{re})$ holds for each agent i. Recalling that if a task was assigned to agent i, then agent i has the cheapest cost for performing it, and thus, the best social welfare $SW(t)$. Let $SW^{'}(t)$ be the second best social welfare, i.e., assign t to the agent who has the second cheapest cost. $SW(o^{-i}(o_{re})) \geq SW_{-i}(o_{re})$ holds because $SW(o^{-i}(o_{re}))$ contains $SW_{-i}(o_{re})$, in addition to the second best social welfare $SW^{'}(t)$ from each task t that was executed by agent i in o_{re}. This guarantees center rationality, and guarantees that the center gets a lower-bound profit of $SW^{'}(t)$ for each successfully executed task t, given that a second cheapest cost exists. □

4 NN-ITA with Limited Resources

In this section, we assume agents with limited resources, which is adequate for scenarios where acquiring additional resources is not possible. For representing resources, we assume that each agent i has a set of NAND (i.e., negated conjunctions) constraints T_i^{rc} defined over T_i to express the agent's resource constraints (e.g., $t_1, t_2 \in T_i$ and $\neg(t_1 \wedge t_2)$ mean that agent i can't execute both t_1 and t_2 because of limited resources, so the agent may be assigned only t_1, only t_2, or none of them). This representation is suitable because we defined T as a *set* of tasks, which - by definition - doesn't allow the repetition of tasks (e.g., if task t_1 is required to be executed twice, then the second copy must appear under a different notation $t_1^{'}$). Under outcome o_d, the assigned tasks to agent i (i.e., $T_i(o)$) must satisfy the agent's resource constraints (i.e., all constraints in T_i^{rc} must be true). Given that the resource constraints are privately known for agent i, these constraints can be under-reported or over-reported. In limited resources ITA, it is possible to achieve truthfulness in ex-post incentive compatible, but we will not present this result because center rationality is lost and due to space limits as well.

Theorem 5. *There is no mechanism that can achieve center rationality for limited resources NN-ITA, even under ex-post incentive compatible.*

Proof. Assume the following example: *1.* $T = \{t_1, t_2, t_3, t_4\}$ with interdependencies between tasks $t_1 \prec t_2$ and $t_3 \prec t_4$, and each task has a reward of 10; *2.* Two agents i and j; *3.* Agent i is the only agent who can perform t_1 for $C_i(t_1) = 4$ and t_3 for $C_i(t_3) = 2$, but has a resource constraint $\neg(t_1 \wedge t_3)$; *4.* Agent j is the only agent who can perform t_2 for $C_j(t_2) = 1$ and t_4 for $C_j(t_4) = 7$, but has a resource constraint $\neg(t_2 \wedge t_4)$; and *5.* Agent j reports truthfully (i.e., ex-post incentive compatibility). The center can assign either $o_d^1 = \{(t_1, i), (t_2, j)\}$, or $o_d^2 = \{(t_3, i), (t_4, j)\}$. Given that this example assumes no second cheapest cost for tasks (i.e., only one agent who can perform each task), any mechanism that guarantees truthfulness in ex-post incentive compatibility must pay each agent the whole reward of the task it executed. If agent i reported truthfully, then the center will choose o_d^1 (i.e., $SW(o_d^1) = 15 > SW(o_d^2) = 11$), and the utility of agent i will be $10 - 4 = 6$. Here, agent i can under-report the ability to perform t_1 (i.e., excludes o_d^1). This makes the center chooses the only remaining outcome o_d^2, and the utility of agent i will be $10 - 2 = 8$. To prevent that from happening, the center must pay agent i an amount more than the reward of t_1, and given that the center pays agent j the whole reward for t_2, then center-rationality is lost. □

This impossibility result finalizes our study, as center rationality is a crucial property for mechanisms proposed for commercial use. Maintaining center rationality as well as achieving truthfulness in dominant strategy for limited resources NN-ITA is possible by imposing assumptions (e.g., cost verification as in [Porter *et al.*, 2008]).

5 Discussion and Related Work

Interdependent Valuations. We stress that outcome failure problems (e.g., task allocation, multiagent planning) are not the only type of problem that involves interdependent valuations (see [Mezzetti, 2004] for other examples), and if tasks are not interdependent (i.e., independent valuations), strategy-proof mechanisms already exist (e.g., [Nisan and Ronen, 2001]). When valuations are interdependent, a Groves mechanism [Groves, 1973] loses its strategy-proofness, because its payment depends on the agents' tentative values. All previous efficient mechanisms for interdependent valuations settings achieve truthfulness at ex-post incentive compatibility. Mezzetti [2004] introduced a two-stage Groves mechanism, which works for any interdependent valuations problem. This mechanism is identical to a Groves mechanism, except for a second reporting phase, where agents report their actual values of the determined outcome, and the Groves payment is made based on these actual values. This second reporting phase can be eliminated under Assumption I, as the center is monitoring the outcome and knows the agents' actual values. Domain specific mechanisms for outcome failure problems can handle failures easily, as agent i can be the *only agent* behind the outcome failure (i.e., other agents are reporting truthfully under ex-post incentive compatibility). In [Porter *et al.*, 2008; Ramchurn *et al.*, 2009], mechanisms were proposed for task allocation, where valuations were interdependent in the first because of the interdependencies between tasks, while in the second because of assuming a trust-based model. In [van der Krogt *et al.*, 2008; Zhang and de Weerdt, 2009], mechanisms were proposed for multiagent planning, where valuations were interdependent because of the interdependencies between the plans executed by different agents. The multiagent planning model is more complicated than an ITA model, as interdependencies between actions are not pre-defined, and agents report their own goals and the goals' associated rewards.

Failure Models. Previous studies assume that an outcome may fail either accidentally (e.g., [Porter *et al.*, 2008]) by assuming that an agent privately knows its probability of success (PoS) of $[0, 1]$ when performing a particular task, or intentionally (e.g., [Zhang and de Weerdt, 2009]) as we assume here (i.e., an agent reports PoS of '1' for a task instead of reporting its true PoS of '0'). Accidental failure models assume that a task may fail even if the agent reported truthfully its PoS, and an agent will attempt a task only once. To extend the work proposed here to consider accidental failures, we need to differentiate between if an agent failed because it can't execute the task at all (where the task must be reassigned to another agent as we did here), and between if the agent can execute the task but failed because there is a PoS (where here the agent must keep trying to execute the task until it succeeds). We can achieve this differentiation by extending Assumption I to allow the center to decide whether an agent attempted to execute a task in the first place or not, and we leave that for future work.

Private Durations. Another way - a study we have under review - to design strategy-proof mechanisms for ITA without using reassignment is to factorize the agent's privately known cost for performing a task into two components: a privately known duration in which the agent can perform that task, and a publicly known unit cost associated with each duration unit. Although here and previous studies [Porter *et al.*, 2008; Ramchurn *et al.*, 2009; Zhang and de Weerdt, 2009] use Assumption I, assuming private durations gives an additional advantage, because if an agent claims the ability to perform a task in a shorter period than its actual capability, then the agent can easily be detected. With private costs, an agent can execute a task for a higher or lower cost than its actual cost without being detected.

6 Conclusions and Future Work

In this study, we proposed a reassignment mechanism that is efficient and strategy-proof when valuations are interdependent. And then, we illustrated the effects of assuming agents with limited resources. Interdependent valuations introduce a lot of complexities to the classical mechanism design problem, which only can be handled by designing domain specific mechanisms. Extending the current model and methods to consider combinatorial values in ITA, and to multiagent planning appear fruitful avenues of pursuit.

References

[Groves, 1973] T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.

[Mas-Colell *et al.*, 1995] A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford Uni. Press, 1995.

[Mezzetti, 2004] C. Mezzetti. Mechanism design with interdependent valuations: Efficiency. *Econometrica*, 72(5), 2004.

[Nisan and Ronen, 2001] N. Nisan and A. Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35, 2001.

[Porter *et al.*, 2008] R. Porter, A. Ronen, Y. Shoham, and M. Tennenholtz. Fault tolerant mechanism design. *Artificial Intelligence*, 172(15):1783–1799, 2008.

[Ramchurn *et al.*, 2009] S. D. Ramchurn, C. Mezzetti, A. Giovannucci, J. A. Rodriguez-Aguilar, R. K. Dash, and N. R. Jennings. Trust-based mechanisms for robust and efficient task allocation in the presence of execution uncertainty. *Journal of Artificial Intelligence Research*, 35:119–159, 2009.

[van der Krogt *et al.*, 2008] R. van der Krogt, M. M. de Weerdt, and Y. Zhang. Of mechanism design and multiagent planning. In *The 18th ECAI*, 2008.

[Zhang and de Weerdt, 2009] Y. Zhang and M. M. de Weerdt. Creating incentives to prevent intentional execution failures. In *IEEE/WIC/ACM*, pages 431–434, 2009.