

Article

On the Hardness of Lying under Egalitarian Social Welfare

Jonathan Carrero ^{1,†} , Ismael Rodríguez ^{1,2,†}  and Fernando Rubio ^{1,2,*,†} 

¹ Departamento Sistemas Informáticos y Computación, Facultad Informática, Universidad Complutense, 28040 Madrid, Spain; joncarre@ucm.es (J.C.); isrodrig@ucm.es (I.R.)

² Instituto de Tecnologías del Conocimiento, Universidad Complutense, 28040 Madrid, Spain

* Correspondence: fernando@sip.ucm.es; Tel.: +34-913-947-629

† These authors contributed equally to this work.

Abstract: When it comes to distributing resources among different agents, there are different objectives that can be maximized. In the case of egalitarian social welfare, the goal is to maximize the utility of the least satisfied agent. Unfortunately, this goal can lead to strategic behaviors on the part of the agents: if they lie about their utility functions, then the dealer might grant them more goods than they would be entitled to. In this work, we study the computational complexity of obtaining the optimal lie in this context. We show that although it is extremely easy to obtain the optimal lie when we do not impose any restrictions on the lies used, the problem becomes Σ_2^P -complete by imposing simple limits on the usable lies. Thus, we prove that we can easily make it hard to lie in the context of egalitarian social welfare.

Keywords: social welfare; complexity; multi-agent systems



Citation: Carrero, J.; Rodríguez, I.; Rubio, F. On the Hardness of Lying under Egalitarian Social Welfare. *Mathematics* **2021**, *9*, 1599. <https://doi.org/10.3390/math9141599>

Academic Editor: Ezequiel López-Rubio

Received: 4 June 2021

Accepted: 2 July 2021

Published: 7 July 2021

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2021 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Economics, Mathematics, and Computer Science are three disciplines that have been shown to provide interesting synergies between each other. For example, basic concepts of Economics can be successfully applied to the distribution of scarce resources in computer systems (see, e.g., [1–3]) or even to design programming paradigms [4], and mathematical techniques are useful in identifying the difficulty and limits to the development of certain economic tasks (see, e.g., [5–9]), and computer systems facilitate the implementation of new economic environments (see, e.g., [10–13]).

An environment of great current interest in which these three disciplines converge is the Multi-Agent Resource Allocation (MARA). Given a set of agents, a set of resources, and some utility functions denoting how much interest each agent has for each resource (or set of resources), the objective is to find the best way to distribute resources among agents. However, how do we define *the best way*? Depending on how we do it, we will get different problems. For example, we may want to maximize the profit of the auctioneer (called the *Winner Determination Problem* [14]), so that each resource is assigned to whoever is most willing to pay. Alternatively, we may want to maximize the utility obtained by the agents. In this case, we may be interested in maximizing the sum of profits obtained by each agent (utilitarian social welfare), or we may want to maximize the utility of the agent that is least benefited (egalitarian social welfare). In our case, we will focus on this last objective, that is, we want the agent benefiting the least to get the best possible profit. Note that this objective is especially important in situations such as delivery of humanitarian aid, whether during specific natural disasters, or in situations of support to vulnerable people.

A fundamental problem that we must face while distributing resources is the risk that agents cheat by lying. For example, if our objective were to maximize the overall utility of the summation of all agents (utilitarian social welfare), then agents could lie by communicating preferences for resources being much higher than the real ones. This way, the optimal distribution would consist in giving all the resources to those who have

exaggerated their lies the most. The Generalized Vickrey Auction (GVA) [15] discourages these lies by forcing each agent to pay what it makes the rest of the agents lose.

Unfortunately, the strategy used in GVA does not work in the case of egalitarian social welfare. Recall that in this case, we want to benefit the less favored agent. Thus, if an agent wants to lie, it should not exaggerate its preferences to the upside, but quite the contrary. That is, the optimal lie consists in saying that any resource would provide it with a very low profit. This way, giving it an acceptable profit requires assigning it a lot of resources. What can we do, then, to discourage these lies? In our previous work [16], we searched for possible mechanisms to make it harder to benefit from lying, and we empirically studied a very simple one consisting in forcing each agent to distribute a fixed total profit among all resources. For instance, suppose we force all agents to receive a 100 total profit if they receive all available resources. This constraint disables the trivial strategy consisting in reducing the declared utility for each resource, because if we reduce the declared utility for some resources, then we will have to increase the declared to others, since the total sum must continue to be 100. In that work, we carried out a series of experiments empirically showing that this simple strategy makes it very difficult for the agents to benefit from lying. In particular, even very effective lies based on some knowledge about the preferences of the other agents proved to be less profitable than saying the truth when that knowledge was not extremely accurate. However, the theoretical difficulty of lying under the egalitarian social welfare was not studied in that work.

Beyond the aforementioned high sensitivity to imprecise information about the preferences of others, in this paper we will show that finding good lies is also very difficult in computational terms, even when assuming totally accurate information. Let us say that the valuations of an agent are *partially locked* if the agent cannot freely lie about the valuations it gives to some specific resources: for them, any fake valuations it chooses must be within some given intervals which include the actual valuations. We will show that the problem of finding out if there exists a lie reaching some given profit under partially locked valuations is Σ_2^P -complete, meaning that the problem is one level over NP-completeness in the Polynomial Hierarchy of complexity (PH). Note that partial locked valuations are expected in real situations. Sometimes, the real interest (or lack of interest) of some agent for some specific resources could be easily estimated just by observing the current (publicly known) needs and situation of the agent. Moreover, for resources whose necessity is mainly constant in time (i.e., not dependent on the current situation), sometimes their valuation for some agent could be easily estimated from previous known interactions between this agent and these resources. In either case, a fake valuation of the agent falling out of some reasonable expected interval would be immediately detected by the community as a bluff. Thus, the lying capability of agents is expected to be partial in real scenarios. It is worth mentioning that reaching the Σ_2^P -completeness does not require forcing the addition of valuations of each agent to be some given constant, although the completeness remains if this condition is added.

This high computational complexity shows that using egalitarian social welfare is feasible even in situations where agents may want to lie, since finding good lies is very difficult not only in terms of sensitivity to knowledge imprecision (as shown in [16]), but also in terms of computational effort.

The rest of the paper is structured as follows. The next section defines the problems and identifies a particularly simple case where good lies are trivial to find. Next, in Section 3 we prove the Σ_2^P -completeness of lying in the more general case. A discussion is presented in Section 4, while conclusions and lines for future work are given in Section 5.

2. Formal Model

Next we present the formal notions and problems that will be considered in the rest of the paper. The main model is basically the same as the one introduced in [16]. However, in [16] we dealt with an experimental environment, while in the current paper we deal with

the theoretical approach, providing a completely new and original proof of the complexity of the problem.

Definition 1. Let $R = \{r_1, \dots, r_m\}$ be a set of resources and $A = \{A_1, \dots, A_n\}$ be a set of agents. The set of possible allocations of R to A is the set $\mathcal{A} = \underbrace{A \times \dots \times A}_m$. Given $\alpha \in \mathcal{A}$ with $\alpha = (\alpha_1, \dots, \alpha_m)$, we say that the allocation α assigns each resource r_j to agent α_j for all $1 \leq j \leq m$.

We will use utility functions to denote the profits given by resources to agents. Given a distribution of resources, a utility function returns a real number representing the utility that an agent assigns to the corresponding distribution of resources. That is, if we use u to represent the utility function of an agent and we have $u(\alpha_1) > u(\alpha_2)$, then the meaning is that the corresponding agent is more interested in distribution α_1 than in distribution α_2 .

Definition 2. A utility function is a function $u : \mathcal{A} \rightarrow \mathbb{Q}^{\geq 0}$.

We say that u depends only on agent A_i if, for all $\alpha = (\alpha_1, \dots, \alpha_m) \in \mathcal{A}$ and $\beta = (\beta_1, \dots, \beta_m) \in \mathcal{A}$ fulfilling $\alpha_j = \beta_j = A_i$ or $\alpha_j \neq A_i \neq \beta_j$ for each $1 \leq j \leq m$, we have $u(\alpha) = u(\beta)$.

If u depends only on agent A_i , then we also say that u is additive for agent A_i if for all $\alpha = (\alpha_1, \dots, \alpha_m) \in \mathcal{A}$ we have $u(\alpha) = \sum_{\{j \mid \alpha_j = A_i\}} u(\underbrace{A_{k_1}, \dots, A_{k_{j-1}}}_{j-1}, A_i, \underbrace{A_{k_{j+1}}, \dots, A_{k_m}}_{m-j})$, where A_k is any arbitrary agent with $A_k \neq A_i$.

Note that if u is additive for A_i , then it is possible to represent u using a vector $P = (p_1, \dots, p_n)$ with $p_i \in \mathbb{Q}^{\geq 0}$, where each element of the vector represents the utility obtained by each individual item. From a formal point of view, we can construct u without ambiguity from P : for all $\alpha = (\alpha_1, \dots, \alpha_m)$, $u(\alpha) = \sum_{\{j \mid \alpha_j = A_i\}} p_j$. Hence, from now on, any additive utility function will be denoted just by its vector $P = (p_1, \dots, p_n)$. Given $r \in \mathbb{Q}^{>0}$, an additive utility function u is r -limited if $\sum_{1 \leq i \leq n} p_i = r$.

If we were interested in considering non-additive utility functions, we would need to deal with the representation of preferences for *bundles* of resources. This can be done in several ways (for instance, extensionally, that is, providing specific outputs for each combination of resources). However, the specific representation is relevant to analyze the complexity of the problem [17]. Hereafter, we will only consider additive utility functions.

We use \mathcal{U} to represent the set of all possible additive utility functions. The utility functions of all agents in a tuple of agents A will be denoted by a tuple $U = (u_1, \dots, u_n) \in \mathcal{U}^n$, where each u_i is the utility function of agent A_i .

Our optimization problem is defined so that we distribute resources in such a way that we try to maximize the utility of the agent that receives less utility. The entity responsible for finding and making this distribution will be called the *auctioneer*. If several allocations of resources provide the same utility to the agent receiving less utility, then the one also giving more utility to the second agent receiving less utility will be preferred, and so on. In order to formalize these preferences, a single number encapsulating the utilities of all agents for each possible allocation of resources together will be defined as follows: we add the utility of each agent for the allocation multiplied by a factor being higher for agents receiving less utility. If these factors are carefully chosen, then the goal of maximizing the utility of the agent with less utility (and, if some agents tie up, the utility of the agent with the second lowest utility, and so on) will be equivalent to maximizing that number. The term $\mathbf{eg}_{A,R,U}(\alpha)$, defined later in Definition 4 will denote this number, and it will be the result of combining the utility values of the agents as defined by function $\mathbf{order_num}_M(w_1, \dots, w_h)$, introduced next in Definition 3.

Actually, we will use these notions to establish a total order among allocations so that, given two possible different allocations, one will always be preferred over the other. If two different allocations provide the same utility to the agent receiving less utility, and also to

the second agent receiving less utility, and so on up to all agents, then the indexes of the agents which these resources are given to will be used to break the tie in some arbitrary way. Note that the suitability of some *fake preferences* to achieve beneficial allocations of resources under the egalitarian social welfare will obviously depend on the allocation found by the auctioneer when these lies are used. By unambiguously defining a (unique) optimum for these allocations of resources, we will be able to unambiguously evaluate the utility of each possible combination of fake preferences. Of course, the expected deviations from these optimal allocations of resources in practice (actually, just finding these allocations will be NP-hard, as we will show later) could affect the suitability of the fake preferences being used to achieve beneficial allocations, in line with the sensitivity to variations observed in [16]. In particular, some fake preferences being beneficial under the actual optimal allocation could not be so if the auctioneer finds and carries out some other sub-optimal allocation.

Definition 3. The numeric order of non-negative numbers w_1, \dots, w_h for base M is defined as:

$$\text{order_num}_M(w_1, \dots, w_h) = \sum_{k=1}^h M^{(h-k)} \cdot w_k$$

Suppose w_1, \dots, w_h are always expected to belong to some closed numeric intervals W_1, \dots, W_h , respectively. Note that, if $M = \max(m_1, \dots, m_h) + 1$, where $m_i = \max(W_i)$, then the numeric order of these parameters for base M will give priority to w_1 , next to w_2 if there is a tie with w_1 , and so on. For instance, let us suppose $w_1, w_2, w_3, w'_1, w'_2, w'_3$ are real numbers in the interval $[1, 10]$. If $11^2 \cdot w_1 + 11 \cdot w_2 + w_3 > 11^2 \cdot w'_1 + 11 \cdot w'_2 + w'_3$; then (w_1, w_2, w_3) is preferred over (w'_1, w'_2, w'_3) according to those priorities, that is, we give priority to the first parameter, next, to the second one if there is a tie with the first, and finally, to the third one if there is a tie with the other two.

Definition 4. Let $A = (A_1, \dots, A_n)$, $R = (r_1, \dots, r_m)$, and $U = (u_1, \dots, u_n)$. For all $\alpha = (A_{b_1}, \dots, A_{b_m}) \in \mathcal{A}$, we define

$$\mathbf{eg}_{A,R,U}(\alpha) = \text{order_num}_M(u_{e_1}(\alpha), \dots, u_{e_n}(\alpha), b_1, \dots, b_m),$$

where $(u_{e_1}(\alpha), \dots, u_{e_n}(\alpha))$ is the ordering of set $\{u_1(\alpha), \dots, u_n(\alpha)\}$ from lowest to highest (where ties $u_i(\alpha) = u_j(\alpha)$ with $i < j$ are solved in any arbitrary way, e.g., by considering $u_i(\alpha)$ lower) and we have $M = \max(\{u_1(\alpha') \mid \alpha' \in \mathcal{A}\} \cup \dots \cup \{u_n(\alpha') \mid \alpha' \in \mathcal{A}\} \cup \{n\}) + 1$.

The egalitarian social welfare optimization problem consists in, given A, R, U , finding α maximizing $\mathbf{eg}_{A,R,U}(\alpha)$. We will denote the (unique) solution of the problem for A, R, U by $\mathbf{egsol}(A, R, U)$.

Note that the definition of the previous problem is not affected by whether partially locked valuations (recall the introduction) are considered or not, because the problem only concerns finding some distribution depending on some given utility functions—regardless of whether these utility functions really reflect the true valuations of the corresponding agents or not.

Typically, only the utility of the agent receiving less utility is considered in the literature in the definition of the previous problem, so no preference is defined between allocations giving the same utility to the agent achieving less utility (note that this is equivalent to considering $\mathbf{eg}_{A,R,U}(\alpha) = \min\{u_j(\alpha) \mid 1 \leq j \leq n\}$). The NP-hardness of this problem is proved in [18]. This problem variant can be trivially polynomially reduced to the variant introduced in Definition 4, which shows the NP-hardness of the latter. Next, we show that this NP-hardness also applies to the problem under the additional assumption that utility functions must be r -limited (note that, in this case, the term M in the previous definition will be $\max\{r, n\} + 1$).

Proposition 1. *Let us consider the variant of the egalitarian social welfare optimization problem given in Definition 4, where an additional input $r \in \mathbb{Q}^{>0}$ is given and all considered utility functions must be r -limited. The resulting problem is also NP-hard.*

Proof. The problem in Definition 4 for the case $\mathbf{eg}_{A,R,U}(\alpha) = \min\{u_j(\alpha) \mid 1 \leq j \leq n\}$ has already been proved to be NP-hard. In particular, in [18] the authors prove it by reducing into it the well-known problem PARTITION: they construct an instance having exactly two agents with the same preferences. Thus, it also shows the NP-hardness of the egalitarian problem with the additional restriction of using instances with two agents whose utility functions are the same (let G denote this particular problem). Hence, we also infer that the NP-hard problem is a more general one where the valuations of each agent are r -limited. This is because the NP-hard problem G can be trivially polynomially reduced to that r -limited problem: we can just set $r = v$, where v is the addition of all preferences of any of the agents. Recall that problem G considers $\mathbf{eg}_{A,R,U}(\alpha) = \min\{u_j(\alpha) \mid 1 \leq j \leq n\}$. Problem G can also be polynomially reduced to a problem taking $\mathbf{eg}_{A,R,U}(\alpha)$ as in Definition 4 and assuming r -limited preferences as follows: we just define the value of $\mathbf{eg}_{A,R,U}(\alpha)$ to be reached, as the least one guaranteeing the agent with less utility gets at least $r/2$ utility (note that achieving this value in $\mathbf{eg}_{A,R,U}(\alpha)$ will imply the other agent also reaches $r/2$). \square

Now we present the problem where an agent has to find its optimal fake utility function, that is, the (probably false) utility function allowing him to obtain the maximum utility when resources are distributed using egalitarian social welfare. Thus, we have to find the utility function satisfying that, if agent A_i communicates it, and the rest of agents communicate the utility functions that agent A_i estimated for them, then the utility that agent A_i obtains (using its *true* utility function) after applying egalitarian social welfare rules is maximized. In Definition 5, $\mathbf{fake}_{A,R,U,i}(f)$ will represent the utility that agent A_i obtains when it communicates to the auctioneer that its utility function is f . As can be expected, in order to compute this term, we need to solve the optimization problem described in Definition 4. Thus, this *maximization* requires optimization of a term that also requires another optimization.

Definition 5. *Given A, R, U as before and $i \in \mathbb{N}$, for all $f \in \mathcal{U}$ let $\mathbf{fake}_{A,R,U,i}(f) = u_i(\alpha_f)$ with $\alpha_f = \mathbf{egsol}(A, R, (u_1, \dots, u_{i-1}, f, u_{i+1}, \dots, u_n))$. The optimal fake utility problem consists in, given A, R, U , and i , finding f maximizing $\mathbf{fake}_{A,R,U,i}(f)$.*

Even though it could seem that this problem requires the optimization of a term whose computation requires performing another optimization, it is easy to check that finding the optimal solution (in particular, without r -limited utility functions or partially locked valuations) is not difficult at all.

Proposition 2. *Let us suppose some allocation of resources provides a non-null utility to all agents. Let the utility functions of each agent A_j estimated by agent A_i (and the utility of agent A_i itself when $j = i$) be $P^j = (p_1^j, \dots, p_n^j)$. The optimal fake utility function for A_i is $P^{f_i} = (c \cdot p_1^i, \dots, c \cdot p_n^i)$, where $c \in \mathbb{Q}^{>0}$ is any positive value such that $\sum_{1 \leq s \leq n} c \cdot p_s^i < p_l^k$ for all $k \neq i$ and l with $p_l^k > 0$.*

Proof. We have to prove that the proposed utility function is actually the optimal one for A_i (let us remark that c always exists). These fake preferences of A_i satisfy that, even when all resources are assigned to A_i , agent A_i is still the agent whose utility turns out to be lower than the utility of any other agent that receives any resource this agent has a non-null preference for. Let U' denote the vector of utility functions defined by preference vector P^{f_i} for agent A_i and preference vectors P^j for the other agents. Note that maximizing $\mathbf{eg}_{A,R,U'}(\alpha)$ in turn maximizes the *fake* utility of agent A_i (i.e., preferences P^{f_i}) in its role of the least satisfied agent *provided that* some non-null utility is given to the other agents

(let us remark that, if this restriction is impossible to satisfy, then there is a contradiction with our initial assumption that some allocation of resources α provides non-null utility to all agents). Moreover, the aim of A_i is maximizing its *true* utility provided that *the same* constraint holds (as allocations not giving a non-null utility to all agents will never be picked by the auctioneer). Note that maximizing the utility achieved with preferences $P^i = (p_1^i, \dots, p_n^i)$ subject to that constraint is equivalent to maximizing the utility achieving with preferences $P'^i = (c \cdot p_1^i, \dots, c \cdot p_n^i)$ subject to the same constraint. Thus, the optimal strategy for A_i consists in sending preferences P'^i to the auctioneer. \square

It is worth noting that this optimal lie $P'^i = (c \cdot p_1^i, \dots, c \cdot p_n^i)$ does not depend at all on the utilities the liar agent estimated the other agents will send to the auctioneer, because this expression is constant with respect to those utilities. Thus, in this case, no estimation is required at all.

3. Proving the Hardness of Lying under Partially Locked Valuations

The result given before in Proposition 2 shows that taking advantage of lying in an egalitarian social welfare allocation is extremely easy if no additional conditions are required, so in this case, this allocation does not have practical usability. Fortunately, partially locked valuations naturally apply in real situations: agents will not have *total* freedom to lie about their preference over all available resources—without triggering obvious distrust. Next we will explicitly consider that the liar could be unable to lie about its preferences over specific resources. In the introduction, we considered that for each resource, an *interval* could denote the set of valuations (including the true one) the liar agent could send to the auctioneer without raising trivial suspicion. However, our capability to denote which valuations are acceptable is not required to be that *rich* in order to make the resulting problem Σ_2^P -hard. Since the hardness is propagated via generalization, hardness results apply to more problem variants (and thus are more interesting) when proved for the most particular problem variants. Thus, here we consider the least general version of the problem we can prove its Σ_2^P -hardness for: the particular case where the liar can provide *any* valuation for some resources, but only the true valuation for others. This is indeed a particularization of a problem version based on intervals, because both situations can be trivially expressed by using intervals $[0, \infty)$ and $[p_i, p_i]$, respectively, where p_i is the true valuation of resource r_i .

Formally, let us consider the problem given in Definition 5 under the additional constraint that the liar cannot lie about some specific resources. For this new problem, optimal solutions cannot be trivially constructed as before. Actually, in this case, solving the problem means making a hard optimization of an expression, and in order to calculate this expression we need to handle another hard optimization, and this *does* make the resulting problem much harder. We show that the resulting (decision) problem is Σ_2^P -complete.

Definition 6. Let A, R, U be as before and $i \in \mathbb{N}$. Let $T \in \mathcal{P}([1..m])$ be a set of indexes denoting the resources for which agent i cannot lie (i.e., it cannot send false preferences to the auctioneer). Given A, R, U, T, i , and a target utility $Q \in \mathbb{Q}^{>0}$ for agent A_i , the fake utility problem under partially locked preferences, denoted by **FPL**, consists in finding out if there exists $f = (f_1, \dots, f_m)$, with $f_j = u_i^j$ for all $j \in T$, such that $\mathbf{fake}_{A,R,U,i}(f) \geq Q$.

Theorem 1. The fake utility problem under partially locked preferences is Σ_2^P -complete.

Proof. First we prove that the problem is in Σ_2^P . Note that our problem can be equivalently stated as finding out whether there *exists* a vector of numbers $f = (f_1, \dots, f_m)$, an allocation of resources $\alpha \in \mathcal{A}$, and a value $l \in \mathbb{Q}^{>0}$ such that, for all possible allocation of resources $\alpha' \in \mathcal{A}$, we have: (a) $\mathbf{eg}_{A,R,U'}(\alpha') < l$; (b) $\mathbf{eg}_{A,R,U'}(\alpha) = l$; (c) $f_j = u_i^j$ for all $j \in T$; and (d) $u_i(\alpha) \geq Q$, where $U' = (u_1, \dots, u_{i-1}, f, u_{i+1}, \dots, u_n)$ (note that only condition (a) depends on the universally quantified variable α'). Hence, we can define the problem

as the search for something of polynomial size such that, for all things of polynomial size, some property checkable in polynomial size holds. Thus, the problem belongs to Σ_2^P .

In order to prove the Σ_2^P -hardness, we will construct a polynomial reduction from a Σ_2^P -hard problem, **QSAT**₂ (also known as **QBF**₂) into **FPL**. This problem consists in checking whether the expression holds, given an expression $\exists x_1, \dots, x_n \forall y_1, \dots, y_m \varphi$, where φ is a propositional logic formula denoted in Disjunctive Normal Form (DNF) depending only on propositional variables $x_1, \dots, x_n, y_1, \dots, y_m$.

Let \bar{x} and \bar{y} abbreviate (x_1, \dots, x_n) and (y_1, \dots, y_m) , respectively. Then we have $\exists \bar{x} \forall \bar{y} \varphi \equiv \exists \bar{x} \neg \forall \bar{y} \neg \varphi \equiv \exists \bar{x} \neg \exists \bar{y} \neg \varphi \equiv \exists \bar{x} \neg \exists \bar{y} \varphi'$, where $\varphi' \equiv \neg \varphi$ is given in Conjunctive Normal Form (CNF). Hereafter, we will only consider that latter expression $\exists \bar{x} \neg \exists \bar{y} \varphi'$, where we assume $\varphi' = c_1 \wedge \dots \wedge c_k$ and $c_i = l_i^1 \vee l_i^2 \vee l_i^3$ for all $1 \leq i \leq k$, and all l_i^j is $x_h, \neg x_h, y_h$, or $\neg y_h$ for some h .

Given this instance of **QSAT**₂, we will create an instance of **FPL** from it such that there exists some \bar{x} , making it impossible to satisfy φ' for all \bar{y} if, in the **FPL** instance, a specific agent reaches a specific target utility in the auctioneer's allocation (i.e., in the egalitarian social welfare allocation) after sending some fake utilities to the auctioneer. That agent, called U_0 , will be able to reach that utility if the auctioneer does *not* manage to find some allocation giving at least some target utility to the agents receiving less utility. By setting its fake preferences in some way, U_0 will force the auctioneer to set variables \bar{x} (actually, the *resources* representing them) in some way, and then the auctioneer will try to set variables \bar{y} so as to satisfy φ' . If it does, then U_0 will *not* reach its target utility with its real preferences, and if it does not, then U_0 will reach it.

The agents, resources, and preferences (i.e., valuations) of the agents for the resources in the constructed **FPL** instance are schematically depicted in Figure 1. Circles denote agents, rectangles denote resources, and arrows show the non-null preferences of agents for resources.

Formally, in addition to agent U_0 , we also consider the following agents:

- U'_0 , *all*, and *all'*.
- For all $1 \leq i \leq n$, we have agents x_i and $\neg x_i$.
- For all $1 \leq i \leq m$, we have agents y_i and $\neg y_i$.
- For all $1 \leq j \leq k$, we have agent c_j .
- For all $1 \leq i \leq n + m$, we have agent col_i .

Besides, the set of resources R consists of the following resources:

- U_0 , U'_0 , *all*, *all'*, and U_0 -*all*.
- For all agents but agents c_j , we have a resource with the same name as the agent.
- For all $1 \leq i \leq n$, we have resources (x_i) , $\langle x_i \rangle$, and $\langle \neg x_i \rangle$. Besides, for all $1 \leq j \leq k$, we have resources $[x_i^j]$ and $[\neg x_i^j]$.
- For all $1 \leq i \leq m$ we have a resource (y_i) , and for all $1 \leq j \leq k$, resources $[y_i^j]$ and $[\neg y_i^j]$.
- For all $1 \leq i \leq k$, we have resources (c_i) and $[c_i]$.

Let $M = n + k$. For the sake of notation simplicity, let us denote the preference of agent a for resource r by $a(r)$. The (non-null) preferences of agents for resources are the following:

- For all $1 \leq i \leq n$, $x_i(x_i) = \neg x_i(\neg x_i) = n - 1$, $x_i(\langle x_i \rangle) = \neg x_i(\langle \neg x_i \rangle) = 1$, and $x_i((x_i)) = \neg x_i((x_i)) = k + 1$. Besides, for all $1 \leq j \leq k$, $x_i([x_i^j]) = \neg x_i([\neg x_i^j]) = 1$.
- For all $1 \leq i \leq m$, $y_i(y_i) = \neg y_i(\neg y_i) = n$, and $y_i((y_i)) = \neg y_i((y_i)) = k$. Besides, for all $1 \leq j \leq k$, $y_i([y_i^j]) = \neg y_i([\neg y_i^j]) = 1$.
- For all $1 \leq i \leq k$ with $c_i = l_i^1 \vee l_i^2 \vee l_i^3$ (note that each l_i^h is $x_h, \neg x_h, y_h$, or $\neg y_h$ for some h) we have $c_i([l_i^l]) = 2M - 2$ for each $1 \leq l \leq 3$. Besides, $c_i((c_i)) = c_i([c_i]) = M - 1$.
- $U_0(U_0) = k + 1$, $U_0(U_0$ - $U'_0) = n$, $U_0(U_0$ - $all) = 0.5$ and, for all $1 \leq i \leq n$, $U_0(\langle x_i \rangle) = U_0(\langle \neg x_i \rangle) = 1$.
- $U'_0(U'_0) = k$, $U'_0(U_0$ - $U'_0) = n + 2$ and, for all $1 \leq i \leq n$, $U'_0(\langle x_i \rangle) = U'_0(\langle \neg x_i \rangle) = 1$.

- $all(all) = n$, $all(U_0_all) = 1$, $all(all_all') = k + 3$ and, for all $1 \leq i \leq k$, $all(c_i) = 1$.
- $all'(all') = M$ and $all'(all_all') = 3M$.
- For all $1 \leq i \leq n + m$, $col_i(col_i) = n$. Besides, for all $1 \leq l \leq k$ we have $col_l([c_l]) = 1$, and in addition, for all $1 \leq j \leq n$ we have $col_i([x_j^l]) = col_i([\neg x_j^l]) = 1$, and for all $1 \leq j \leq m$ we have $col_i([y_j^l]) = col_i([\neg y_j^l]) = 1$.

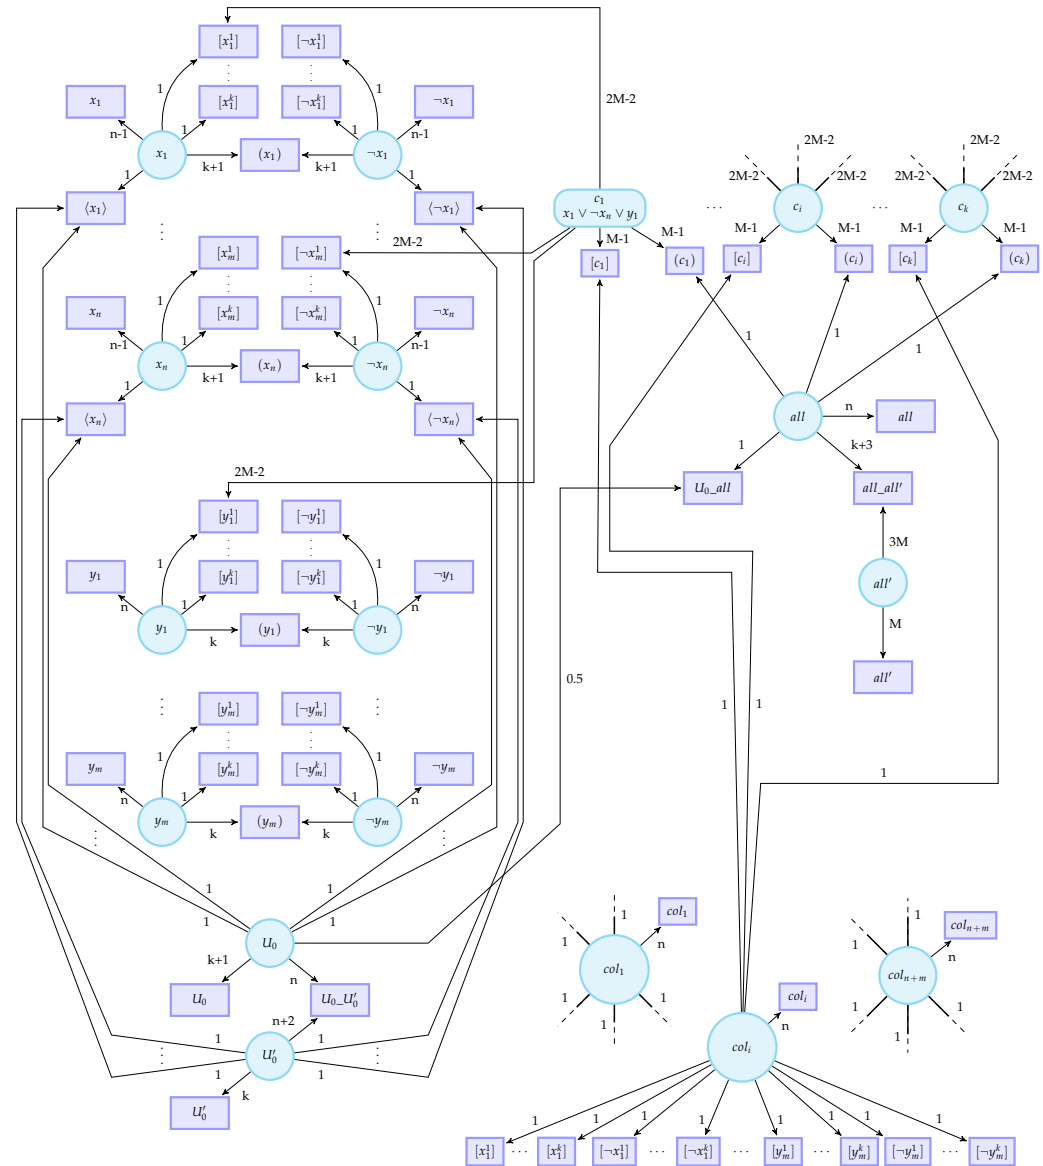


Figure 1. Scheme of the polynomial reduction.

All agents receive 0 utility for any other resource.

The set T of resource agent U_0 cannot show fake preferences for the following: $R \setminus \{\langle x_1 \rangle, \langle \neg x_1 \rangle, \dots, \langle x_n \rangle, \langle \neg x_n \rangle\}$. Thus, agent U_0 can underrate or overrate its preferences for resources in $\{\langle x_1 \rangle, \langle \neg x_1 \rangle, \dots, \langle x_n \rangle, \langle \neg x_n \rangle\}$ as much as it wants.

Finally, Q , the required *real* utility to be reached by U_0 from the allocation of resources formed by the auctioneer, is set to $M + 1.5$. This completes the instance of **FPL** constructed from the original **QSAT**₂ instance. Let us show that the reply to one instance is *yes* if it is *yes* to the other instance.

The **FPL** instance will simulate the **QSAT**₂ instance as follows. Each assignment of resources to agents in **FPL** will represent a *valuation* of propositional variables $x_1, \dots, x_n, y_1, \dots, y_n$: we will consider that the propositional variable x_i is set to \top when agent x_i gets re-

source $\langle x_i \rangle$, whereas a valuation where the propositional variable x_i is set to \perp will be represented when resource $\langle x_i \rangle$ is assigned to agent $\neg x_i$ (and the same for agents y_i and $\neg y_i$ and resource $\langle y_i \rangle$).

Let us show that, regardless of the fake preferences set by user U_0 for resources $\langle x_1 \rangle, \langle \neg x_1 \rangle, \dots, \langle x_n \rangle, \langle \neg x_n \rangle$, an allocation giving all agents at least M utility (*pretend* utility in the case of agent U_0 , that is, utility according to the fake preferences sent to the auctioneer) can always be reached, so only these allocations must be considered.

Note that each agent a is the only one interested in the resource with the same name a , so any optimal allocation of resources will give each resource a to agent a . Given this, each agent x_i (respectively, $\neg x_i$) has two possibilities to reach M utility. One of them consists in getting resource $\langle x_i \rangle$. The other one requires getting resource $\langle x_i \rangle$ (resp. $\langle \neg x_i \rangle$) and all k resources $[x_i^j]$ (resp. $[\neg x_i^j]$). The first case forces its “twin” agent $\neg x_i$ (resp. x_i) to get $\langle \neg x_i \rangle$ (resp. $\langle x_i \rangle$) and all k resources $[\neg x_i^j]$ (resp. $[x_i^j]$) to reach M itself, whereas the latter case forces its twin agent to get resource $\langle x_i \rangle$. A similar argument applies to each pair of agents y_i and $\neg y_i$, although no $\langle y_i \rangle$ or $\langle \neg y_i \rangle$ resources exist in this case. We conclude that, in any allocation of resources where all agents achieve at least M utility, if the allocation represents setting the propositional variable x_i to \top (resp. to \perp), then agents $x_i, \neg x_i$ will consume (receive) all resources of the forms $x_i, \neg x_i, \langle x_i \rangle, \langle \neg x_i \rangle, [x_i^j], [\neg x_i^j]$, but all k resources of the form $[x_i^j]$ (resp. all k resources of the form $[\neg x_i^j]$). A similar argument applies to the assignment of the propositional variable y_i to \top or \perp : the corresponding agents will consume all the corresponding resources (in this case, without resources $\langle y_i \rangle$ or $\langle \neg y_i \rangle$) but all k resources of the form $[y_i^j]$ or all k resources of the form $[\neg y_i^j]$, respectively.

These k resources left available for each variable x_i or y_i will be used by the auctioneer to try to satisfy each disjunctive clause c_j in ϕ' . In the **FPL** instance, this will translate into giving $2M - 2$ utility to the agent c_j representing clause c_j when it receives some resource representing one of the literals of the clause. Each agent c_j representing a clause $c_j = l_j^1 \vee l_j^2 \vee l_j^3$ will achieve $2M - 2$ utility by receiving any of the resources $[x_i^j], [\neg x_i^j], [y_i^j], [\neg y_i^j]$ representing the specific variable valuation required by any of its literals l_j^1, l_j^2 , or l_j^3 . For instance, if $c_5 = x_2 \vee \neg y_7 \vee \neg x_{10}$, then agent c_5 receives $2M - 2$ utility by receiving any of these resources: $[x_2^5], [\neg y_7^5], [\neg x_{10}^5]$. Thus, giving $2M - 2$ utility to some agent c_j this way means satisfying the logical clause c_j , and giving $2M - 2$ to all of them this way implies satisfying ϕ' . As we will see, achieving this will give the auctioneer access to an allocation increasing the utilities of the agents receiving less utility (thus, preferable), but will also give U_0 less real utility, preventing it from getting its Q target utility.

Each agent c_j will be able to reach at least M utility in a possible second way: by getting *both* resources $[c_j]$ and $\langle c_j \rangle$, agent c_j will also receive a $2M - 2$ utility ($M - 1$ utility for each of them). This alternative case will mean clause c_j was *not* satisfied (no resource l_j^1, l_j^2 , or l_j^3 was given to agent c_j). Since satisfying ϕ' will be preferable to the auctioneer, the auctioneer will try to prevent this case.

Note that, regardless of whether all clauses c_j are satisfied or not, making all agents $x_i, \neg x_i, y_i, \neg y_i$, and c_j achieve at least M utility does not require giving these agents more than $(n + m + 1) \cdot k$ resources of the kinds $[x_i], [\neg x_i], [y_i], [\neg y_i], [c_i]$: each pair of agents x_i and $\neg x_i$ or y_i and $\neg y_i$ will use k of them (there are $n + m$ pairs), and the remaining k resources will be required by the k agents c_j . Note that we are not counting in that $(n + m + 1) \cdot k$ expression the number of necessary resources of kinds $x_i, \neg x_i, \langle x_i \rangle, \langle \neg x_i \rangle, \langle x_i \rangle, y_i, \neg y_i, \langle y_i \rangle, \langle \neg y_i \rangle$. If all clauses c_j are satisfied, then no resource of the form $[c_j]$ will be necessary to make all of these agents reach at least M utility, and all $(n + m + 1) \cdot k$ necessary resources of the mentioned kinds will be of the forms $[x_i], [\neg x_i], [y_i], [\neg y_i]$. On the contrary, in the completely opposite case where no c_i is satisfied, $(n + m) \cdot k$ resources of the forms $[x_i], [\neg x_i], [y_i], [\neg y_i]$ will be necessary, as well as all k resources of the form $[c_i]$.

In fact, all the *remaining* $(n + m) \cdot k$ resources of these forms $[x_i], [\neg x_i], [y_i], [\neg y_i], [c_i]$ (that is, the resources of these kinds we do not need to assign to agents $x_i, \neg x_i, y_i, \neg y_i$,

and c_j to let all of them reach M utility) will in turn be needed by agents col_i . All of these $(n + m) \cdot k$ resources will be required to let these agents col_i reach at least M utility (in fact, exactly M): each agent needs k resources to reach M utility, and there are $n + m$ of them.

Agent all can get at least M utility in two possible ways. On the one hand, if all agents c_j get $2M - 2$ utility by receiving one resource $[x_i], [\neg x_i], [y_i], [\neg y_i]$ satisfying clause c_j , then no agent c_j will need any additional resource $[c_j]$ or (c_j) to achieve at least M utility. Resources of kind $[c_j]$ will be needed by agents col_i as mentioned before, but all resources (c_i) will be free to be assigned to agent all , and this agent will achieve exactly M utility just by receiving all of them. Let (a) be this case.

On the other hand, if some agent c_j does not get any resource $[x_i], [\neg x_i], [y_i], [\neg y_i]$, then it will need to take *both* resources $[c_j]$ and (c_j) to achieve at least M utility, particularly the $2M - 2$ utility, as in the other case (this time, this means clause c_j is *not* being satisfied). Thus, not all resources (c_j) will be available to agent all , and agent all will not be able to achieve M utility just by taking all resources (c_j) . Let (b) be this case.

By the arguments given in previous paragraphs, all allocations of resources formed by the auctioneer will be such that agents $x_i, \neg x_i, y_i, \neg y_i, col_i$ get exactly M utility and all agents c_j get exactly $2M - 2$ utility. As we will see, the auctioneer will always manage to give at least M utility to the remaining four agents, U_0, U'_0, all , and all' . Thus, the resource distribution chosen by the auctioneer will totally depend on the numeric utilities reached by these four agents.

The case (a) mentioned before will happen only if it is possible to satisfy all clauses, that is, if it is possible to satisfy φ' . As mentioned earlier, this case will enable the auctioneer to achieve a better distribution of resources for the agents according to the egalitarian social welfare. In particular, all agents U_0, U'_0, all , and all' will achieve at least $M + 1$ utility. On the contrary, in case (b), happening when it is impossible to satisfy all clauses, at least one of these four agents will reach less than $M + 1$ utility, making this case less desirable to the auctioneer. However, this case will be more profitable to U_0 according to its *real* preferences, as it will reach its target utility $Q = M + 1.5$ only in this case.

Let us suppose that the answer to the $QSAT_2$ instance is *yes*. Then, there exists some \bar{x} making φ' false for any \bar{y} . Let us see that the answer to the constructed instance of **FPL** is *yes*. Let the fake preferences of U_0 be the ones representing the valuation of \bar{x} making φ' false for any \bar{y} . That is, if x_i is \top (respectively, \perp) in this valuation, then U_0 pretends to have preferences $U_0(\langle x_i \rangle) = 1$ (resp. 0) and $U_0(\langle \neg x_i \rangle) = 0$ (resp. 1), instead of its real preferences (recall the real preferences are 1 in both cases). Let us study the allocations of resources the auctioneer could form in this case, and let us show that the actual allocation of resources chosen by the auctioneer gives $M + 1.5$ *real* utility to U_0 . Recall that, in all cases, resources U_0, U'_0, all , and all' are assigned to the agents with the same names, giving them $k + 1, k, n$, and M utility, respectively. We have the following cases:

- (1) The auctioneer gives to U_0 the resources of types $\langle x_i \rangle, \langle \neg x_i \rangle$ agent U_0 receives 1 utility from. Thus, U_0 receives n utility from them, adding up to $M + 1$ utility after also counting resource U_0 . In order to give enough utility to agent U'_0 , no resource of types $\langle x_i \rangle, \langle \neg x_i \rangle$ will be available for this agent, so resource $U_0 - U'_0$ must be assigned to agent U'_0 , making it reach $M + 2$ utility after resource U'_0 is also counted. Even after counting resource all , agent all cannot reach M utility just by taking all resources (c_j) not assigned to agents c_j , because $k - 1$ of them are available at most (recall that φ' cannot be satisfied in the valuation of \bar{x} preferred by U_0). If resource all_all' is given to agent all' , then agent all' reaches utility $4M$ after counting resource all' . In this case, agent all reaches at most $M - 1$ utility if resource U_0_all is given to agent U_0 , and at most M utility when it is given to agent U_0 . In the former case, the utilities of U_0, U'_0, all, all' are at most $M + 1.5, M + 2, M - 1, 4M$, respectively, whereas in the latter case, these utilities are at most $M + 1, M + 2, M$, and $4M$. Alternatively, if resource all_all' is given to agent all , then agents U_0, U'_0, all, all' achieve utilities $M + 1.5, M + 2, M + 3 + d, M$ for some $d \geq 0$ if resource U_0_all

is given to agent U_0 (recall that agent all could also receive some resources (c_j)), and utilities $M + 1, M + 2, M + 4 + d, M$ if that resource U_{0_all} is given to agent all . Out of these four possible allocations in this case, the third one (with utilities $M + 1.5, M + 2, M + 3 + d, M$) is the best one according to the egalitarian social welfare: the first allocation does not give at least M utility to all agents, and only the third one gives $M, M + 1.5$, and $M + 2$ utilities to the three agents with less utility. Hence, the auctioneer chooses it, and this way, U_0 reaches $Q = M + 1.5$ utility (both pretend and real).

- (2) The auctioneer does *not* allocate the resources to form the specific valuation of variables x_i demanded by the fake preferences of U_0 , although U_0 receives at least one of these resources. Thus, U_0 does not receive all the resources $\langle x_i \rangle$ and $\langle \neg x_i \rangle$ it gives 1 utility to, and at most, it receives $n - 1$ of them, achieving from them at most $n - 1$ utility. Let us see that, in any allocation of resources reached by agents U_0, U'_0, all, all' in this case, the multiset of their utilities will always be worse than the multiset $\{M + 1.5, M + 2, M + 3 + d, M\}$ reached in the allocation chosen in case (1), as we saw before. Thus, that allocation picked in case (1) will be preferred by the auctioneer over any allocation in this case.

Let us suppose that, by giving to U_0 up to $n - 1$ of the resources of types $\langle x_i \rangle$ and $\langle \neg x_i \rangle$ it wants and at least one of them, the auctioneer can fulfill φ' (recall that this is impossible if all n of them are given to U_0). Alternatively, if the auctioneer could not fulfill φ' , then the possibilities listed next would just be reduced. If the only resources of types $\langle x_i \rangle$ and $\langle \neg x_i \rangle$ agent U_0 receives are them, and the remaining ones (at most $n - 1$ of them) are given to U'_0 (which does not have any preference for them, as all of them give U'_0 a 1 utility), then resource $U_{0_U'_0}$ must be given to U'_0 as well, because it needs it to reach at least M utility. By counting these resources as well as resources U_0 and U'_0 , agents U_0 and U'_0 reach at most M utility and at least $M + 3$ utility, respectively.

If agent all gets all k resources (c_j) and agent all' receives resource all_all' , then, after counting resources all and all' , these agents achieve M and $4M$ utility, respectively. If the remaining resource (i.e., U_{0_all}) is given to U_0 , then the multiset of utilities of these four agents U_0, U'_0, all, all' is, at best, $\{M + 0.5, M + 3 + d, M, 4M\}$ for some $d \geq 0$, and if resource U_{0_all} is given to agent all , then this multiset is, at best, $\{M, M + 3 + d, M + 1, 4M\}$. These multisets of utilities for these four agents are less desirable than the multiset $\{M + 1.5, M + 2, M + 3, M\}$ seen in case (1).

Alternatively, if agent all receives resource all_all' , then the multiset of utilities of these four agents will be no more attractive than $\{M + 0.5, M + 3 + d, M + 3 + e, M\}$ or $\{M, M + 3 + d, M + 4 + e, M\}$ (for some $d \geq 0$ and $e \geq 0$) depending on whether U_{0_all} is given to agent U_0 or to agent all , respectively. Both multisets of utilities are again worse than $\{M + 1.5, M + 2, M + 3, M\}$.

- (3) Finally, we consider the case where U_0 receives none of the resources $\langle x_i \rangle$ and $\langle \neg x_i \rangle$. In this case, U_0 clearly needs resource $U_{0_U'_0}$ to reach at least M utility, and some combination of n resources of types $\langle x_i \rangle$ and $\langle \neg x_i \rangle$ will be given to U'_0 . Let us suppose it is a combination representing some \bar{x} letting φ' hold for some \bar{y} (if it is not, then just some cases listed next will be impossible). By receiving the mentioned resources as well as resources U_0 and U'_0 , agents U_0 and U'_0 get $M + 1$ and M utility, respectively. On the one hand, if resource all_all' is given to agent all then the multiset of utilities of agents U_0, U'_0, all, all' will be $\{M + 1.5, M, M + 3 + d, M\}$ or $\{M + 1, M, M + 4 + d, M\}$, depending on whether U_{0_all} is given to U_0 or to all . On the other hand, if resource all_all' is given to agent all' , then the corresponding multisets in these two cases will be no more attractive than $\{M + 1.5, M, M, 4M\}$ or $\{M + 1, M, M + 1, 4M\}$, respectively. All of these multisets are again less attractive than $\{M + 1.5, M + 2, M + 3, M\}$.

Let us suppose that the answer to the $QSAT_2$ instance is *no*. Then, there does not exist \bar{x} making φ' false for any \bar{y} . Let us see that the answer to the constructed instance of **FPL** is

no. That is, the allocation of resources formed by the auctioneer will not give $Q = M + 1.5$ real utility to U_0 in any case. We consider the following cases:

- (i) Under the fake preferences set by agent U_0 for resources of types $\langle x_i \rangle$ and $\langle \neg x_i \rangle$, the auctioneer can give some or all of these resources to U_0 in such a way that U_0 gets at least $n - 1$ fake utility from them. Then, assigning these resources to U_0 gives it $M + d$ fake utility (after also counting the effect of resource U_0), where $d \geq 0$ is the excess of utility over $n - 1$ given by these resources. Since less than n resources $\langle x_i \rangle$ and $\langle \neg x_i \rangle$ will remain to be given to U'_0 , making U'_0 reach at least M utility implies giving $U_0 _ U'_0$ to U'_0 . This way, U'_0 reaches $M + 2 + e$ utility, where e is the number of resources $\langle x_i \rangle$ and $\langle \neg x_i \rangle$ given to U'_0 . Since we are assuming that some \bar{y} can satisfy φ' no matter what \bar{x} is chosen, the auctioneer can manage to leave all k resources (c_j) unassigned to agents c_j (i.e., it can satisfy all clauses c_j), meaning that agent *all* can receive all of these k resources and reach M utility with them.

On the one hand, if resource *all_all'* is given to agent *all'*, then it obtains $4M$ utility. Thus, the utilities of U_0 , U'_0 , *all*, *all'* are given by $\{M + 0.5 + d, M + 2 + e, M, 4M\}$ or $\{M + d, M + 2 + e, M + 1, 4M\}$, depending on whether $U_0 _ all$ is assigned to U_0 or to *all*, respectively. On the other hand, if *all_all'* is given to agent *all*, then the previous two cases turn into allocations with utility multisets $\{M + 0.5 + d, M + 2 + e, M + 3 + k, M\}$ and $\{M + d, M + 2 + e, M + 4 + k, M\}$, respectively.

Note that, in addition to the previous four possible allocations, other four possible allocations arise by *not* giving any $\langle x_i \rangle$ or $\langle \neg x_i \rangle$ resource to agent U_0 . In this case, U'_0 can achieve M utility by receiving n resources of form $\langle x_i \rangle$ and $\langle \neg x_i \rangle$, which in turn implies giving $U_0 _ U'_0$ to agent U_0 to let it also reach at least M utility. The translation of the previous four cases into this local rearrangement between agents U_0 and U'_0 gives rise to the following utility multisets, respectively:

$\{M + 1.5, M, M, 4M\}$,
 $\{M + 1, M, M + 1, 4M\}$,
 $\{M + 1.5, M, M + 3 + k, M\}$, and $\{M + 1, M, M + 4 + k, M\}$.

If $d > 0$, then, out of these eight possible allocations, the second one, giving $\{M + d, M + 2 + e, M + 1, 4M\}$, is the only one reaching strictly more than M utility for all four agents, so it is chosen by the auctioneer. Since U_0 cannot receive more than n resources $\langle x_i \rangle$ or $\langle \neg x_i \rangle$ in any allocation, and neither $U_0 _ all$ nor $U_0 _ U'_0$ are given to U_0 in this allocation, U_0 does *not* reach $M + 1.5$ real utility in this case, as required.

If $d = 0$, then, out of these eight possibilities, again the second one, with utilities multiset $\{M, M + 2 + e, M + 1, 4M\}$, is the preferred one for the auctioneer (note that this is the only one where the three lowest utilities are M , $M + 1$, and at least $M + 2$), so it is again chosen by the auctioneer. For the same reasons as before, again U_0 does not reach $M + 1.5$ real utility.

- (ii) Under the fake preferences set by agent U_0 for resources of types $\langle x_i \rangle$ and $\langle \neg x_i \rangle$, the auctioneer can give some or all of these resources to U_0 in such a way that U_0 gets at least $n - 2$ and less than $n - 1$ fake utility from them. By similar reasoning as before, this time, the first four allocation cases considered in the previous case yield the following utility multisets:

$\{M - 0.5 + d, M + 2 + e, M, 4M\}$,
 $\{M - 1 + d, M + 2 + e, M + 1, 4M\}$,
 $\{M - 0.5 + d, M + 2 + e, M + 3 + k, M\}$, and $\{M - 1 + d, M + 2 + e, M + 4 + k, M\}$,

where $0 \leq d < 1$. In addition, the last four allocations considered in case (i) yield exactly the same utility multisets as in that case. Note that the second and fourth allocations do not provide at least M utility to all four agents. The remaining allocations in our set of eight possible allocations give at least M utility to the agent getting less utility, although only the sixth one provides $M + 1$ utility to the second agent with less utility (note that $M - 0.5 + d < M + 1$). Thus, the auctioneer picks it. Since U_0 does not obtain resource $U_0 _ all$ in this allocation, U_0 does not achieve $M + 1.5$ real utility in this case.

- (iii) Under the fake preferences set by agent U_0 for resources of types $\langle x_i \rangle$ and $\langle \neg x_i \rangle$, the auctioneer cannot give some or all of these resources to U_0 in such a way that U_0 gets at least $n - 2$ fake utility from them. Then, the first four allocations considered in the previous cases (where U_0 is given some combination of $\langle x_i \rangle$ and $\langle \neg x_i \rangle$ resources) cannot be better than in these previous cases. Thus, the allocation preferred by the auctioneer is again the sixth allocation, and again, U_0 does not get $M + 1.5$ real utility.

We conclude that the answer for some QSAT_2 instance is *yes* if the answer to the corresponding FPL instance is *yes*. It is easy to see that the FPL instance can be built from the original QSAT_2 instance in polynomial time with respect to the size of the QSAT_2 instance, so we have a polynomial reduction from QSAT_2 into FPL , and the Σ_2^P -hardness of FPL holds. \square

A simple consequence of the previous result is the Σ_2^P -completeness of another problem variant where, for each resource, an interval denotes the set of valuations the liar agent can give to it. As we saw at the beginning of this section, the Σ_2^P -hardness of this variant is inferred by simple generalization, as it generalizes the problem in Definition 6. On the other hand, in order to prove the inclusion of this variant in class Σ_2^P , it suffices to change condition (c) in the first paragraph of the previous proof by a condition requiring that, for all j , f_j is within the interval specified for the j -th resource.

Another possible variant consists in accepting only r -limited utility functions. The Σ_2^P -completeness is preserved in this case as well.

Theorem 2. *Let us consider that, in the problem given in Definition 6, utility functions must be r -limited (i.e., for some given $r \in \mathbb{Q}^{>0}$ and for all agents, the addition of preferences for all resources must be r). The resulting problem is Σ_2^P -complete.*

Proof. The proof of the inclusion in class Σ_2^P is basically the first paragraph in the proof of Theorem 1 (we just have to add the condition $\sum_{1 \leq j \leq m} f_j = r$). Besides, we can prove the Σ_2^P -hardness by using almost the same proof as in Theorem 1—just a few changes must be introduced to include the r -limitation.

In the FPL instance constructed there, we set r to any value higher than the addition of preferences shown by any agent for all resources. An additional agent z and an additional resource z are added so that $z(z) = r$ and, for all other agent a , we set $a(z) = r - u$, where u is the addition of preferences shown by agent a for all other resources. Agent z only has a non-null preference for resource z . Thus, this resource will be given to agent z in any allocation conforming the egalitarian social welfare, so agent z and resource z will be transparent to the rest of the setting.

In addition, the set T of resources agent U_0 cannot show fake preferences, for it is redefined as follows: $T = R \setminus \{z, \langle x_1 \rangle, \langle \neg x_1 \rangle, \dots, \langle x_n \rangle, \langle \neg x_n \rangle\}$. Since resource z will always be given to agent z , agent U_0 can underrate or overrate its preferences for resources in $\{\langle x_1 \rangle, \langle \neg x_1 \rangle, \dots, \langle x_n \rangle, \langle \neg x_n \rangle\}$ as much as it wants, in particular by overrating or underrating z accordingly, respectively, to keep the addition of preferences for all resources being equal to r . \square

4. Discussion

A previous work [16] showed the high sensitivity to imprecise information when lies are introduced in the egalitarian social welfare allocation: profitable lies are probably worse than saying the truth if the preferences of other agents are not exactly those assumed by the liar agent. Together with the Σ_2^P -completeness we have proved in this paper, both results show that taking profit from lying in the egalitarian social welfare is not as easy as one might think at a first glance: not only very precise information about the neighbours is required to compose useful lies, but finding these useful lies from that information is also very hard from the computational point of view (actually, one level over NP-completeness in PH). Note that the computational difficulty to find exact solutions introduces a new

level of sensitivity to imprecision, although, this time, not imprecise information about the other agents, but imprecise results if we try to sub-optimally solve the problem to cope with its intractability. The utility of a lie dramatically depends on whether the allocation delivered by the auctioneer is *exactly* the one previously predicted by the liar agent: a different allocation, giving slightly more utility to the least satisfied agent, would be preferable to the auctioneer, and both allocations could give a very different profit to the liar agent. Facing a Σ_2^P -complete problem means making a hard optimization over a function whose computation requires another hard optimization, so any imprecision in the latter optimization will dramatically mislead the former one. This brings an additional layer of difficulty to find useful lies in practice.

All in all, the egalitarian social welfare allocation has some features making it more resilient to lies than one might initially expect. Further research is required to assess this resilience in real case studies.

Regarding the limits of our current research, let us remark that we are working with partially locked valuations. That is, we use the (reasonable) assumption that agents cannot freely lie about *all* their preferences about all the resources. Although this assumption is quite realistic, it would also be interesting to explore the complexity of the problem if this assumption does not hold, and we can only assume that preferences are r -limited.

5. Conclusions and Future Work

In this paper, we have proved the Σ_2^P -completeness of the problem of finding some fake preferences such that, if they are sent to the auctioneer, then the egalitarian social welfare allocation will make the liar agent reach some given target (real) utility. This result assumes that the valuations of the resources the liar agent can send to the auctioneer are partially locked. This Σ_2^P -completeness also extends to the case where, in addition, the sum of valuations for all resources must be equal to some given constant.

Our future work will consist in studying this allocation in realistic scenarios, particularly implemented as smart contracts. This will allow us to improve the practical relevance of our study, dealing with practical examples in a context where the whole process can be verified by any partner.

We also want to analyze if the problem remains Σ_2^P -complete even if partially locked valuations are fully removed from the problem definition (i.e., the liar agent can lie about any resource without any restrictions).

Let us remark that in our current work, we are only analyzing the complexity of the decision problem. In contrast, in our previous work [16] we dealt with experiments trying to find suboptimal (but useful) lies in different contexts. In this sense, an interesting line of future work would be to analyze, from a formal point of view, the difficulty of finding good approximations under different scenarios. That is, we would like to find the approximation class the problem belongs to (see, e.g., [19,20]). Moreover, we are interested in exploring alternative heuristic methods (like River Formation Dynamics [21]) that could be used to find suboptimal solutions.

Author Contributions: Conceptualization, I.R. and F.R.; methodology, I.R. and F.R.; validation, J.C., I.R. and F.R.; formal analysis, J.C., I.R. and F.R.; investigation, J.C., I.R. and F.R.; resources, J.C.; writing—original draft preparation, J.C., I.R. and F.R.; writing—review and editing, J.C., I.R. and F.R.; visualization, J.C.; supervision, I.R. and F.R.; project administration, I.R. and F.R.; funding acquisition, F.R. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by projects PID2019-108528RB-C22, and by Comunidad de Madrid as part of the program S2018/TCS-4339 (BLOQUES-CM) co-funded by EIE Funds of the European Union.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Not applicable.

Conflicts of Interest: The authors declare no conflict of interest.

Abbreviations

The following abbreviations are used in this manuscript:

CNF	Conjunctive normal form
DNF	Disjunctive normal form
FPL	Fake utility problem under partial truthful limited preferences
MARA	Multiagent resource allocation
PH	Polynomial hierarchy
QSAT	Quantified satisfaction problem

References

1. Stonebraker, M.; Aoki, P.M.; Litwin, W.; Pfeffer, A.; Sah, A.; Sidell, J.; Staelin, C.; Yu, A. Mariposa: A wide-area distributed database system. *Int. J. Very Large Data Bases* **1996**, *5*, 48–63. [\[CrossRef\]](#)
2. Buyya, R.; Abramson, D.; Venugopal, S. The grid economy. *Proc. IEEE* **2005**, *93*, 698–714. [\[CrossRef\]](#)
3. León, X.; Trinh, T.A.; Navarro, L. Using economic regulation to prevent resource congestion in large-scale shared infrastructures. *Future Gener. Comput. Syst.* **2010**, *26*, 599–607. [\[CrossRef\]](#)
4. Miller, M.S.; Drexler, K.E. Markets and computation: Agoric open systems. *Ecol. Comput.* **1988**, *1*, 133–176.
5. Maymin, P.Z. Markets are efficient if and only if $P = NP$. *Algorithmic Financ.* **2011**, *1*, 1–11. [\[CrossRef\]](#)
6. Rodríguez, I.; Rubio, F.; Rabanal, P. Automatic media planning: Optimal advertisement placement problems. In Proceedings of the 2016 IEEE Congress on Evolutionary Computation (CEC), Vancouver, BC, Canada, 24–29 July 2016; pp. 5170–5177.
7. Rodríguez, I.; Rabanal, P.; Rubio, F. How to make a best-seller: Optimal product design problems. *Appl. Soft Comput.* **2017**, *55*, 178–196. [\[CrossRef\]](#)
8. Chen, S.H.; Kaboudan, M.; Du, Y.R. *The Oxford Handbook of Computational Economics and Finance*; Oxford University Press: Oxford, UK, 2018.
9. Monaco, G.; Moscardelli, L.; Velaj, Y. On the performance of stable outcomes in modified fractional hedonic games with egalitarian social welfare. In Proceedings of the AAMAS’19, Montreal, QC, Canada, 13–17 May 2019; pp. 873–881.
10. Schafer, J.B.; Konstan, J.A.; Riedl, J. E-commerce recommendation applications. *Data Min. Knowl. Discov.* **2001**, *5*, 115–153. [\[CrossRef\]](#)
11. Núñez, M.; Rodríguez, I.; Rubio, F. A tutoring system supporting experimentation with virtual macroeconomic environments. In *International Conference on Artificial Intelligence: Methodology, Systems, and Applications*; Springer: Berlin/Heidelberg, Germany, 2004; pp. 361–370.
12. Nakamoto, S. Bitcoin: A Peer-to-Peer Electronic Cash System. 2008. Available online: <https://bitcoin.org/bitcoin.pdf> (accessed on 2 July 2021).
13. De Filippi, P. What blockchain means for the sharing economy. *Harv. Bus. Rev. Digit. Artic.* **2017**, *15*, 2–5.
14. Lehmann, D.; Müller, R.; Sandholm, T. The winner determination problem. In *Combinatorial Auctions*; MIT Press: Cambridge, UK, 2006; pp. 297–318.
15. Ausubel, L.M. *A Generalized Vickrey Auction*; University of Maryland: College Park, MD, USA, 1999.
16. Carrero, J.; Rodríguez, I.; Rubio, F. Measuring the benefits of lying in MARA under egalitarian social welfare. In Proceedings of the 2020 IEEE International Conference on Systems, Man, and Cybernetics (SMC), Toronto, ON, Canada, 11–14 October 2020; pp. 559–566.
17. Nguyen, N.; Nguyen, T.; Roos, M.; Rothe, J. Computational complexity and approximability of social welfare optimization in multiagent resource allocation. *Auton. Agents Multi-Agent Syst.* **2014**, *28*, 256–289. [\[CrossRef\]](#)
18. Roos, M.; Rothe, J. Complexity of Social Welfare Optimization in Multiagent Resource Allocation. In Proceedings of the AAMAS’10, Toronto, ON, Canada, 10–14 May 2010; pp. 641–648.
19. Paschos, V.T. An overview on polynomial approximation of NP-hard problems. *Yugosl. J. Oper. Res.* **2009**, *19*, 3–40. [\[CrossRef\]](#)
20. Muñoz, A.; Rubio, F. Evaluating genetic algorithms through the approximability hierarchy. *J. Comput. Sci.* **2021**, *53*, 101388. [\[CrossRef\]](#)
21. Rabanal, P.; Rodríguez, I.; Rubio, F. Applications of river formation dynamics. *J. Comput. Sci.* **2017**, *22*, 26–35. [\[CrossRef\]](#)