

Mechanisms for Fair Allocation Problems: No-Punishment Payment Rules in Verifiable Settings

Gianluigi Greco

*Dipartimento di Matematica e Informatica
Università della Calabria
I-87036 Rende, Italy*

GGRECO@MAT.UNICAL.IT

Francesco Scarcello

*DIMES
Università della Calabria
I-87036 Rende, Italy*

SCARCELLO@DIMES.UNICAL.IT

Abstract

Mechanism design is considered in the context of fair allocations of indivisible goods with monetary compensation, by focusing on problems where agents' declarations on allocated goods can be verified before payments are performed. A setting is considered where verification might be subject to errors, so that payments have to be awarded under the presumption of innocence, as incorrect declared values do not necessarily mean manipulation attempts by the agents. Within this setting, a mechanism is designed that is shown to be truthful, efficient, and budget-balanced. Moreover, agents' utilities are fairly determined by the Shapley value of suitable coalitional games, and enjoy highly desirable properties such as equal treatment of equals, envy-freeness, and a stronger one called individual-optimality. In particular, the latter property guarantees that, for every agent, her/his utility is the maximum possible one over any alternative optimal allocation.

The computational complexity of the proposed mechanism is also studied. It turns out that it is $\#P$ -complete so that, to deal with applications with many agents involved, two polynomial-time randomized variants are also proposed: one that is still truthful and efficient, and which is approximately budget-balanced with high probability, and another one that is truthful in expectation, while still budget-balanced and efficient.

1. Introduction

Whenever the outcome of some social choice process depends on the information collected from a number of self-interested agents, strategic issues may come into play. Indeed, agents may find it convenient to misreport their *types*, i.e., the relevant information they own as private knowledge, so that the (global) best possible solution can be missed. In these cases, *mechanism design* techniques can be used as solution approaches, which augment combinatorial algorithms with appropriate monetary payments, aimed at motivating all agents to truthfully report their private types (see, e.g., Nisan, Roughgarden, Tardos, & Vazirani, 2007; Shoham & Leyton-Brown, 2009).

On the class of social choice *utilitarian* problems, agent types encode (monetary) valuations over the set of all solutions and the goal is to compute a solution maximizing the *social welfare*, i.e., the sum of agents' true evaluations. A prominent role in mechanism design for problems of this class is played by the Vickrey-Clarke-Grove (VCG) paradigm (Vick-

ery, 1961; Clarke, 1971; Groves, 1973), which is a general method for designing *truthful* mechanisms, i.e., mechanisms where truth-telling is a dominant strategy for each agent. In particular, VCG mechanisms are *efficient*. That is, they guarantee that a solution maximizing the social welfare is actually computed. However, they are not *budget-balanced*, i.e., the algebraic sum of the monetary transfers is not always zero and mechanisms from this class can run into deficit. In fact, this is a well-known drawback of VCG mechanisms (see, e.g., Archer & Tardos, 2007), but it is essentially the best one can hope to do, given classical impossibility theorems (Green & Laffont, 1977; Hurwicz, 1975) stating that no truthful mechanism can be designed to be always efficient and budget-balanced.

In many practical applications, however, payments to agents can be performed *after* the final outcome is known, so that some kind of *verification* on reported types might be possible. This additional power is not considered in the classical mechanism-design setting and, in fact, whenever verification is allowed, some impossibility results might no longer hold. Mechanisms with verification have been introduced by Nisan and Ronen (2001), who considered verification for a task scheduling problem: We have some agents declaring the amount of time they need to solve each task, and the goal is to have all tasks being solved, by minimizing the completion time of the last-solved one (hence, the make-span). In this context, payments are computed after the actual task release times have been observed, so that we have, for instance, the ability to “punish” some agent whose declared ability has been verified to be different than its actual performance in the process.

Compared to standard mechanisms (see, e.g., Nisan et al., 2007), those with verification have received considerably less attention in the literature (see, e.g., Nisan & Ronen, 2001; Auletta, De Prisco, Penna, & Persiano, 2009; Penna & Ventre, 2012a; Krysta & Ventre, 2010; Ferrante, Parlato, Sorrentino, & Ventre, 2009; Penna & Ventre, 2012b; Auletta, De Prisco, Penna, Persiano, & Ventre, 2006; Auletta, Penna, Persiano, & Ventre, 2011). In particular, such works consider a verification ability that is *partial*, in the sense that agents’ reporting is restricted to true types plus certain specific kinds of deviations (e.g., values that are lower than the true ones) and verification is focused on detecting such lies only.

An extension of the above model has been recently proposed by Caragiannis, Elkind, Szegedy, and Yu (2012), who assume no a-priori restrictions on the agents’ reported types, within a setting where an agent cheating on her/his type will be caught with some probability that may depend on her/his true type, the reported type, or both. In fact, despite the different facets of the verification power, most of the mechanisms with verification proposed in the literature share the idea of providing incentives to truthfully report private types by exploiting the intimidation of punishing those agents that will be caught lying. Moreover, while budget limits have been considered in some approaches (see, e.g., Nisan & Ronen, 2001), no mechanism with verification has been designed to be budget-balanced, with the focus being on truthfulness and efficiency.

In this paper, we consider instead a budget-balanced mechanism based on a model of verification where there is no restriction on the possible declarations (hence, arbitrary deviations are possible), and nevertheless no punishment can be used after the verification process. This design constraint has been guided by real-world applications where it clearly emerges that a punishing approach would hardly be acceptable by agents, unless a clear proof of a deliberate malevolent behavior can be exhibited. Moreover, even in this case the punishment should be proportional to the amount of discrepancy between declared and

verified values that can be attributed to a malevolent behavior. The resulting setting shares the spirit of the work by Feige and Tennenholtz (2011), where it is observed that possible discrepancies between agents' declarations and third-party verified values are more often due to different reasons, in particular to the fact that agents, while not being malevolent, might still be unable to accurately collect and/or report information about their valuations.

In more detail, Feige and Tennenholtz (2011) considered a scheduling problem on a single machine where each agent reports the length of her/his job and the scheduler needs to finish as many jobs as possible by a given deadline. Differently from earlier literature, it is assumed that agents are uncertain of their own job lengths, for instance, because of their limited computational resources. Moving from the observation that mechanisms with verification are often designed in way that performs well when agents have accurate information about their private features, but might perform arbitrarily bad when agents are uncertain of this information, Feige and Tennenholtz then proposed the use of the “forgiving” mechanisms, where punishments are not used to enforce truthfulness. These mechanisms are applied over two models of uncertainty: One that is probabilistic in nature, and another (called “qualitative private input”) where there is no quantitative model explaining to which extent the agents can trust their estimate, and the preference of an agent over various lotteries might be even inconsistent with any probability distribution.

In this paper we follow the work by Feige and Tennenholtz (2011) and in particular their qualitative model of uncertainty. Moreover, in addition to their “subjective” perspective of the problem, where uncertainty is inherent to private inputs, we also take into account the dual (“objective”) perspective, where discrepancies between declared and verified values might due to errors that can occur in the verification process. Indeed, verification can be practically implemented by *sensing* some parameters that become observable after the mechanism is performed, and sensing is clearly affected by errors (it is unrealistic to assume that it can be carried out with arbitrary precision).

In fact, no matter the perspective (objective vs subjective) from which the problem discussed above is analyzed, an intrinsic limit of mechanisms with verification has clearly emerged: Whenever an agent misreports her/his type and this is detected by the verifier, “punishing” this agent might be effective in mathematical studies, but very inappropriate in real life situations in which uncertainty is inherent. Accordingly, we will therefore assume that only a limited use of the verification power given at hand can be made. In particular, the goal of the paper is to design mechanisms that are *not based on punishments* (while nonetheless resulting to be truthful, efficient, and budget-balanced) and that are *tolerant* of measurement errors and uncertain inputs, in the sense that “small” errors should determine small deviations from the outcome we would have obtained with no errors at all.

1.1 Mechanisms for Fair Division with Monetary Compensation

We consider mechanisms with verification in the context of *fair allocation* problems (see, e.g., Moulin, 2003; Young, 1994; Thomson, 2011). We assume that we are given a set of indivisible goods to be allocated to a set of agents. Each agent is equipped with a private preference relation, which is encoded as a real-valued function (basically, a monetary valuation) over all possible goods—formal definitions are in Section 2. An agent can have allocated more than one good, in which case her/his evaluation is additive over them. More-

over, goods are indivisible, i.e., each good can be allocated to one agent at most. However, monetary transfers are allowed, in terms of both payments charged to agents and monetary compensations provided to them. The goal is to find an efficient allocation, that is, an allocation maximizing the total value of the allocated goods, by designing rules guaranteeing that certain desirable properties are achieved, such as truthfulness and *individual rationality*, i.e., no agent is ever worse off than she/he would have been without participating in the mechanism. Moreover, we want to obtain outcomes that are “politically” acceptable. That is, agents should perceive the designed mechanism as a *fair* one (see, e.g., Brandt, Conitzer, & Endriss, 2012), independently of the rules leading them to be honest. For instance, it is desirable that no agent *envies* the allocation of any other agent, or that the selected outcome is *Pareto efficient*, i.e., there must be no different allocation that is preferred by all agents and strictly preferred by at least one of them.

The model and, in particular, properties of fair allocations with indivisible objects and monetary transfers have been studied, e.g., by Svensson (1983), Beviá (1998), Maskin (1987), Tadenuma and Thomson (1993), Meertens, Potters, and Reijnierse (2002), Tadenuma and Thomson (1991), Alkan, Demange, and Gale (1991), Willson (2003), Su (1999), Yang (2001), Quinzii (1984), and Sakai (2007). Moreover, procedures to compute fair allocations have been proposed by Aragonés (1995), Klijn (2000), Haake, Raith, and Su (2002), Brams and Kilgour (2001), Potthoff (2002), and Abdulkadiroğlu, Sönmez, and Ünver (2004).

None of the approaches listed above, however, can guarantee the elicitation of honest preferences from the agents. In fact, the question of designing truthful and fair mechanisms has been recently considered as well (Andersson & Svensson, 2008; Andersson, 2009; Svensson, 2009; Yengin, 2012; Ohseto, 2004; Porter, Shoham, & Tennenholtz, 2004; Shioura, Sun, & Yang, 2006). In these approaches, while budget limits are sometimes enforced and mechanisms are defined that cannot run into deficit, budget-balance is never guaranteed. Indeed, this comes again with no surprise, given that no truthful mechanism can be simultaneously fair (e.g., envy-free or Pareto-efficient) and budget-balanced (see, e.g., Tadenuma & Thomson, 1995; Alcalde & Barberà, 1994; Andersson, Svensson, & Ehlers, 2010).

To circumvent this impossibility, approaches have been studied that focus on weaker notions of truthfulness. For instance, Andersson et al. (2010) and Pathak (2013) consider a notion of degree of manipulability which can be used to compare the ease of manipulation in allocation mechanisms, whereas the notion of *weak* strategy-proofness is considered by Lindner (2010), i.e., cheating agents are always risking an actual loss, and are never guaranteed to cheat successfully.

In this paper, we depart from the settings studied in all such earlier approaches, because we are interested in applications where a form of verification is available to the mechanism at the time of deciding monetary compensations among agents. In particular, we assume that valuations as well as allocation scenarios are determined by objective properties of goods and agents that can be observed and measured by a verifier, after an allocation is performed and payments are to be computed. Note that only information on allocated goods can be verified and hence used by the mechanism. In this framework, classical impossibility results no longer hold. Indeed, we propose mechanisms for allocation problems that enjoy a number of highly desirable properties, being in particular truthful, efficient, budget-balanced, individually rational, and fair, even though agents with verified incorrect declarations are not punished. Observe that having this kind of *a-posteriori* knowledge

at payment time is quite common to many applications. We also point out that in some cases a thorough verification could also be performed in advance, in order to get the best performances independently of agents' declarations. However, in practice this is not done because of either money or time restrictions, so that it is more convenient to allocate goods on the basis of agents' declarations (especially if a mechanism makes them honest enough). Anyway, our results can be used even when the full information is known to the mechanism, to provide a fair division enjoying a number of desirable properties listed below.

Appendix A reports a number of examples of possible applications of the proposed framework, including the real-world case of the Italian research-assessment program, which first motivated this work.

For completeness, we leave the section by recalling that our work, as well as the above mentioned related literature, deals with a setting where monetary transfers are allowed. In fact, fair division of indivisible goods without money transfers has also attracted attention in the literature. For instance, this topic has been studied by Bouveret and Lang (2008) from the points of view of compact representation (for expressing preference relations) and computational complexity (of reasoning about efficiency and fairness concepts in the resulting framework), and by Lipton, Markakis, Mossel, and Saberi (2004) from the point of view of defining approximation schemes for envy-freeness. Finally, it is relevant to point out that, in our paper and in the papers discussed above, allocations are assumed to be computed in a centralized way. However, it might be relevant in some cases to adopt decentralized approaches, based on successive negotiations of goods (and of money) between groups of agents. The reader interested in distributed negotiation frameworks is referred to the work by Sandholm (1998), Dunne, Wooldridge, and Laurence (2005), Dunne (2005), and Endriss, Maudet, Sadri, and Toni (2006), and to the references therein.

1.2 Contributions

In this paper, we study allocation problems in a strategic setting where agents can misreport their private types, and we study mechanisms with verification from both the algorithmic and the computational complexity viewpoint.

1.2.1 ALGORITHMIC ISSUES

We show that in the given setting none of the classical impossibility theorems discussed above holds. In particular, we exhibit a payment rule \mathbf{p}^ξ that turns any optimal allocation algorithm, i.e., any algorithm computing an optimal allocation given the reported types, into a mechanism with verification such that:

- ▶ The mechanism is truthful. This is shown by pointing out a number of properties of allocation problems which are of interest on their own.
- ▶ The mechanism is efficient, budget-balanced, individually rational, envy-free, and Pareto efficient.
- ▶ The payment rule is indifferent w.r.t. the values (possibly misreports) declared for goods that do not occur in the allocation being selected (and hence that are not verified).

- ▶ For each agent, her/his utility (when truth-telling) is the maximum one over all possible allocations. In particular, the utility is indifferent w.r.t. the specific choice of allocated goods in optimal allocations. Note that this is a strong fairness property, which immediately entails envy-freeness and Pareto-efficiency.
- ▶ Verification is not used to force truthfulness by just punishing those agents whose reported values are found different from the verified ones, so that the mechanism is “forgiving” in the sense recently discussed by Feige and Tennenholtz (2011). Moreover, the mechanism is shown to be tolerant of discrepancies emerging between declared types and verified/true ones. That is, all its properties hold at the equilibrium where all agents report their true types, and are also preserved approximatively in case of discrepancies, with a guarantee that is within a constant factor from the “distance” between declared types and verified/true ones. Note that this is generally not possible in mechanisms based on punishment approaches where, to enforce truthfulness, the punishment might be disproportional to the harm done by misreporting (cf. Feige & Tennenholtz, 2011).
- ▶ Agents’ utilities are distributed according to the *Shapley value* of two suitably-associated *coalitional games*—see, e.g., (Nisan et al., 2007), for a comprehensive introduction to sharing problems and coalitional games. In fact, the Shapley value is a prototypical solution concept for fair division with monetary compensations,¹ and its desirable properties in (games associated with) allocation problems have been extensively studied in the literature (e.g., Moulin, 1992; Maniquet, 2003; Mishra & Rangarajan, 2007).

Note that the Shapley value has been studied in mechanism-design contexts too, where emphasis has been given to the pricing problem for a service provider (Moulin & Shenker, 2001; Moulin, 1999; Jain & Vazirani, 2001): The cost of providing a service is a function of the sets of customers, and the goal is that of determining which customers (and at what price) have to receive it. The model gives rise to a *cross-monotonic* cost-sharing game, where Shapley-value based sharing mechanisms can be defined that are truthful and budget-balanced, and which achieve the lowest worst-case loss of efficiency over all utility profiles (Moulin & Shenker, 2001). With this respect, our pricing rule \mathbf{p}^ξ can abstractly be viewed as a witness that, whenever (partial) verification is possible, Shapley-valued based mechanisms may also be implemented with no loss of efficiency at all.

1.2.2 COMPLEXITY ISSUES

Computing an optimal allocation on the basis of the reported types is an easy task, which can be carried out via adaptations of classical matching algorithms. However, one might suspect that computing payments is not computationally-efficient, as it is based on the computation of a Shapley value. This is indeed a challenging task that involves iterating over all possible subsets of agents. We analyze these issues, and we provide the following contributions:

- ▶ We show that computing the Shapley value for allocation problems is inherently intractable, in fact, $\#P$ -complete. Note that $\#P$ -hardness results for problems involving

1. Depending on the application, solutions concepts different from the Shapley value might be more appropriate. For instance, in bankruptcy problems, the *nucleolus* is considered as the most appropriate solution concept for fair distribution (Aumann & Maschler, 1985).

Shapley value computation have been proven in the literature, for instance, for *weighted voting games* (Deng & Papadimitriou, 1994), *minimum spanning-tree games* (Nagamochi, Zeng, Kabutoya, & Ibaraki, 1997), and games associated with *normative systems* (Ågotnes, van der Hoek, Tennenholtz, & Wooldridge, 2009). Moreover, #P-hardness results have been established for the *Banzhaf power index*, which is a solution concept closely related to the Shapley value (see, e.g., Bachrach & Rosenschein, 2009, 2008; Bachrach, Zuckerman, Wooldridge, & Rosenschein, 2013).

- Therefore, in order to deal also with scenarios involving a large number of agents, two modified rules, $\hat{\mathbf{p}}^\xi$ and $\bar{\mathbf{p}}^\xi$, are presented, which allow us to employ a fully polynomial-time randomized approximation scheme for the Shapley value computation. The resulting polynomial-time mechanisms retain most of the properties of \mathbf{p}^ξ . In particular, the mechanism based on $\hat{\mathbf{p}}^\xi$ is universally truthful, efficient, and with high-probability approximately budget-balanced. Instead, the mechanism based on $\bar{\mathbf{p}}^\xi$ is truthful in expectation, but it is always efficient and budget-balanced.

1.2.3 ORGANIZATION

The rest of the paper is organized as follows. Section 2 illustrates the formal framework and the basic concepts to design mechanisms with verification, whose desirable properties are illustrated in Section 3. The payment rule \mathbf{p}^ξ is defined in Section 4, and its connections with coalitional games are pointed out in Section 5. Rules $\hat{\mathbf{p}}^\xi$ and $\bar{\mathbf{p}}^\xi$ are defined in Section 6, where computational issues are dealt with. A comparison with related works is reported in Section 7, and a few concluding remarks are discussed in Section 8. Finally, Appendix A illustrates a real-world case study and further application examples of the notions presented in the paper.

2. Formal Framework

In this section, we define a formal framework for studying allocation problems based on mechanism design tools. In particular, we focus on mechanisms equipped with a verification ability that meets the “no-punishment” perspective.

2.1 Allocation Scenarios

We focus on allocation problems where goods from a set G have to be allocated to a set of agents $\mathcal{A} = \{1, \dots, n\}$ in such a way that the overall value of the allocated goods is maximum over all feasible allocations, that is, the *social welfare* is maximized. More precisely, an *allocation* is a function $\pi : \mathcal{A} \rightarrow 2^G$ mapping each agent $i \in \mathcal{A}$ into a non-empty set of goods $\pi(i) \subseteq G$ such that $\pi(i) \cap \pi(j) = \emptyset$, for each $j \neq i$. Moreover, we are given a vector of upper-bound constraints $\zeta = (\zeta_1, \dots, \zeta_n)$ that specifies the maximum number of goods ζ_i that can be assigned to any agent $i \in \mathcal{A}$. The tuple $S = \langle \mathcal{A}, G, \zeta \rangle$ is called an *allocation scenario*. A mapping $\pi : \mathcal{A} \rightarrow 2^G$ is a *feasible allocation* for the scenario S if it is an allocation of goods in G to agents in \mathcal{A} such that, for each agent $i \in \mathcal{A}$, $|\pi(i)| \leq \zeta_i$ holds. Observe that it is not required that all goods from G are allocated to the agents.

Note that in most applications the maximum number of goods ζ_i that can be assigned to agent i represents some ability of i (e.g., the maximum number of tasks that she/he

can execute), hence we next represent it as a function $\zeta_i = f_u(\chi_i)$, where χ_i is an objective property of the agent (e.g., her/his speed) and f_u is a public-knowledge computable function. Note that the setting is more general than earlier approaches in the literature for which upper bounds are fixed and independent of agents' features.

Moreover, in this paper we assume that the value of each good g for agent i is determined by some objective property λ_g of the good, as well as by the property χ_i of the agent. Formally, we assume that good valuations are encoded by a *valuation vector* $\mathbf{w} = (w_1, \dots, w_n)$ where, for each $i \in \mathcal{A}$, i 's valuation function assigns to any good $g \in G$ a real value $w_i(g) = f_v(\lambda_g, \chi_i)$, for some public-knowledge computable function f_v .²

The idea is that a verifier, best described in the next section, should be able to measure, after an allocation π of goods to agents has been performed, the objective property λ_g of every allocated good g and the objective property χ_i of every agent. Therefore, by using the (public) function f_v , it is then possible to compute the values of all those goods that have been assigned to some agent. We believe that this assumption about valuations as functions of objective properties of goods and agents holds for many applications of allocation problems, in particular for those where the social welfare is to be maximized (since this is typically a measurable value). We provide some examples in Appendix A, including the real-world application about the evaluation of the research activities in Italy, which originally motivated the present work.

Let us fix an allocation scenario $S = \langle \mathcal{A}, G, \zeta \rangle$ and a valuation vector $\mathbf{w} = (w_1, \dots, w_n)$. Let π be an allocation. Define $w_i(\pi) = \sum_{g \in \pi(i)} w_i(g)$, for each $i \in \mathcal{A}$, and denote by $\text{val}(\pi, \mathbf{w})$ the overall value $\sum_{i \in \mathcal{A}} w_i(\pi)$. We say that π is *optimal* (for S) w.r.t. \mathbf{w} if it is a feasible allocation for S and there is no feasible allocation π' for S such that $\text{val}(\pi', \mathbf{w}) > \text{val}(\pi, \mathbf{w})$. The value of an optimal allocation for S w.r.t. \mathbf{w} is denoted by $\text{opt}(S, \mathbf{w})$.

An *allocation algorithm* is a function A mapping each allocation scenario S and each vector \mathbf{w} to a feasible allocation $A(S, \mathbf{w})$ for S . The algorithm is *optimal* if $A(S, \mathbf{w})$ is an optimal allocation w.r.t. \mathbf{w} , for any given pair (S, \mathbf{w}) .

2.2 Strategic Issues and Verification

We consider a classical setting for mechanism design where optimal allocations have to be computed in a context where neither the agent-depending upper bounds ζ nor the valuation vector \mathbf{w} is known to the allocation algorithm. Therefore, even having an optimal algorithm A at hand, we do not have enough information to find an optimal allocation, in general.

In fact, we assume as usual that each agent $i \in \mathcal{A}$ privately knows certain features, called i 's *type*, determining both the maximum number of goods ζ_i that can be allocated to her/him, as well as the function w_i encoding her/his preferences over the allocations. Note that, in our setting, the type of agent i naturally consists of her/his characterizing property χ_i plus the property λ_g , for any good g (s)he is interested in. In Section 3.2, we also consider what happens when agent i may have a subjective, possibly incorrect, perception of such properties (including her/his characterizing property χ_i).

2. Note that, for the sake of simplicity, we use two functions f_u and f_v for all agents. However, nothing changes in the paper if we consider a slightly more general version where different agents may have different public functions.




Figure 1: Running example in Section 2.

We assume that the type of agent i is taken from a set Θ_i of available types, and we denote by Θ the cartesian product $\Theta_1 \times \dots \times \Theta_n$ of all possible agents' types. Then, we consider *direct revelation mechanisms* where agents are asked to report such types to let the mechanism compute an allocation. While doing so, agents are self-interested, and strategic issues come into play.

For each agent $i \in \mathcal{A}$, we hereinafter assume that $t_i \in \Theta_i$ always denotes the *true* type of agent i , i.e., the type owned as private knowledge, and that $d_i \in \Theta_i$ is her/his *declared* type. Then, $\mathbf{t} = (t_1, \dots, t_n)$ and $\mathbf{d} = (d_1, \dots, d_n)$ are the vectors of true and declared types, respectively. In general, it can happen that \mathbf{d} does not coincide with \mathbf{t} , if agents find convenient to misreport their types.

For any (true or declared) type vector $\theta = (\theta_1, \dots, \theta_n) \in \Theta$, we denote hereinafter by $S_\theta = \langle \mathcal{A}, G, \zeta_\theta \rangle$ the *allocation scenario* for θ , that is, the scenario where the vector $\zeta_\theta = (\zeta_1, \dots, \zeta_n)$ is such that, for each $i \in \mathcal{A}$, ζ_i is the upper bound determined by i 's type θ_i . Similarly, we denote by $\mathbf{w}_\theta = (w_1, \dots, w_n)$ the valuation vector such that w_i is the valuation function determined by i 's type θ_i , for each $i \in \mathcal{A}$.

Example 2.1. Let us consider two agents, a_1 and a_2 , the type vector $\mathbf{t} = (t_1, t_2)$, and the allocation scenario $S_{\mathbf{t}} = \langle \mathcal{A}, G, \zeta_{\mathbf{t}} \rangle$ illustrated in Figure 1(I) by means of an intuitive graphical notation, with $\mathcal{A} = \{a_1, a_2\}$, $G = \{g_1, \dots, g_8\}$, and $\zeta_{\mathbf{t}} = (3, 3)$. Moreover, consider the valuation vector \mathbf{w} with $\mathbf{w}_{\mathbf{t}} = (w_{a_1}, w_{a_2})$ such that, if there is an edge connecting a_1 (resp., a_2) with g_j in Figure 1(I), then $w_{a_1}(g_j)$ (resp., $w_{a_2}(g_j)$) is the value associated with it. Otherwise, i.e., if there is no edge connecting a_1 (resp., a_2) with g_j , then $w_{a_1}(g_j)$ (resp., $w_{a_2}(g_j)$) is some negative number, say -1 . Given this setting, it is easily seen that an optimal allocation for $S_{\mathbf{t}}$ w.r.t. $\mathbf{w}_{\mathbf{t}}$ is the allocation π^* where $\pi^*(a_1) = \{g_1, g_2, g_4\}$ and $\pi^*(a_2) = \{g_5, g_7, g_8\}$ —see Figure 1(II). Note that $w_{a_1}(\pi^*) = 25$ and $w_{a_2}(\pi^*) = 26$.

Consider now the allocation $\hat{\pi}^*$ of Figure 1(III). Note that $\hat{\pi}^*$ is another optimal allocation. However, we have $w_{a_1}(\hat{\pi}^*) = 26 > w_{a_1}(\pi^*)$ and $w_{a_2}(\hat{\pi}^*) = 25 < w_{a_2}(\pi^*)$. \triangleleft




Figure 2: Strategic manipulation.

Example 2.2. Consider again the setting of Example 2.1, where the type vector \mathbf{t} is private knowledge of the agents. Moreover, assume that the vector \mathbf{d} of declared types is such that $S_{\mathbf{d}} = S_{\mathbf{t}}$ (the allocation scenario is the correct one) and that the vector $\mathbf{w}_{\mathbf{d}}$ is the one illustrated in Figure 2(I), where negative edges are omitted. Basically, agent a_2 truthfully reports her/his type, while agent a_1 reports a type for which the values of the goods g_2 and g_3 are underestimated. Therefore, $\mathbf{w}_{\mathbf{d}} \neq \mathbf{w}_{\mathbf{t}}$ holds. The only optimal allocation $\bar{\pi}^*$ for $S_{\mathbf{d}}$ w.r.t. $\mathbf{w}_{\mathbf{d}}$ is shown in Figure 2(II). It emerges that, because of the declarations of agent a_1 , it is not convenient to include g_2 and g_3 in $\bar{\pi}^*$. In fact, $\bar{\pi}^*$ coincides with the optimal allocation $\hat{\pi}^*$ depicted in Figure 1(III). Hence, by misreporting the type, agent a_1 has now the guarantee that the overall value of the goods assigned to her/him is 26. Instead, by truthfully reporting the type, a_1 might risk that the allocation $\hat{\pi}^*$ is selected, where the overall value of the goods assigned to her/him is only 25.

Finally, consider a slight variant of the problem instance where the actual value of good g_7 is 6 (instead of 8) for a_2 . Then, the above egoistic behavior of agent a_1 also leads to an allocation that is no longer optimal. Indeed, due to the low declared values for g_2 and g_3 , the good g_7 is selected and allocated to a_2 in the unique (wrong) optimal allocation, whose total value is now 49 (instead of 51). \triangleleft

In this paper, we focus on allocation problems where some kind of *verification* on reported types is possible, because objective properties of agents and goods can be observed and measured (we say *verified*, hereinafter) by a third-party verifier after an allocation π is performed. Recall that in general only a subset of goods is actually assigned to some agent by π . Denote the set $\bigcup_{i \in \mathcal{A}} \pi(i) \subseteq G$ of allocated goods by $\text{img}(\pi)$. Thus, in our model, for each good $g \in \text{img}(\pi)$, its objective property λ_g can be verified, while nothing can be said for non-allocated and hence non-observed goods. Moreover, for each agent $i \in \mathcal{A}$ participating in the mechanism, i 's property χ_i that determines the upper-bound constraint and i 's preferences over good allocations is verifiable, too. Thus, by using such a verifier, the correct allocation scenario and a restriction over $\text{img}(\pi)$ of the valuation functions can be determined, as formalized next.

Definition 2.3. Let \mathbf{t} be the vector of true types. Then, the *verifier* \mathbf{v} (for \mathbf{t}) is the function mapping any allocation π to the pair $S_{\mathbf{v}(\pi)}, \mathbf{w}_{\mathbf{v}(\pi)}$ such that:

- (1) $S_{\mathbf{v}(\pi)}$ is the actual allocation scenario $S_{\mathbf{t}}$ (in particular, the “true” upper bound $f_u(\chi_i)$ is computed by the verifier for each agent i); and
- (2) $\mathbf{w}_{\mathbf{v}(\pi)} = (w_{v_1}, \dots, w_{v_n})$ is the vector such that, for each agent $i \in \mathcal{A}$, $w_{v_i} : \text{img}(\pi) \rightarrow \mathbb{R}$ is the function assigning to each good $g \in \text{img}(\pi)$ its actual value $f_v(\lambda_g, \chi_i)$ for agent

i. Observe that, by definition of the framework, w_{v_i} coincides with the restriction over $\text{img}(\pi)$ of the valuation function $w_{t_i} : G \rightarrow \mathbb{R}$ determined by i 's (true) type t_i . \square

It is worthwhile noting that the verifier is in general unable to discover whether some agent misreported her/his type, because goods that are not allocated do not undergo any measurement process. We pinpoint that this is actually the case in practical applications, where measurements over non-allocated goods may be technically unfeasible or simply too expensive (in money or time). We refer the reader to Appendix A.1 (in particular, subsections A.1.1–A.1.3) for an exemplification of these notions in the real-world case study of the Italian research evaluation (VQR).

Example 2.4. Consider again the setting of Example 2.2, and recall that agent a_1 finds it convenient to underestimate the true values of g_2 and g_3 . However, since g_2 and g_3 are not selected in $\bar{\pi}^*$, as we can see in Figure 2(II), then there is no way to discover that a_1 has actually misreported her/his type. \triangleleft

In some cases, the constraints in the allocation scenario depend only on the specific application and not on agents' types, so that item (1) above is immaterial. More generally, however, the proposed setting allows us to model classes of problems where types play a role even in the definitions of upper-bound constraints.

For completeness, we remark that all results in the paper can easily be shown to hold even for allocation problems where every agent must get a minimum number of goods greater than one (defined by the specific application). However, for the sake of presentation we prefer to keep the standard setting where any non-empty set of goods can be assigned to each agent, as long as her/his upper bound constraint is met.

2.3 Payment Rules and Mechanisms with Verification

In order to encourage agents to truthfully report their private types, we design mechanisms where monetary transfers can be performed, after the verification process.

For the sake of presentation, let us assume that $S_{\mathbf{t}} = \langle \mathcal{A}, G, \zeta_{\mathbf{t}} \rangle$ is an allocation scenario (recall that \mathbf{t} denotes the vector of true types), that \mathbf{d} is a vector of declared types with an associated allocation scenario $S_{\mathbf{d}} = \langle \mathcal{A}, G, \zeta_{\mathbf{d}} \rangle$, and that \mathbf{v} is the verifier (for \mathbf{t}).

A *payment rule* \mathbf{p} is defined as a vector of functions (p_1, \dots, p_n) , with $p_i(\pi, \mathbf{w}_{\mathbf{d}}, \mathbf{v})$ being some amount of money that is given to agent i , on the basis of a given allocation π , the vector of declared valuations $\mathbf{w}_{\mathbf{d}}$, and the verifier \mathbf{v} . Observe that, with this notation, any negative value $p_i(\pi, \mathbf{w}_{\mathbf{d}}, \mathbf{v})$ means that some amount of money is charged to i . Let $\mathbf{w}_{\mathbf{t}}$ be the vector (w_1, \dots, w_n) . Then, i 's (quasi-linear) *utility* under \mathbf{p} , sometimes called *individual welfare*, is defined as the value $u_{i,\mathbf{p}}(\pi, \mathbf{w}_{\mathbf{d}}, \mathbf{v}) = w_i(\pi) + p_i(\pi, \mathbf{w}_{\mathbf{d}}, \mathbf{v})$. As the verifier \mathbf{v} for \mathbf{t} is always understood, $u_{i,\mathbf{p}}(\pi, \mathbf{w}_{\mathbf{d}}, \mathbf{v})$ and $p_i(\pi, \mathbf{w}_{\mathbf{d}}, \mathbf{v})$ are simply denoted by $u_{i,\mathbf{p}}(\pi, \mathbf{w}_{\mathbf{d}})$ and $p_i(\pi, \mathbf{w}_{\mathbf{d}})$, respectively. Moreover, whenever the payment rule is also understood from the context, i 's utility is simply denoted as $u_i(\pi, \mathbf{w}_{\mathbf{d}})$.

As payments can be computed after the verification process, to define the amount of money $p_i(\pi, \mathbf{w}_{\mathbf{d}})$ to be paid to agent i , we exploit the verifier \mathbf{v} . Accordingly, it is desirable that goods that are not allocated, and hence not verified, play no role in the definition of the payments. This latter property is formalized below.

▷ *verifiability*: Let \mathbf{d}' and \mathbf{d}'' be two type vectors, and let $\mathbf{w}_{\mathbf{d}'} = (w'_1, \dots, w'_n)$ and $\mathbf{w}_{\mathbf{d}''} = (w''_1, \dots, w''_n)$ be their associated valuations. Moreover, let π be any mapping that is a feasible allocation for both scenarios $S_{\mathbf{d}'}$ and $S_{\mathbf{d}''}$. Then, for each $i \in \mathcal{A}$, $p_i(\pi, \mathbf{w}_{\mathbf{d}'}) = p_i(\pi, \mathbf{w}_{\mathbf{d}''})$ whenever $w'_i(g) = w''_i(g)$ holds for every allocated good $g \in \text{img}(\pi)$. Therefore, \mathbf{d}' and \mathbf{d}'' are undistinguishable as far as the computation of the payments is concerned, even if they differ on some unallocated goods. That is, the payment rule depends only on goods subject to the verifier evaluation.

A *mechanism with verification* is a pair (A, \mathbf{p}) , where A is an allocation algorithm and \mathbf{p} is a payment rule that can exploit the verifier \mathbf{v} . The mechanism (A, \mathbf{p}) can be viewed as consisting of the following two-phases: First, agents report a declaration vector \mathbf{d} , and a feasible allocation $\pi = A(\mathbf{w}_{\mathbf{d}})$ for $S_{\mathbf{d}}$ is computed. Second, $\mathbf{v}(\pi)$ is made available, and payments under a given rule \mathbf{p} are calculated with respect to the allocation π and the valuations $\mathbf{w}_{\mathbf{d}}$, by exploiting the knowledge of $\mathbf{v}(\pi)$. Our goal is to design a payment rule \mathbf{p} guaranteeing that declared types in \mathbf{d} lead to an allocation π maximizing the social welfare, i.e., such that π is an optimal allocation for $S_{\mathbf{t}}$ w.r.t. $\mathbf{w}_{\mathbf{t}}$. This might be problematic as, in our setting, even the fact that π is a feasible allocation for $S_{\mathbf{t}}$ is not guaranteed, because the allocation scenario depends on the types of the agents and we might have $S_{\mathbf{t}} \neq S_{\mathbf{d}}$.

In order to accomplish the above goal, we use an optimal allocation algorithm A , and we need that \mathbf{p} encourages agents to truthfully report their private types. Formally, for any type vector $\boldsymbol{\theta} = (\theta_1, \dots, \theta_n) \in \Theta$ and for any type $\bar{\theta}_i \in \Theta_i$, with $i \in \mathcal{A}$, let $(\bar{\theta}_i, \boldsymbol{\theta}_{-i})$ be the type vector $(\theta_1, \dots, \theta_{i-1}, \bar{\theta}_i, \theta_{i+1}, \dots, \theta_n) \in \Theta$. Then, we shall consider the following concept of truthful mechanism.

Definition 2.5. Let (A, \mathbf{p}) be a mechanism with verification, and let i be any agent in \mathcal{A} . We say that $\bar{\theta}_i$ is a *dominant strategy* of agent i w.r.t. (A, \mathbf{p}) if, for each type vector $\boldsymbol{\theta} \in \Theta$, $u_i(A(\mathbf{w}_{(\bar{\theta}_i, \boldsymbol{\theta}_{-i})}), \mathbf{w}_{(\bar{\theta}_i, \boldsymbol{\theta}_{-i})}) \geq u_i(A(\mathbf{w}_{\boldsymbol{\theta}}), \mathbf{w}_{\boldsymbol{\theta}})$ holds. The mechanism (A, \mathbf{p}) is *truthful* if, for each $i \in \mathcal{A}$, t_i is a dominant strategy. \square

Example 2.6. Consider again the setting discussed in Example 2.2, and the trivial payment rule \mathbf{p}° where no payment is actually performed. Consider again the optimal allocation π^* (for $S_{\mathbf{t}}$ w.r.t. $\mathbf{w}_{\mathbf{t}}$) depicted in Figure 1(II), and note that $u_{a_1}(\pi^*, \mathbf{w}_{\mathbf{t}}) = 25$. Instead, for the allocation $\bar{\pi}^*$ depicted in Figure 2(II), we have $u_{a_1}(\bar{\pi}^*, \mathbf{w}_{\mathbf{t}}) = 26$, with $\bar{\pi}^*$ being the unique optimal allocation for $S_{\mathbf{d}} = S_{\mathbf{t}}$ w.r.t. $\mathbf{w}_{\mathbf{d}}$, where $\mathbf{d} = (d_{a_1}, \mathbf{t}_{-a_1})$.

Therefore, any mechanism (A, \mathbf{p}°) , with A being an optimal allocation algorithm such that $A(\mathbf{w}_{\mathbf{t}}) = \pi^*$, is not truthful. More generally, for each optimal allocation algorithm A , an example witnessing that (A, \mathbf{p}°) is not truthful can be easily defined by suitably adapting the above. Hence, non-trivial payment rules are necessary to encourage agents to truthfully report their types. \triangleleft

A comparison of our approach to verification with existing ones is reported in Section 7.

3. Properties of (Truthful) Mechanisms

In this section, we discuss desiderata for mechanisms that lead to fair allocations and that are tolerant of uncertain inputs. Note that the design of mechanisms that are able to deal with these two issues is of interest even in settings that are not strategic, i.e., even

when we are granted that all agents truthfully report their types. Exemplifications of the proposed notions are reported Appendix A.1.4, which completes the description of the case study regarding the Italian research evaluation (VQR). Further examples are described in Appendix A.2 and Appendix A.3.

3.1 Fairness Issues and Further Desirable Properties

Let (A, \mathbf{p}) be any truthful mechanism with verification. In the paper, we focus on a number of (ex-post) properties of such a mechanism, to be checked at the equilibrium \mathbf{t} where agents truthfully report their private types.

- ▷ *(allocative) efficiency*: $A(\mathbf{w}_{\mathbf{t}})$ is an optimal allocation for $S_{\mathbf{t}}$ w.r.t. $\mathbf{w}_{\mathbf{t}}$. That is, the social welfare is maximized.
- ▷ *individual rationality*: $u_i(A(\mathbf{w}_{(t_i, \boldsymbol{\theta}_{-i})}), \mathbf{w}_{(t_i, \boldsymbol{\theta}_{-i})}) \geq 0$, for each agent $i \in \mathcal{A}$ and for each type vector $\boldsymbol{\theta}_{-i}$ for the agents in $\mathcal{A} \setminus \{i\}$. Hence, voluntary participation of each agent to take part in the allocation problem is encouraged (independently of whether the other agents actually report their true types or not).
- ▷ *(strong) budget-balance*: $\sum_{i \in \mathcal{A}} p_i(A(\mathbf{w}_{\mathbf{t}}), \mathbf{w}_{\mathbf{t}}) = 0$. In other words, there is no transfer of money out or into the scenario.
- ▷ *envy-freeness*: for each pair of agents $i, j \in \mathcal{A}$, and for each feasible allocation π for $S_{\mathbf{t}}$ such that $\pi(i) \cap A(\mathbf{w}_{\mathbf{t}})(j) \neq \emptyset$, $u_i(A(\mathbf{w}_{\mathbf{t}}), \mathbf{w}_{\mathbf{t}}) \geq u_i(\pi, \mathbf{w}_{\mathbf{t}})$.
- ▷ *Pareto-efficiency*: there is no feasible allocation π for $S_{\mathbf{t}}$ such that: (1) $u_i(\pi, \mathbf{w}_{\mathbf{t}}) \geq u_i(A(\mathbf{w}_{\mathbf{t}}), \mathbf{w}_{\mathbf{t}})$, for each agent $i \in \mathcal{A}$, and (2) there is an agent $j \in \mathcal{A}$ with $u_j(\pi, \mathbf{w}_{\mathbf{t}}) > u_j(A(\mathbf{w}_{\mathbf{t}}), \mathbf{w}_{\mathbf{t}})$. That is, $A(\mathbf{w}_{\mathbf{t}})$ is not Pareto-dominated by any other allocation.
- ▷ *equal treatment of equals*: for each pair of agents $i, j \in \mathcal{A}$ such that $\zeta_i = \zeta_j$ and $w_i = w_j$, with $\mathbf{w}_{\mathbf{t}} = (w_1, \dots, w_n)$, it must be the case that $u_i(A(\mathbf{w}_{\mathbf{t}}), \mathbf{w}_{\mathbf{t}}) = u_j(A(\mathbf{w}_{\mathbf{t}}), \mathbf{w}_{\mathbf{t}})$.
- ▷ *individual optimality*: $u_i(A(\mathbf{w}_{\mathbf{t}}), \mathbf{w}_{\mathbf{t}}) \geq u_i(\pi, \mathbf{w}_{\mathbf{t}})$, for each $i \in \mathcal{A}$ and each feasible allocation π for $S_{\mathbf{t}}$.

Note that all the above properties make sense even in settings that are not strategic, but where the goal is to compute a fair allocation. Indeed, in a setting that is not strategic, agents report their true types, even without any monetary incentive. However, without such payments, fairness cannot be achieved in general: just think, for instance, about a scenario with two agents, a_1 and a_2 , and one good g having the same value for both agents. In this case, no matter which optimal allocation is considered (where g is assigned to either a_1 or a_2), one of the two agents would envy the other. This makes it clear that payment rules play not only the role to encourage agents to reports their true types, but they are also crucial to induce agents to perceive a given allocation as a fair one. In fact, all the above properties, but the last, have been classically considered in the context of fair allocation problems, also in absence of strategic issues.

Here, we have additionally considered *individual optimality*, which is readily seen to imply both envy-freeness and Pareto-efficiency. It also entails that there is a unique possible

vector of utilities for agents. In particular, this means that agents' utilities are not sensible to possible alternative allocations, and hence are independent of the specific set of allocated goods selected by the optimal algorithm A .

Example 3.1. Consider the trivial payment rule \mathbf{p}° discussed in Example 2.6. Consider the optimal allocation π^* depicted in Figure 1(II), and compare it with the allocation $\hat{\pi}^*$ of Figure 1(III). Recall that $\hat{\pi}^*$ is another optimal allocation (for S_t w.r.t. \mathbf{w}_t). Moreover, under \mathbf{p}° , $u_{a_1}(\pi^*, \mathbf{w}_t) = 25$ and $u_{a_1}(\hat{\pi}^*, \mathbf{w}_t) = 26$ hold. Therefore, while from the optimization perspective the choice between π^* and $\hat{\pi}^*$ is immaterial, a_1 might have good arguments to complain if π^* is selected in place of $\hat{\pi}^*$. In fact, \mathbf{p}° is not “individually optimal”. \triangleleft

Individual optimality is definitely a very desirable requirement, but we still miss something. Indeed, notice that this property is trivially satisfied by the fully “uniform” payment rule, which guarantees that each agent gets the same utility, no matter of her/his valuation of goods. Of course, this is not desirable in general. Rather, meritocracy should be somehow addressed, so that a true fair rule should reflect the actual “contribution” of each agent to the overall value of the allocation.

Let π be an allocation for S_t , not necessarily an optimal one w.r.t. \mathbf{w}_t , and define the *marginal contribution* of a non-empty set $\mathcal{C} \subseteq \mathcal{A}$ of agents to π w.r.t. \mathbf{w}_t as the value:

$$\text{marg}_{\pi, \mathbf{w}_t}(\mathcal{C}) = \text{opt}(\langle \mathcal{A}, \text{img}(\pi), \zeta_t \rangle, \mathbf{w}_t) - \text{opt}(\langle \mathcal{A} \setminus \mathcal{C}, \text{img}(\pi), \zeta_t \rangle, \mathbf{w}_t). \quad (1)$$

In words, the marginal contribution $\text{marg}_{\pi, \mathbf{w}_t}(\mathcal{C})$ of the agents in \mathcal{C} assesses the loss of the overall value of π we would register if the agents in \mathcal{C} were not part of the problem.

Example 3.2. Consider the setting of Example 2.1 and the optimal allocation π^* of Figure 1(II). Note that $\text{marg}_{\pi^*, \mathbf{w}_t}(\{a_1\}) = \text{marg}_{\pi^*, \mathbf{w}_t}(\{a_2\}) = 51 - 26 = 25$. Therefore, the two agents (viewed as singletons) have the same marginal contribution, and defining a payment rule that leads to utilities equally sharing the overall value of 51 is a natural option.

Consider now a different setting, where true types induce the same allocation scenario S_t and the same vector \mathbf{w}_t , but for the valuation of good g_8 for agent a_2 , which is now 110 instead of 10. In this case, π^* would still remain an optimal allocation, with overall value 151. Moreover, the marginal contribution of a_1 is not affected by the modification, while the marginal contribution of a_2 would become $151 - 26 = 125$. This witnesses that a_2 contributes more than a_1 , and a payment rule leading to equally sharing the overall value can be no longer perceived as a fair one. \triangleleft

The intuition conveyed by the above example can be then formalized via the following requirement.

\triangleright *marginality:* For each non-empty set $\mathcal{C} \subseteq \mathcal{A}$ of agents, $\sum_{i \in \mathcal{C}} u_i(A(\mathbf{w}_t), \mathbf{w}_t) \geq \text{marg}_{A(\mathbf{w}_t), \mathbf{w}_t}(\mathcal{C})$. Hence, each group of agents gets at least its own marginal contribution to the given allocation.

3.2 Sensing and Errors

We conclude the presentation of our (strategic) setting for fair allocation problems by illustrating some subtle issues arising with the process of verification. The starting point of our

discussion is the observation that verifiers can be practically implemented by *sensing* some parameters (in our setting the parameters χ_i , for each agent i , and λ_g , for each allocated good g) that become observable after the allocation is performed and that, in real-world applications, sensing is subject to errors; for instance, because of the limited precision of the measurement instruments. Therefore, it is unrealistic to assume that the verifier is always able to *exactly* discover (i.e., with arbitrary precision) the actual upper bounds in the scenario S_t and the valuation vector \mathbf{w}_t , and it might be problematic to decide whether an observed discrepancy between verified values and declared ones is due to a strategic behavior or to such sensing errors. In fact, sensing troubles arise even in settings where all relevant information is available as public knowledge that can be acquired via sensing the environment, i.e., even by getting rid of any strategic consideration.

As discussed in the Introduction, this issue has been pointed out in the recent work by Feige and Tennenholtz (2011), though from a slightly different perspective. There, it is observed that mechanisms with verification are often designed in way that performs well when agents have accurate information about their private inputs, but might perform arbitrarily bad when agents are uncertain of their private features. Uncertainty might be again the result of “hardware” measurement errors, or due to the limited computational resources employed by agents for identifying the declared valuations. For instance, in our setting, the type of an agent i consists of her/his characterizing property χ_i plus the property λ_g , for each good g (s)he is interested in. According to the perspective of Feige and Tennenholtz (2011), agent i might only have an estimate of this type. For instance, the agent might not have enough computational resources to precisely determine such properties λ_g for all goods g , so that i ’s type actually represent just a subjective perception of them.

In the light of the above observations, it clearly emerges that “punishing” agents might be effective in mathematical studies, but very inappropriate in real life situations in which uncertainty is inherent due to measurements errors or uncertain inputs. Therefore, in addition to the requirements discussed so far, another desirable property is for the mechanism to use no punishment (or to be “forgiving”, in the sense of Feige & Tennenholtz, 2011).

- ▷ *no punishment*: For each type vector $\theta \in \Theta$, for each feasible allocation π for S_θ , and for each agent $i \in \mathcal{A}$, it is the case that $p_i(\pi, \mathbf{w}_\theta) = p_i(\pi, \mathbf{w}_{(t_i, \theta_{-i})})$. That is, discrepancies between the given type (possibly declared) and the true/verified one do not have any impact on the payment to agent i . In other words, we may think of payments being always computed under the presumption of innocence, where incorrect declared values do not mean manipulation attempts by the agents.

Moreover, if we admit that sensing errors (or uncertain inputs) might occur, then it is relevant to quantitatively assess their impact, too. Ideally, we would like to deal with mechanisms that can tolerate such errors, in the sense that “small” errors should determine small deviations from the outcome we would have obtained with no errors at all. Note that this is generally not possible in mechanisms based on punishment approaches where, to enforce truthfulness, the punishment might be disproportional to the harm done by misreporting (cf. Feige & Tennenholtz, 2011). We next formalize our final desideratum.

For any type vector θ , for any set $\mathcal{C} \subseteq \mathcal{A}$ of agents, and for any set $G' \subseteq G$ of goods, let us define $\langle \mathcal{C}, G', \zeta_\theta \rangle$ as the restriction of the scenario $\langle \mathcal{A}, G, \zeta_\theta \rangle$ where only agents in \mathcal{C}

and goods in G' are considered.³ Moreover, given two type vectors θ and $\bar{\theta}$, we denote by $\Theta[\theta, \bar{\theta}]$ the set of all type vectors of the form (X_1, \dots, X_n) , where $X_i \in \{\theta_i, \bar{\theta}_i\}$, for each $i \in \mathcal{A}$. Then, the ‘distance’ $dist_{\mathbf{w}}(\theta, \bar{\theta})$ between θ and $\bar{\theta}$ under the valuation vector \mathbf{w} (or just $dist(\theta, \bar{\theta})$, if \mathbf{w} is understood from the context) is defined by looking at the worst possible impact that type vectors in $\Theta[\theta, \bar{\theta}]$ may have on the optimal solutions computed over all possible restrictions of the given setting:

$$dist_{\mathbf{w}}(\theta, \bar{\theta}) = \max_{\mathcal{C} \subseteq \mathcal{A}, G' \subseteq G, \theta', \theta'' \in \Theta[\theta, \bar{\theta}]} |\text{opt}(\langle \mathcal{C}, G', \zeta_{\theta'} \rangle, \mathbf{w}_{\theta'}) - \text{opt}(\langle \mathcal{C}, G', \zeta_{\theta''} \rangle, \mathbf{w}_{\theta''})|.$$

Now, recall that in a truthful mechanism it is always convenient for the agents to report the true types. Assume however that agents declare a type vector $\bar{\mathbf{t}}$ different from their true type vector \mathbf{t} .⁴ As a consequence, we get a revealed setting $S_{\bar{\mathbf{t}}}$ and a valuation vector $\mathbf{w}_{\bar{\mathbf{t}}} = (\bar{w}_1, \dots, \bar{w}_n)$, while the available verifier \mathbf{v} discloses information about the scenario $S_{\mathbf{t}}$ and the vector $\mathbf{w}_{\mathbf{t}} = (w_1, \dots, w_n)$ (restricted to allocated goods). In standard mechanisms design settings (and in particular under punishment approaches), no guarantee on any property could be given in this case, as mechanisms are designed to be analyzed when reports are truthful. Instead, we would like here that mechanisms are tolerant of sensing errors, as formalized below.

▷ *error tolerance*: There is a constant $c \geq 0$ such that, for each type vector $\bar{\mathbf{t}}$ and for each agent $i \in \mathcal{A}$, $|u_i(A(\mathbf{w}_{\mathbf{t}}), \mathbf{w}_{\mathbf{t}}) - u_i(A(\mathbf{w}_{\bar{\mathbf{t}}}), \mathbf{w}_{\bar{\mathbf{t}}})| \leq c \times dist(\mathbf{t}, \bar{\mathbf{t}})$.

Intuitively, under an error tolerant mechanism, the consequences of errors over good allocation outcomes produce a linear ‘distorting effect’ over agents’ utilities (and, in turn on the various properties of the mechanism). In particular, the above property is stated without any assumption about how sensing errors come into play. Indeed, the notion of $dist(\mathbf{t}, \bar{\mathbf{t}})$ formalizes these errors from a global perspective. For instance, we do not require that errors affect uniformly the valuations of agents, and it might be well the case that errors are biased towards some specific agent.

4. Mechanisms with Verification for Allocation Problems

In this section, we introduce a mechanism with verification for allocation problems and start its analysis, by preliminary proving some properties that hold over optimal allocations.

4.1 General Properties of Allocation Problems

Let $\theta \in \Theta$ be any given type vector, and consider the allocation scenario $S_{\theta} = \langle \mathcal{A}, G, \zeta_{\theta} \rangle$, where $\zeta_{\theta} = (\zeta_1, \dots, \zeta_n)$, together with the valuations given by $\mathbf{w}_{\theta} = (w_1, \dots, w_n)$. We start

-
3. Note the little abuse of notation: the vector ζ in $\langle \mathcal{C}, G', \zeta \rangle$ should be in fact its restriction over \mathcal{C} . However, to keep the notation simple, we just write ζ , as no confusion may arise. Similarly, any valuation vector \mathbf{w} for \mathcal{A} will be transparently considered as a valuation vector for any subset of agents $\mathcal{C} \subseteq \mathcal{A}$ —we just get rid of the unused components associated with agents in $\mathcal{A} \setminus \mathcal{C}$.
 4. The property is discussed from the perspective of uncertain inputs. The adaptation to the case of verification errors (or to the case when both types of errors occur) is an easy task, as it is mainly a matter of different interpretation of concepts.




Figure 3: One-good version of the allocation problem in Example 4.1, with two allocations and their associated update graph, as defined in the proof of Theorem 4.4. In the graphical representation, crossing lines represent the edges of the bipartite cliques connecting the two groups of virtual agents with the goods they are interested in.

by observing that the optimization problem used to allocate goods to agents can be equivalently reformulated in such a way that precisely one good can be allocated to each agent. Intuitively, we may replace each agent i by ζ_i fresh agents with the same valuations as i . We remark that such an equivalence is just used for combinatorial optimization purposes, i.e., without affecting any game theoretic issue.

Let us now formalize the above intuition. First, we denote by S_{θ}^1 the *one-good* version $\langle \mathcal{A}^1, G, \mathbf{1} \rangle$ of the scenario S_{θ} , where:

- \mathcal{A}^1 is the set of agents $\bigcup_{i \in \mathcal{A}} \text{clones}(i)$ such that for each agent $i \in \mathcal{A}$, $\text{clones}(i)$ is a set of ζ_i fresh agents;
- $\mathbf{1}$ is the vector where all components are 1.

Moreover, we denote by \mathbf{w}_{θ}^1 the vector for agents in \mathcal{A}^1 where, for each agent $c \in \mathcal{A}^1$, the component w_c^1 associated with c is such that $w_c^1 = w_i$, with i being the agent in \mathcal{A} for which $c \in \text{clones}(i)$ holds. Thus, in the allocation problem S_{θ}^1 and by considering the vector \mathbf{w}_{θ}^1 , each “clone” $c \in \text{clones}(i)$ gets exactly one good, and has the same valuations as agent i in \mathbf{w}_{θ} .

Example 4.1. Consider the scenario $S_t = \langle \mathcal{A}, \{g_1, \dots, g_8\}, \zeta_t \rangle$, with $\mathcal{A} = \{a_1, a_2\}$, and the vector $\mathbf{w}_t = (w_{a_1}, w_{a_2})$ discussed in Example 2.1. The one-good version is the scenario $S_t^1 = \langle \mathcal{A}^1, \{g_1, \dots, g_8\}, \mathbf{1} \rangle$ shown Figure 3(I). Note that the set of agents in this scenario is $\mathcal{A}^1 = \{(a_1)_1, (a_1)_2, (a_1)_3, (a_2)_1, (a_2)_2, (a_2)_3\}$, where $\text{clones}(a_i) = \{(a_i)_1, (a_i)_2, (a_i)_3\}$, for each $i \in \{1, 2\}$. Indeed, recall that $\zeta_t = (3, 3)$. Finally, the vector \mathbf{w}_t^1 is such that $w_{(a_1)_h}^1 = w_{a_1}$ (resp., $w_{(a_2)_h}^1 = w_{a_2}$), for each $h \in \{1, 2, 3\}$. \triangleleft

Now, consider the scenario $\langle \mathcal{C}, G', \zeta_\theta \rangle$, i.e., the restriction of $\langle \mathcal{A}, G, \zeta_\theta \rangle$ where only agents in \mathcal{C} and goods in G' are considered, and let $\pi_{\mathcal{C}}^1$ be an allocation for its one-good version $\langle \mathcal{C}, G', \zeta_\theta \rangle^1$. Consider the function $\pi_{\mathcal{C}} : \mathcal{C} \rightarrow 2^{G'}$ such that $\pi_{\mathcal{C}}(i) = \bigcup_{c \in \text{clones}(i)} \pi_{\mathcal{C}}^1(c)$, for each $i \in \mathcal{A}$. Note that $|\pi_{\mathcal{C}}(i)| \leq \zeta_i$, for each $i \in \mathcal{A}$. Therefore, $\pi_{\mathcal{C}}$ is a feasible allocation for $\langle \mathcal{C}, G', \zeta_\theta \rangle$, denoted by $\zeta\text{-good}(\pi_{\mathcal{C}}^1)$. Moreover, by construction, $\text{val}(\pi_{\mathcal{C}}, \mathbf{w}_\theta) = \text{val}(\pi_{\mathcal{C}}^1, \mathbf{w}_\theta^1)$. Conversely, any feasible allocation $\bar{\pi}_{\mathcal{C}}$ for $\langle \mathcal{C}, G', \zeta_\theta \rangle$ is associated with the non-empty set $\text{one-good}(\bar{\pi}_{\mathcal{C}})$ of all those allocations $\bar{\pi}_{\mathcal{C}}^1$ for $\langle \mathcal{C}, G', \zeta_\theta \rangle^1$ such that $\bar{\pi}_{\mathcal{C}} = \zeta\text{-good}(\bar{\pi}_{\mathcal{C}}^1)$, also called the one-good forms of $\bar{\pi}_{\mathcal{C}}$. The following is immediate by definition of one-good version and forms.

Fact 4.2. *Let $\theta \in \Theta$ be any given type vector, let S_θ^1 be the one-good version of S_θ , and let $\pi_{\mathcal{C}}^1$ be an allocation for S_θ^1 . Then, $\pi_{\mathcal{C}}^1$ is an optimal allocation for S_θ^1 w.r.t. \mathbf{w}_θ^1 if, and only if, $\zeta\text{-good}(\pi_{\mathcal{C}}^1)$ is an optimal allocation for S_θ w.r.t. \mathbf{w}_θ .*

Example 4.3. Consider again the setting of Example 4.1, and the allocation π such that $\pi(a_1) = \{g_1, g_2, g_3\}$ and $\pi(a_2) = \{g_5, g_7, g_8\}$. Note that π is not an optimal allocation w.r.t. the valuation vector \mathbf{w}_t (reported in Figure 1(I)), because $\text{val}(\pi, \mathbf{w}_t) = 50$ while $\text{opt}(S_t, \mathbf{w}_t) = 51$ (see, e.g., the optimal allocation π^* in Figure 1(II)). Now, notice that the allocation depicted in Figure 3(II) is indeed an associated one-good form allocation, which is actually not an optimal allocation for S_t^1 w.r.t. \mathbf{w}_t^1 , by Fact 4.2. \triangleleft

We are now in the position of stating a property that holds on any optimal allocation π . The property is in fact of interest of its own, i.e., independently of its application to the study of fair allocation problems. In words, it tells us that, whenever we are interested in allocating goods to any subset of agents, we may safely consider only goods in $\text{img}(\pi)$, rather than the whole set G . In our case, it is a basic technical ingredient for showing a number of key properties because, intuitively, it allows us to get rid of alternative (optimal) allocations, possibly based on non-allocated goods in $G \setminus \text{img}(\pi)$.

Theorem 4.4. *Let $\theta \in \Theta$ be any given type vector, let π be an optimal allocation for $\langle \mathcal{A}, G, \zeta_\theta \rangle$ w.r.t. \mathbf{w}_θ , and let $\mathcal{C} \subseteq \mathcal{A}$ be a set of agents. Then, every optimal allocation for $\langle \mathcal{C}, \text{img}(\pi), \zeta_\theta \rangle$ w.r.t. \mathbf{w}_θ is an optimal allocation for $\langle \mathcal{C}, G, \zeta_\theta \rangle$ w.r.t. \mathbf{w}_θ .*

Proof. Let $\mathcal{C} \subseteq \mathcal{A}$ be any set of agents, and assume that $\eta_{\mathcal{C}}$ is an optimal allocation for $\langle \mathcal{C}, \text{img}(\pi), \zeta_\theta \rangle$ w.r.t. \mathbf{w}_θ . We shall show that $\eta_{\mathcal{C}}$ is an optimal allocation for the unrestricted problem $\langle \mathcal{C}, G, \zeta_\theta \rangle$ w.r.t. \mathbf{w}_θ , too.

To this end, consider any optimal allocation $\lambda_{\mathcal{C}}$ for the problem $\langle \mathcal{C}, G, \zeta_\theta \rangle$ where all goods in G are available to the agents in \mathcal{C} . We next prove that $\text{val}(\eta_{\mathcal{C}}, \mathbf{w}_\theta) = \text{val}(\lambda_{\mathcal{C}}, \mathbf{w}_\theta)$. This clearly follows from the optimality of $\eta_{\mathcal{C}}$ if $\text{img}(\lambda_{\mathcal{C}}) \subseteq \text{img}(\pi)$ holds. Therefore, to be strictly better than $\eta_{\mathcal{C}}$, $\lambda_{\mathcal{C}}$ must allocate some good in $G \setminus \text{img}(\pi)$. Assume thus by contradiction that $\text{val}(\eta_{\mathcal{C}}, \mathbf{w}_\theta) < \text{val}(\lambda_{\mathcal{C}}, \mathbf{w}_\theta)$, and hence $\text{img}(\lambda_{\mathcal{C}}) \not\subseteq \text{img}(\pi)$, which entails that $\text{img}(\pi) \subset G$. Consider two allocations $\eta_{\mathcal{C}}^1 \in \text{one-good}(\eta_{\mathcal{C}})$ and $\lambda_{\mathcal{C}}^1 \in \text{one-good}(\lambda_{\mathcal{C}})$, and observe first that: $\text{val}(\eta_{\mathcal{C}}^1, \mathbf{w}_\theta^1) = \text{val}(\eta_{\mathcal{C}}, \mathbf{w}_\theta) < \text{val}(\lambda_{\mathcal{C}}, \mathbf{w}_\theta) = \text{val}(\lambda_{\mathcal{C}}^1, \mathbf{w}_\theta^1)$.

Let L be the set of agents whose good-assignment are the same according to these allocations, i.e., $L = \{c \in \mathcal{C}^1 \mid \lambda_{\mathcal{C}}^1(c) = \eta_{\mathcal{C}}^1(c)\}$. Then, define $\Delta(\eta_{\mathcal{C}}^1, \lambda_{\mathcal{C}}^1) = (\mathcal{C}^1 \setminus L \cup \{s, t\}, E)$ to be the directed graph, called *update* graph for $\eta_{\mathcal{C}}^1$ w.r.t. $\lambda_{\mathcal{C}}^1$, whose nodes are the agents in \mathcal{C}^1 that change their goods in the two allocations plus two distinguished nodes s and t , and whose edges in E are defined as follows:

- There is an edge from agent c to agent c' if $\lambda_c^1(c') = \eta_c^1(c) \neq \emptyset$;
- There is an edge from s to agent c' if there is no agent c such that $\lambda_c^1(c') = \eta_c^1(c) \neq \emptyset$;
- There is an edge from agent c to t if there is no agent c' such that $\lambda_c^1(c') = \eta_c^1(c) \neq \emptyset$;
- No further edges are in E .

For an example construction, consider Figure 3(IV) showing the update graph for the allocation shown in Figure 3(II) w.r.t. the allocation shown in Figure 3(III).

As each agent gets one good in η_c^1 and λ_c^1 , each node in $\Delta(\eta_c^1, \lambda_c^1)$ but s and t has exactly one incoming edge and one outgoing edge. Moreover, by construction, s has no incoming edge, and t has no outgoing edge. Thus, the update graph consists of a number of paths from s to t and a number of cycles, all of them being disjoint from each other.

Let $\{\tau_1, \dots, \tau_h\}$ be the set of all possible paths from s to t or cycles in $\Delta(\eta_c^1, \lambda_c^1)$, and for a path or a cycle $\tau_i = \alpha_1, \dots, \alpha_m$, let $\mathbf{agents}(\tau_i)$ be the set $\{\alpha_1, \dots, \alpha_m\} \setminus \{s, t\}$. In addition, let us fix the following notation: For any function $\pi : X \rightarrow 2^G$ over some given domain X , let $\pi[X']$ denote the restriction of π over the (sub-)domain $X' \subseteq X$. Moreover, given two functions $\pi_1 : X_1 \rightarrow 2^{G_1}$ and $\pi_2 : X_2 \rightarrow 2^{G_2}$ over the two domains X_1 and X_2 , respectively, such that $X_1 \cap X_2 = \emptyset$, let $\pi_1 \uplus \pi_2 : X_1 \cup X_2 \rightarrow 2^{G_1 \cup G_2}$ be the function such that $(\pi_1 \uplus \pi_2)[X_1] = \pi_1$ and $(\pi_1 \uplus \pi_2)[X_2] = \pi_2$.

By the construction of the update graph, note that λ_c^1 can be expressed in terms of the disjoint paths/cycles τ_1, \dots, τ_h by the following expression:

$$\eta_c^1[\mathcal{C}^1 \setminus \bigcup_{i=1}^h \mathbf{agents}(\tau_i)] \uplus \lambda_c^1[\bigcup_{i=1}^h \mathbf{agents}(\tau_i)].$$

Observe now that, because $\mathbf{val}(\eta_c^1, \mathbf{w}^1) < \mathbf{val}(\lambda_c^1, \mathbf{w}^1)$, there must exist a set of agents $\mathbf{agents}(\tau_k)$, associated with some disjoint path/cycle τ_k , with $1 \leq k \leq h$, such that the value of the goods allocated to these agents according to λ_c^1 is greater than the corresponding value for the same agents obtained with η_c^1 . Then, consider the function $\pi_{\tau_k} = \eta_c^1[\mathcal{C}^1 \setminus \mathbf{agents}(\tau_k)] \uplus \lambda_c^1[\mathbf{agents}(\tau_k)]$, and note first that π_{τ_k} is an allocation for $\langle \mathcal{C}, G, \zeta_{\theta} \rangle^1$. In particular, note that, for each agent $c \in \mathcal{C}^1$, $|\pi_{\tau_k}(c)| = 1$ holds, because this constraint actually holds on η_c^1 and λ_c^1 , and because of the definition of the operator ‘ \uplus ’. Moreover, by the choice of τ_k , we also have $\mathbf{val}(\pi_{\tau_k}, \mathbf{w}_{\theta}^1) > \mathbf{val}(\eta_c^1, \mathbf{w}_{\theta}^1)$.

Note that if τ_k were either a cycle or a path of the form $s, \alpha_2, \dots, \alpha_{m-1}, t$ such that $\lambda_c^1(\alpha_2) \subseteq \mathbf{img}(\pi)$, then $\mathbf{img}(\pi_{\tau_k}) \subseteq \mathbf{img}(\pi)$ would hold. Indeed, by definition of the edges of the update graph, only the first node of a path (that is, α_2 in case τ_k is a path) may be such that $\lambda_c^1(\alpha_2) \setminus \mathbf{img}(\pi) \neq \emptyset$. However, as observed above, this is impossible because $\mathbf{val}(\pi_{\tau_k}, \mathbf{w}_{\theta}^1) > \mathbf{val}(\eta_c^1, \mathbf{w}_{\theta}^1)$ would contradict the optimality of η_c^1 , and hence the optimality of η_c , by Fact 4.2. Therefore, we can conclude that τ_k is a path of the form $s, \alpha_2, \dots, \alpha_{m-1}, t$ with $\pi_{\tau_k}(\alpha_2) = \lambda_c^1(\alpha_2) = \{g'\} \in G \setminus \mathbf{img}(\pi)$. That is, the allocation π_{τ_k} (over the agents in \mathcal{C}^1) is such that $\mathbf{img}(\pi_{\tau_k}) = \{g'\} \cup \mathbf{img}(\eta_c^1) \setminus \eta_c^1(\alpha_{m-1})$. In particular, observe that $\eta_c^1(\alpha_{m-1}) \subseteq \mathbf{img}(\pi) \setminus \mathbf{img}(\pi_{\tau_k})$ holds, by definition of the edges of the update graph and since there is an edge from α_{m-1} to t .

Let us now come back to the optimal allocation π for $\langle \mathcal{A}, G, \zeta_{\theta} \rangle$ w.r.t. \mathbf{w}_{θ} , and let π^1 be an (optimal) allocation in $\mathbf{one-good}(\pi)$. Let $A \subseteq \mathcal{C}^1$ be a set of agents with $\alpha_2 \in A$

Input:	A type vector $\boldsymbol{\theta} \in \Theta$, and a feasible allocation π for $S_{\boldsymbol{\theta}}$;
Assumption:	The verifier \mathbf{v} (for \mathbf{t}) is available, with $\mathbf{v}(\pi) = (v_1, \dots, v_n)$;
Notation:	$S_{\boldsymbol{\theta}} = \langle \mathcal{A}, G, \zeta_{\boldsymbol{\theta}} \rangle$, $\mathbf{w}_{\boldsymbol{\theta}} = (w_1, \dots, w_n)$, $\mathbf{w}_{\mathbf{v}(\pi)} = (w_{v_1}, \dots, w_{v_n})$;
<ol style="list-style-type: none"> 1. Let \mathbb{C} denote the set of all possible subsets of \mathcal{A}; 2. For each $i \in \mathcal{A}$ and $\mathcal{C} \in \mathbb{C}$, 3. Compute an optimal allocation $\pi_{\mathcal{C},i}$ for $\langle \mathcal{C}, \text{img}(\pi), \zeta_{(v_i, \boldsymbol{\theta}_{-i})} \rangle$ w.r.t. $\mathbf{w}_{(v_i, \boldsymbol{\theta}_{-i})}$; 4. For each agent $i \in \mathcal{A}$, 5. For each set $\mathcal{C} \in \mathbb{C}$ such that $i \in \mathcal{C}$, 6. Let $\Delta_{\mathcal{C},i}^1(\pi, \boldsymbol{\theta}) := w_{v_i}(\pi_{\mathcal{C},i}) + \sum_{j \in \mathcal{C} \setminus \{i\}} w_j(\pi_{\mathcal{C},i})$; (=val($\pi_{\mathcal{C},i}$, $\mathbf{w}_{(v_i, \boldsymbol{\theta}_{-i})}$)); 7. Let $\Delta_{\mathcal{C},i}^2(\pi, \boldsymbol{\theta}) := \sum_{j \in \mathcal{C} \setminus \{i\}} w_j(\pi_{\mathcal{C} \setminus \{i\}, i})$; (=val($\pi_{\mathcal{C} \setminus \{i\}, i}$, $\mathbf{w}_{(v_i, \boldsymbol{\theta}_{-i})}$)); 8. Let $\xi_i(\pi, \mathbf{w}_{\boldsymbol{\theta}}) := \sum_{\mathcal{C} \in \mathbb{C}} \frac{(\mathcal{A} - \mathcal{C})! (\mathcal{C} - 1)!}{ \mathcal{A} !} (\Delta_{\mathcal{C},i}^1(\pi, \boldsymbol{\theta}) - \Delta_{\mathcal{C},i}^2(\pi, \boldsymbol{\theta}))$; 9. Define $p_i^{\xi}(\pi, \mathbf{w}_{\boldsymbol{\theta}}) := \xi_i(\pi, \mathbf{w}_{\boldsymbol{\theta}}) - w_{v_i}(\pi)$; 	

 Figure 4: Payment rule \mathbf{p}^{ξ} .

such that the set of goods $\bigcup_{c \in A} \pi^1(c)$ allocated to these agents according to π^1 is equal to $\bigcup_{c \in A} \pi_{\tau_k}(c) \setminus \{g'\} \cup G''$, where $G'' \subseteq \text{img}(\pi) \setminus \text{img}(\pi_{\tau_k})$ and $|G''| \leq 1$. Note that a set A having this property in fact exists: just start with $\{\alpha_2\}$ and then add agents from $\text{agents}(\tau_k)$ until some c is found with $\pi^1(c) \subseteq \text{img}(\pi) \setminus \text{img}(\pi_{\tau_k})$.

Consider then $\bar{\pi} = \pi_{\tau_k}[A] \uplus \pi^1[\mathcal{A}^1 \setminus A]$ and note that $\bar{\pi}$ is well-defined, because the construction of the set A guarantees that no good allocated according to $\pi_{\tau_k}[A]$ can be allocated by π^1 to agents in $\mathcal{A}^1 \setminus A$, and vice-versa. Moreover, $\bar{\pi}$ is a feasible allocation for $\langle \mathcal{A}, G, \zeta_{\boldsymbol{\theta}} \rangle^1$. Indeed, for each agent $c \in \mathcal{A}^1$, $|\bar{\pi}(c)| = 1$ holds, because this constraint actually holds on π_{τ_k} and π^1 . Eventually, since π^1 is an optimal allocation for $\langle \mathcal{A}, G, \zeta_{\boldsymbol{\theta}} \rangle^1$ w.r.t. $\mathbf{w}_{\boldsymbol{\theta}}^1$, $\text{val}(\pi^1, \mathbf{w}_{\boldsymbol{\theta}}^1) \geq \text{val}(\bar{\pi}, \mathbf{w}_{\boldsymbol{\theta}}^1)$ holds. Thus, by construction of $\bar{\pi}$, we get $\text{val}(\pi^1[A], \mathbf{w}_{\boldsymbol{\theta}}^1) \geq \text{val}(\bar{\pi}[A], \mathbf{w}_{\boldsymbol{\theta}}^1) = \text{val}(\pi_{\tau_k}[A], \mathbf{w}_{\boldsymbol{\theta}}^1)$.

Finally, let $\bar{\pi}'_{\mathcal{C}} = \pi_{\tau_k}[\mathcal{C}^1 \setminus A] \uplus \pi^1[A]$ and note that $\bar{\pi}'_{\mathcal{C}}$ is a feasible allocation for $\langle \mathcal{C}, G, \zeta_{\boldsymbol{\theta}} \rangle^1$, because of the arguments used above to show that $\bar{\pi}$ is a feasible allocation. Moreover, observe that $\text{val}(\bar{\pi}'_{\mathcal{C}}, \mathbf{w}_{\boldsymbol{\theta}}^1) \geq \text{val}(\pi_{\tau_k}, \mathbf{w}_{\boldsymbol{\theta}}^1) > \text{val}(\eta_{\mathcal{C}}^1, \mathbf{w}_{\boldsymbol{\theta}}^1)$ and $\text{img}(\bar{\pi}'_{\mathcal{C}}) \subseteq \text{img}(\pi)$. For this latter, just recall that $\alpha_2 \in A$ is the only agent in \mathcal{C}^1 such that $\pi_{\tau_k}(\alpha_2) \setminus \text{img}(\pi) \neq \emptyset$. Again, this entails that $\eta_{\mathcal{C}}^1$ is not optimal w.r.t. $\mathbf{w}_{\boldsymbol{\theta}}^1$ and hence, by Fact 4.2, $\eta_{\mathcal{C}}$ is also not optimal for $\langle \mathcal{C}, \text{img}(\pi), \zeta_{\boldsymbol{\theta}} \rangle$ w.r.t. $\mathbf{w}_{\boldsymbol{\theta}}$. Contradiction. \square

The result immediately entails the following two corollaries.

Corollary 4.5. *For each optimal allocation π for $\langle \mathcal{A}, G, \zeta_{\boldsymbol{\theta}} \rangle$ w.r.t. $\mathbf{w}_{\boldsymbol{\theta}}$ and for each set $\mathcal{C} \subseteq \mathcal{A}$ of agents, $\text{opt}(\langle \mathcal{C}, \text{img}(\pi), \zeta_{\boldsymbol{\theta}} \rangle, \mathbf{w}_{\boldsymbol{\theta}}) = \text{opt}(\langle \mathcal{C}, G, \zeta_{\boldsymbol{\theta}} \rangle, \mathbf{w}_{\boldsymbol{\theta}})$.*

Corollary 4.6. *Let π be an optimal allocation for $\langle \mathcal{A}, G, \zeta_{\boldsymbol{\theta}} \rangle$ w.r.t. $\mathbf{w}_{\boldsymbol{\theta}}$, and let π' be any feasible allocation for $\langle \mathcal{A}, G, \zeta_{\boldsymbol{\theta}} \rangle$, hence with $\text{val}(\pi, \mathbf{w}_{\boldsymbol{\theta}}) \geq \text{val}(\pi', \mathbf{w}_{\boldsymbol{\theta}})$. Then, for each set $\mathcal{C} \subseteq \mathcal{A}$ of agents, $\text{opt}(\langle \mathcal{C}, \text{img}(\pi), \zeta_{\boldsymbol{\theta}} \rangle, \mathbf{w}_{\boldsymbol{\theta}}) \geq \text{opt}(\langle \mathcal{C}, \text{img}(\pi'), \zeta_{\boldsymbol{\theta}} \rangle, \mathbf{w}_{\boldsymbol{\theta}})$.*

4.2 The Design of a Truthful Mechanism

Consider the payment rule \mathbf{p}^ξ defined in Figure 4: We are given a type vector $\boldsymbol{\theta} \in \Theta$, and a feasible allocation π for $S_{\boldsymbol{\theta}}$ that selects some goods $\text{img}(\pi) \subseteq G$ for the agents in \mathcal{A} . Moreover, we use the verifier \mathbf{v} (for the vector of true types \mathbf{t}) that, given π , is able to compute the actual scenario $S_{\mathbf{v}(\pi)}$ and the valuation vector $\mathbf{w}_{\mathbf{v}(\pi)} = (w_{v_1}, \dots, w_{v_n})$ over the allocated goods in $\text{img}(\pi)$. Note that, for the sake of presentation, in this section it is convenient to look at the output of the verifier as a list of equivalent types $\mathbf{v}(\pi) = (v_1, \dots, v_n)$ where v_i , for each agent i , is such that i 's upper bound constraint and i 's goods valuation over $\text{img}(\pi)$ are those computed by the verifier \mathbf{v} . Then, as usual, $(v_i, \boldsymbol{\theta}_{-i})$ denotes the type vector obtained from $\boldsymbol{\theta}$ by replacing i 's type θ_i with the verified type v_i . In particular, $\mathbf{w}_{(v_i, \boldsymbol{\theta}_{-i})}$ denotes the valuation vector (defined over $\text{img}(\pi)$) where the function w_{v_i} is used in place of the valuation function declared in θ_i .

In the first three steps, the payment rule associates an optimal allocation $\pi_{\mathcal{C},i}$ for $\langle \mathcal{C}, \text{img}(\pi), \zeta_{(v_i, \boldsymbol{\theta}_{-i})} \rangle$ w.r.t. $\mathbf{w}_{(v_i, \boldsymbol{\theta}_{-i})}$ with each set $\mathcal{C} \in \mathbb{C}$ of agents and each agent $i \in \mathcal{A}$, where \mathbb{C} is the powerset of \mathcal{A} , i.e., the set of all possible subsets of agents. Then, for each agent $i \in \mathcal{A}$, the rule computes the value $\xi_i(\pi, \mathbf{w}_{\boldsymbol{\theta}})$ in step 8, by means of a formula that depends on the valuations associated with the allocations $\pi_{\mathcal{C},i}$ and $\pi_{\mathcal{C} \setminus \{i\},i}$, for each $\mathcal{C} \in \mathbb{C}$. In particular, it defines two terms ($\Delta_{\mathcal{C},i}^1(\pi, \boldsymbol{\theta})$ and $\Delta_{\mathcal{C},i}^2(\pi, \boldsymbol{\theta})$), which evaluate the allocations $\pi_{\mathcal{C},i}$ and $\pi_{\mathcal{C} \setminus \{i\},i}$, respectively, w.r.t. the valuation vector $\mathbf{w}_{(v_i, \boldsymbol{\theta}_{-i})}$. Actually, for the term $\Delta_{\mathcal{C},i}^2(\pi, \boldsymbol{\theta})$, i 's valuation is immaterial as $\pi_{\mathcal{C} \setminus \{i\},i}$ is an allocation over $\mathcal{C} \setminus \{i\}$. Finally, the payment $p_i^\xi(\pi, \mathbf{w})$ is defined in step 9 as the difference between $\xi_i(\pi, \mathbf{w}_{\boldsymbol{\theta}})$ and $w_{v_i}(\pi)$.

Note that the payment rule depends only on the values of the goods in $\text{img}(\pi)$, so that it is verifiable, according to the definition provided in Section 2.3. Moreover note that, as far as paying agent i is concerned, the rule depends only on the values given by w_{v_i} over allocated goods, i.e., by those values returned by the verifier, rather than by w_i . Thus, i 's declaration is immaterial as far as the computation of the payment is concerned, and the next fact easily follows.

Fact 4.7 (no punishment). *For each type vector $\boldsymbol{\theta} \in \Theta$, for each feasible allocation π for $S_{\boldsymbol{\theta}}$, and for each agent $i \in \mathcal{A}$, it is the case that $p_i^\xi(\pi, \mathbf{w}_{\boldsymbol{\theta}}) = p_i^\xi(\pi, \mathbf{w}_{(t_i, \boldsymbol{\theta}_{-i})})$.*

Moreover, note that the idea underlying the definition of \mathbf{p}^ξ is that, after verification is performed, the utility function will precisely coincide with the “bonus” $\xi_i(\pi, \mathbf{w}_{\boldsymbol{\theta}})$, hence sharing the spirit⁵ of the approach by Nisan and Ronen (2001).

Lemma 4.8. *For each type vector $\boldsymbol{\theta} \in \Theta$ and for each feasible allocation π for $S_{\boldsymbol{\theta}}$, it is the case that $u_i(\pi, \mathbf{w}_{\boldsymbol{\theta}}) = \xi_i(\pi, \mathbf{w}_{\boldsymbol{\theta}})$.*

By exploiting this characterization, we can now show the first crucial result on the payment rule \mathbf{p}^ξ , i.e., that the mechanism (A, \mathbf{p}^ξ) is truthful, provided that A is any arbitrary optimal allocation algorithm.

To get a high-level intuition of the proof below observe that $\xi_i(\pi, \mathbf{w}_{\boldsymbol{\theta}})$ depends on two groups of terms, with $\Delta_{\mathcal{C},i}^2(\pi, \boldsymbol{\theta})$ being basically independent on the given agent i . Thus,

5. We say the “spirit”, because the peculiar form of $\xi_i(\pi, \mathbf{w}_{\boldsymbol{\theta}})$ does not formally fit the framework considered by Nisan and Ronen (2001).

the goal of agent i is to maximize the terms of the form $\Delta_{\mathcal{C},i}^1(\pi, \boldsymbol{\theta})$ defined in step 6 as the valuations of optimal allocations computed (in step 3) by considering i 's verified type, by focusing on goods in $\text{img}(\pi)$ (so that verified values coincide with true ones), and by considering subsets of the whole set \mathcal{A} of agents. The salient machinery is then provided by Corollary 4.5, according to which it will be always convenient for agent i to report its true type. Indeed, if π is an optimal allocation computed via A based on the true type of agent i , then Corollary 4.5 guarantees that $\Delta_{\mathcal{C},i}^1(\pi, \boldsymbol{\theta})$ will get the maximum possible value over all possible allocations for scenarios obtained by considering subsets of agents, i.e., independently on the allocation π being actually selected. This can be done without strategically interacting with the other agents. Therefore, by designing the payment rule in a way that depends only on the values returned by the verifier, not only we end up with a verifiable rule using no punishment, but we also obtain a truthful mechanism based on it.

Theorem 4.9 (truthfulness). *Let A be any optimal allocation algorithm. Then, the mechanism with verification (A, \mathbf{p}^ξ) is truthful.*

Proof. We have to show that, for each agent $i \in \mathcal{A}$ and each reported type vector \mathbf{d} , the following holds: $u_i(A(\mathbf{w}_{(t_i, \mathbf{d}_{-i})}), \mathbf{w}_{(t_i, \mathbf{d}_{-i})}) \geq u_i(A(\mathbf{w}_{\mathbf{d}}), \mathbf{w}_{\mathbf{d}})$; hence, by Lemma 4.8, that $\xi_i(A(\mathbf{w}_{(t_i, \mathbf{d}_{-i})}), \mathbf{w}_{(t_i, \mathbf{d}_{-i})}) \geq \xi_i(A(\mathbf{w}_{\mathbf{d}}), \mathbf{w}_{\mathbf{d}})$.

Consider the construction reported in Figure 4 for the two cases of $\boldsymbol{\theta} = \mathbf{d}$ and $\boldsymbol{\theta} = (t_i, \mathbf{d}_{-i})$, and let $\pi = A(\mathbf{w}_{\mathbf{d}})$ and $\pi' = A(\mathbf{w}_{(t_i, \mathbf{d}_{-i})})$ be the corresponding optimal allocations (for the scenarios $S_{\mathbf{d}}$ and $S_{(t_i, \mathbf{d}_{-i})}$, respectively) received as input by the payment rule in Figure 4. For any set $\mathcal{C} \in \mathbb{C}$ of agents, and for any agent $i \in \mathcal{A}$, let $\pi_{\mathcal{C},i}$ (resp., $\pi'_{\mathcal{C},i}$) be the allocation computed at step 3. Note that this step is well defined, because such an optimal allocation always exists. Indeed, just note that there exist feasible allocations for any scenario $S_{\boldsymbol{\theta}}$, e.g., any allocation $\bar{\pi} : \mathcal{C} \rightarrow 2^{\text{img}(\pi)}$ such that $\bar{\pi}(j) = \{g\}$ for some good $g \in \pi(j)$. That is, any allocation that trivially satisfies every upper-bound constraint, since it assigns one good to each agent.

We now show that the following two properties hold, for each $\mathcal{C} \in \mathbb{C}$ and $i \in \mathcal{A}$:

- (A) $\Delta_{\mathcal{C},i}^1(\pi', (t_i, \mathbf{d}_{-i})) \geq \Delta_{\mathcal{C},i}^1(\pi, \mathbf{d})$, and
- (B) $\Delta_{\mathcal{C},i}^2(\pi', (t_i, \mathbf{d}_{-i})) = \Delta_{\mathcal{C},i}^2(\pi, \mathbf{d})$.

In order to prove (A), observe that by step 6, $\Delta_{\mathcal{C},i}^1(\pi', (t_i, \mathbf{d}_{-i})) = \text{val}(\pi_{\mathcal{C},i}, \mathbf{w}_{(v_i, \mathbf{d}_{-i})}) = \text{val}(\pi_{\mathcal{C},i}, \mathbf{w}_{(t_i, \mathbf{d}_{-i})})$, where the last equality holds because $\pi'_{\mathcal{C},i} \subseteq \text{img}(\pi')$, so that the true valuation is disclosed by the verifier. Then, observe that $\pi'_{\mathcal{C},i}$ is an optimal allocation for $\langle \mathcal{C}, \text{img}(\pi'), \zeta_{(v_i, \mathbf{d}_{-i})} \rangle = \langle \mathcal{C}, \text{img}(\pi'), \zeta_{(t_i, \mathbf{d}_{-i})} \rangle$ w.r.t. $\mathbf{w}_{(v_i, \mathbf{d}_{-i})}$ and, hence, w.r.t. $\mathbf{w}_{(t_i, \mathbf{d}_{-i})}$, and $\pi' = A(\mathbf{w}_{(t_i, \mathbf{d}_{-i})})$ is an optimal allocation for $S_{(t_i, \mathbf{d}_{-i})} = \langle \mathcal{A}, G, \zeta_{(t_i, \mathbf{d}_{-i})} \rangle$ w.r.t. $\mathbf{w}_{(t_i, \mathbf{d}_{-i})}$. Thus, by Corollary 4.5, we get the following expression:

$$\Delta_{\mathcal{C},i}^1(\pi', (t_i, \mathbf{d}_{-i})) = \text{val}(\pi'_{\mathcal{C},i}, \mathbf{w}_{(t_i, \mathbf{d}_{-i})}) = \text{opt}(\langle \mathcal{C}, G, \zeta_{(t_i, \mathbf{d}_{-i})} \rangle, \mathbf{w}_{(t_i, \mathbf{d}_{-i})}). \quad (2)$$

Similarly, $\Delta_{\mathcal{C},i}^1(\pi, \mathbf{d}) = \text{val}(\pi_{\mathcal{C},i}, \mathbf{w}_{(v_i, \mathbf{d}_{-i})}) = \text{val}(\pi_{\mathcal{C},i}, \mathbf{w}_{(t_i, \mathbf{d}_{-i})})$ holds. Thus, we can use Equation 2, in order to get $\Delta_{\mathcal{C},i}^1(\pi', (t_i, \mathbf{d}_{-i})) = \text{opt}(\langle \mathcal{C}, G, \zeta_{(t_i, \mathbf{d}_{-i})} \rangle, \mathbf{w}_{(t_i, \mathbf{d}_{-i})}) \geq \text{val}(\pi_{\mathcal{C},i}, \mathbf{w}_{(t_i, \mathbf{d}_{-i})})$. This shows that (A) holds.

Let us now focus on (B). By step 7, we have $\Delta_{\mathcal{C},i}^2(\pi', (t_i, \mathbf{d}_{-i})) = \mathbf{val}(\pi'_{\mathcal{C} \setminus \{i\},i}, \mathbf{w}_{(v_i, \mathbf{d}_{-i})})$ whereas $\Delta_{\mathcal{C},i}^2(\pi, \mathbf{d}) = \mathbf{val}(\pi_{\mathcal{C} \setminus \{i\},i}, \mathbf{w}_{(v_i, \mathbf{d}_{-i})})$. Recall that $\pi_{\mathcal{C} \setminus \{i\},i}$ (resp., $\pi'_{\mathcal{C} \setminus \{i\},i}$) is an optimal allocation for $\langle \mathcal{C} \setminus \{i\}, \mathbf{img}(\pi), \zeta_{(v_i, \mathbf{d}_{-i})} \rangle$ (resp., $\langle \mathcal{C} \setminus \{i\}, \mathbf{img}(\pi'), \zeta_{(v_i, \mathbf{d}_{-i})} \rangle$) w.r.t. $\mathbf{w}_{(v_i, \mathbf{d}_{-i})}$. Then, because of the fact that i 's evaluation is immaterial here, we have that $\pi_{\mathcal{C} \setminus \{i\},i}$ (resp., $\pi'_{\mathcal{C} \setminus \{i\},i}$) is an optimal allocation for $\langle \mathcal{C} \setminus \{i\}, \mathbf{img}(\pi), \zeta_{\mathbf{d}} \rangle$ (resp., $\langle \mathcal{C} \setminus \{i\}, \mathbf{img}(\pi'), \zeta_{(t_i, \mathbf{d}_{-i})} \rangle$) w.r.t. $\mathbf{w}_{\mathbf{d}}$ (resp., $\mathbf{w}_{(t_i, \mathbf{d}_{-i})}$). Then, we recall that π (resp., π') is an optimal allocation for $S_{\mathbf{d}}$ (resp., $S_{(t_i, \mathbf{d}_{-i})}$) w.r.t. $\mathbf{w}_{\mathbf{d}}$ (resp., $\mathbf{w}_{(t_i, \mathbf{d}_{-i})}$). Thus, by Corollary 4.5, $\Delta_{\mathcal{C},i}^2(\pi', (t_i, \mathbf{d}_{-i})) = \mathbf{opt}(\langle \mathcal{C} \setminus \{i\}, G, \zeta_{(t_i, \mathbf{d}_{-i})} \rangle, \mathbf{w}_{(t_i, \mathbf{d}_{-i})})$. Moreover, we get:

$$\Delta_{\mathcal{C},i}^2(\pi, \mathbf{d}) = \mathbf{opt}(\langle \mathcal{C} \setminus \{i\}, G, \zeta_{\mathbf{d}} \rangle, \mathbf{w}_{\mathbf{d}}). \quad (3)$$

Eventually, $\Delta_{\mathcal{C},i}^2(\pi', (t_i, \mathbf{d}_{-i})) = \mathbf{opt}(\langle \mathcal{C} \setminus \{i\}, G, \zeta_{(t_i, \mathbf{d}_{-i})} \rangle, \mathbf{w}_{(t_i, \mathbf{d}_{-i})}) = \mathbf{opt}(\langle \mathcal{C} \setminus \{i\}, G, \zeta_{\mathbf{d}} \rangle, \mathbf{w}_{\mathbf{d}})$ holds, as i 's valuation is immaterial, and we get (B) by Equation 3. \square

4.3 Further Properties of Truthful Strategies

Let us now analyze some relevant properties that hold whenever agents choose the dominant strategy of truthfully reporting their private types. The first property is a useful characterization of agents' utilities.

Theorem 4.10. *For each optimal allocation π for $S_{\mathbf{t}} = \langle \mathcal{A}, G, \zeta_{\mathbf{t}} \rangle$ w.r.t. $\mathbf{w}_{\mathbf{t}}$, and for each agent $i \in \mathcal{A}$, it holds that:*

$$u_i(\pi, \mathbf{w}_{\mathbf{t}}) = \sum_{\mathcal{C} \in \mathbb{C}} \frac{(|\mathcal{A}| - |\mathcal{C}|)! (|\mathcal{C}| - 1)!}{|\mathcal{A}|!} \left(\mathbf{opt}(\langle \mathcal{C}, G, \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}}) - \mathbf{opt}(\langle \mathcal{C} \setminus \{i\}, G, \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}}) \right).$$

Proof. By Lemma 4.8, we know that $u_i(\pi, \mathbf{w}_{\mathbf{t}}) = \xi_i(\pi, \mathbf{w}_{\mathbf{t}})$. Then, for each set $\mathcal{C} \in \mathbb{C}$ of agents, and for each agent $i \in \mathcal{A}$, consider the expressions $\Delta_{\mathcal{C},i}^1(\pi, \mathbf{t})$ and $\Delta_{\mathcal{C},i}^2(\pi, \mathbf{t})$ defined in step 6 and step 7, respectively, of the mechanism in Figure 4. Note that, because of the properties of the verifier \mathbf{v} stated in Definition 2.3 and the fact that the payment rule considers only goods that are allocated via π , $\Delta_{\mathcal{C},i}^1(\pi, \mathbf{t}) = \mathbf{val}(\pi_{\mathcal{C},i}, \mathbf{w}_{(v_i, \mathbf{t}_{-i})}) = \mathbf{val}(\pi_{\mathcal{C},i}, \mathbf{w}_{\mathbf{t}})$ and $\Delta_{\mathcal{C},i}^2(\pi, \mathbf{t}) = \mathbf{val}(\pi_{\mathcal{C},i}, \mathbf{w}_{(v_i, \mathbf{t}_{-i})}) = \mathbf{val}(\pi_{\mathcal{C} \setminus \{i\},i}, \mathbf{w}_{\mathbf{t}})$ hold, where $\pi_{\mathcal{C},i}$ and $\pi_{\mathcal{C} \setminus \{i\},i}$ are optimal allocations for $\langle \mathcal{C}, \mathbf{img}(\pi), \zeta_{\mathbf{t}} \rangle$ w.r.t. $\mathbf{w}_{\mathbf{t}}$ and for $\langle \mathcal{C} \setminus \{i\}, \mathbf{img}(\pi), \zeta_{\mathbf{t}} \rangle$ w.r.t. $\mathbf{w}_{\mathbf{t}}$, respectively. Thus, $\Delta_{\mathcal{C},i}^1(\pi, \mathbf{t}) = \mathbf{opt}(\langle \mathcal{C}, \mathbf{img}(\pi), \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}})$ and $\Delta_{\mathcal{C},i}^2(\pi, \mathbf{t}) = \mathbf{opt}(\langle \mathcal{C} \setminus \{i\}, \mathbf{img}(\pi), \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}})$. It follows that:

$$u_i(\pi, \mathbf{w}_{\mathbf{t}}) = \sum_{\mathcal{C} \in \mathbb{C}} \frac{(|\mathcal{A}| - |\mathcal{C}|)! (|\mathcal{C}| - 1)!}{|\mathcal{A}|!} \left(\mathbf{opt}(\langle \mathcal{C}, \mathbf{img}(\pi), \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}}) - \mathbf{opt}(\langle \mathcal{C} \setminus \{i\}, \mathbf{img}(\pi), \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}}) \right). \quad (4)$$

Recall now by Corollary 4.5 that, for each optimal allocation π for $\langle \mathcal{A}, G, \zeta_{\mathbf{t}} \rangle$ w.r.t. $\mathbf{w}_{\mathbf{t}}$ and for each set $\mathcal{C} \in \mathbb{C}$ of agents, $\mathbf{opt}(\langle \mathcal{C}, \mathbf{img}(\pi), \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}}) = \mathbf{opt}(\langle \mathcal{C}, G, \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}})$. Therefore, $\Delta_{\mathcal{C},i}^1(\pi, \mathbf{t}) = \mathbf{opt}(\langle \mathcal{C}, G, \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}})$ and $\Delta_{\mathcal{C},i}^2(\pi, \mathbf{t}) = \mathbf{opt}(\langle \mathcal{C} \setminus \{i\}, G, \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}})$. By using these equalities, the result follows from Equation 4. \square

As agents' utilities are completely independent of the particular optimal allocation π , every agent gets precisely the same utility in every optimal allocation.

Corollary 4.11. *Let π and π' be two optimal allocations for $S_{\mathbf{t}}$ w.r.t. $\mathbf{w}_{\mathbf{t}}$. Then, $u_i(\pi, \mathbf{w}_{\mathbf{t}}) = u_i(\pi', \mathbf{w}_{\mathbf{t}})$ holds, for each $i \in \mathcal{A}$.*

Example 4.12. Consider the scenario $S = \langle \mathcal{A}, G, \zeta_{\mathbf{t}} \rangle$, with $\mathcal{A} = \{a_1, a_2\}$ and $G = \{g_1, \dots, g_8\}$, the valuation vector $\mathbf{w}_{\mathbf{t}} = (w_{a_1}, w_{a_2})$ discussed in Example 2.1, and the allocation π^* illustrated in Figure 1(I). Then, we have:

$$\begin{aligned} u_{a_1}(\pi^*, \mathbf{w}_{\mathbf{t}}) &= \frac{1}{2}(\text{opt}(\langle \{a_1, a_2\}, G, \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}}) - \text{opt}(\langle \{a_2\}, G, \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}})) + \\ &\quad \frac{1}{2}(\text{opt}(\langle \{a_1\}, G, \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}}) - \text{opt}(\langle \{\}, G, \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}})) + \\ &\quad \frac{1}{2}(\text{opt}(\langle \{a_2\}, G, \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}}) - \text{opt}(\langle \{a_2\}, G, \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}})) = \\ &\quad \frac{1}{2}(51 - 26) + \frac{1}{2}(26 - 0) + \frac{1}{2}(26 - 26) = \frac{51}{2}. \end{aligned}$$

For instance, note that $\text{opt}(\langle \{a_1\}, G, \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}}) = 26$, as we can allocate g_1, g_4 , and g_5 to a_1 , if (s)he is the only agent in the scenario.

Similarly, we get $u_{a_2}(\pi^*, \mathbf{w}_{\mathbf{t}}) = \frac{51}{2}$. That is, the two agents will share precisely half of the total value, based on our payment scheme. In fact, by looking at the allocation π^* in Figure 1(II), one might naïvely suppose that a_2 contributed more than a_1 to the overall value associated with π^* . However, this is only due to the specific allocation considered, and not to the actual values of the goods of the two agents. Indeed, the fairness of the utility values resulting from our payment rule suddenly appears when considering the existence of the alternative allocation $\hat{\pi}^*$ in Figure 1(II), which is symmetric w.r.t. π^* and where it seems that a_1 contributed more than a_2 to the overall value: As a matter of fact, the two agents are completely interchangeable over optimal allocations, and this is correctly reflected by our payment scheme (without the need of actually looking at $\hat{\pi}^*$). In particular, from Corollary 4.11, the agents are indifferent w.r.t. the specific optimal allocation being selected, and hence in this case they equally divide all the available value between themselves. \triangleleft

Further basic properties of the mechanism are pointed out next.

Theorem 4.13 (basic properties). *Let A be any optimal allocation algorithm. Then, the mechanism (A, \mathbf{p}^{ξ}) is efficient and guarantees an equal treatment of the equals. Moreover, if all valuations are non-negative, then (A, \mathbf{p}^{ξ}) is individually rational.*

Proof. It is clear that (A, \mathbf{p}^{ξ}) satisfies efficiency and equal treatment of equals.

Let now $i \in \mathcal{A}$ be an agent and recall from Lemma 4.8 that, for each type vector $\theta \in \Theta$ and allocation π for S_{θ} , $u_i(\pi, \mathbf{w}_{\theta}) = \xi_i(\pi, \mathbf{w}_{\theta})$ holds. Consider the payment rule \mathbf{p}^{ξ} defined in Figure 4, and observe that $\xi_i(\pi, \mathbf{w}_{\theta})$ is defined as a weighted summation, over all coalitions $\mathcal{C} \in \mathcal{C}$, of terms having the form $\Delta_{\mathcal{C},i}^1(\pi, \theta) - \Delta_{\mathcal{C},i}^2(\pi, \theta)$. In particular, all weights are positive and we claim that $\Delta_{\mathcal{C},i}^1(\pi, \theta) - \Delta_{\mathcal{C},i}^2(\pi, \theta) \geq 0$ holds. Indeed, just check that, by definition, $\Delta_{\mathcal{C},i}^1(\pi, \theta) = \text{opt}(\langle \mathcal{C}, \text{img}(\pi), \zeta_{\theta} \rangle, \mathbf{w}_{\theta})$ and $\Delta_{\mathcal{C},i}^2(\pi, \theta) = \text{opt}(\langle \mathcal{C} \setminus \{i\}, \text{img}(\pi), \zeta_{\theta} \rangle, \mathbf{w}_{\theta})$. So, under the hypothesis that all valuations are non-negative, we have that $\text{opt}(\langle \mathcal{C}, \text{img}(\pi), \zeta_{\theta} \rangle, \mathbf{w}_{\theta}) - \text{opt}(\langle \mathcal{C} \setminus \{i\}, \text{img}(\pi), \zeta_{\theta} \rangle, \mathbf{w}_{\theta}) \geq 0$ holds, for each $\mathcal{C} \subseteq \mathcal{A}$ and agent $i \in \mathcal{A}$. Thus, $u_i(\pi, \mathbf{w}_{\theta}) \geq 0$, hence the fact that (A, \mathbf{p}^{ξ}) is individually rational trivially follows (even independently of the allocation π , and of whether agent i is truth-telling). \square

Moreover, the mechanism is also tolerant of sensing errors (or uncertain inputs), in the sense of Section 3.2.

Theorem 4.14 (error tolerance). *Let A be any optimal allocation algorithm. Then, the mechanism (A, \mathbf{p}^ξ) is such that for each type vector $\bar{\mathbf{t}}$ and each agent $i \in \mathcal{A}$, $|u_i(A(\mathbf{w}_{\mathbf{t}}), \mathbf{w}_{\mathbf{t}}) - u_i(A(\mathbf{w}_{\bar{\mathbf{t}}}), \mathbf{w}_{\bar{\mathbf{t}}})| \leq 3 \times \text{dist}(\mathbf{t}, \bar{\mathbf{t}})$.*

Proof. Consider the construction reported in Figure 4 for the two cases of $\theta = \mathbf{t}$ and $\theta = \bar{\mathbf{t}}$, and let $\pi = A(\mathbf{w}_{\mathbf{t}})$ and $\bar{\pi} = A(\mathbf{w}_{\bar{\mathbf{t}}})$ be the corresponding allocations received as its input by the payment rule (optimal w.r.t. $\mathbf{w}_{\mathbf{t}}$ and $\mathbf{w}_{\bar{\mathbf{t}}}$, respectively). Moreover, for any set $\mathcal{C} \in \mathbb{C}$ of agents and for any agent $i \in \mathcal{A}$, let $\pi_{\mathcal{C},i}$ and $\bar{\pi}_{\mathcal{C},i}$ be the corresponding allocations computed at step 3. Then, we get:

- $\Delta_{\mathcal{C},i}^1(\pi, \mathbf{t}) = \text{opt}(\langle \mathcal{C}, G, \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}})$, by Equation 2 in the proof of Theorem 4.9;
- $\Delta_{\mathcal{C},i}^1(\bar{\pi}, \bar{\mathbf{t}}) = \text{opt}(\langle \mathcal{C}, \text{img}(\bar{\pi}), \zeta_{(t_i, \bar{t}_{-i})} \rangle, \mathbf{w}_{(t_i, \bar{t}_{-i})})$, by step 3 in Figure 4 and the fact the true valuation of agent i is disclosed by the verifier because $\bar{\pi}_{\mathcal{C},i} \subseteq \text{img}(\bar{\pi})$.
- $\Delta_{\mathcal{C},i}^2(\pi, \mathbf{t}) = \text{opt}(\langle \mathcal{C} \setminus \{i\}, G, \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}})$, by Equation 3 in the proof of Theorem 4.9;
- $\Delta_{\mathcal{C},i}^2(\bar{\pi}, \bar{\mathbf{t}}) = \text{opt}(\langle \mathcal{C} \setminus \{i\}, G, \zeta_{\bar{\mathbf{t}}} \rangle, \mathbf{w}_{\bar{\mathbf{t}}})$, again by Equation 3;

Note now that $|\Delta_{\mathcal{C},i}^1(\bar{\pi}, \bar{\mathbf{t}}) - \text{opt}(\langle \mathcal{C}, \text{img}(\bar{\pi}), \zeta_{\bar{\mathbf{t}}} \rangle, \mathbf{w}_{\bar{\mathbf{t}}})| \leq \text{dist}(\mathbf{t}, \bar{\mathbf{t}})$. Moreover, since $\bar{\pi}$ is an optimal allocation w.r.t. $\mathbf{w}_{\bar{\mathbf{t}}}$, by Corollary 4.5, $\text{opt}(\langle \mathcal{C}, \text{img}(\bar{\pi}), \zeta_{\bar{\mathbf{t}}} \rangle, \mathbf{w}_{\bar{\mathbf{t}}}) = \text{opt}(\langle \mathcal{C}, G, \zeta_{\bar{\mathbf{t}}} \rangle, \mathbf{w}_{\bar{\mathbf{t}}})$ and, hence, $|\Delta_{\mathcal{C},i}^1(\bar{\pi}, \bar{\mathbf{t}}) - \text{opt}(\langle \mathcal{C}, G, \zeta_{\bar{\mathbf{t}}} \rangle, \mathbf{w}_{\bar{\mathbf{t}}})| \leq \text{dist}(\mathbf{t}, \bar{\mathbf{t}})$ holds. Then, observe that $|\text{opt}(\langle \mathcal{C}, G, \zeta_{\bar{\mathbf{t}}} \rangle, \mathbf{w}_{\bar{\mathbf{t}}}) - \Delta_{\mathcal{C},i}^1(\pi, \mathbf{t})| \leq \text{dist}(\mathbf{t}, \bar{\mathbf{t}})$. Therefore, we conclude that $|\Delta_{\mathcal{C},i}^1(\bar{\pi}, \bar{\mathbf{t}}) - \Delta_{\mathcal{C},i}^1(\pi, \mathbf{t})| \leq 2 \times \text{dist}(\mathbf{t}, \bar{\mathbf{t}})$.

Similarly, $|\Delta_{\mathcal{C},i}^2(\bar{\pi}, \bar{\mathbf{t}}) - \Delta_{\mathcal{C},i}^2(\pi, \mathbf{t})| \leq \text{dist}(\mathbf{t}, \bar{\mathbf{t}})$ holds. Hence, by putting it all together, we derive $|(\Delta_{\mathcal{C},i}^1(\bar{\pi}, \bar{\mathbf{t}}) - \Delta_{\mathcal{C},i}^2(\bar{\pi}, \bar{\mathbf{t}})) - (\Delta_{\mathcal{C},i}^1(\pi, \mathbf{t}) - \Delta_{\mathcal{C},i}^2(\pi, \mathbf{t}))| \leq 3 \times \text{dist}(\mathbf{t}, \bar{\mathbf{t}})$.

In the light of the above expression and by step 7 in Figure 4, we have that

$$|\xi_i(\bar{\pi}, \mathbf{w}_{\bar{\mathbf{t}}}) - \xi_i(\pi, \mathbf{w}_{\mathbf{t}})| \leq \sum_{\mathcal{C} \in \mathbb{C}} \frac{(|\mathcal{A}| - |\mathcal{C}|)! (|\mathcal{C}| - 1)!}{|\mathcal{A}|!} 3 \times \text{dist}(\mathbf{t}, \bar{\mathbf{t}}).$$

To conclude the proof, observe that $\sum_{\mathcal{C} \in \mathbb{C}} \frac{(|\mathcal{A}| - |\mathcal{C}|)! (|\mathcal{C}| - 1)!}{|\mathcal{A}|!} = 1$ holds, so that $|\xi_i(\bar{\pi}, \mathbf{w}_{\bar{\mathbf{t}}}) - \xi_i(\pi, \mathbf{w}_{\mathbf{t}})| \leq 3 \times \text{dist}(\mathbf{t}, \bar{\mathbf{t}})$. Therefore, by Lemma 4.8, we get $|u_i(A(\mathbf{w}_{\bar{\mathbf{t}}}), \mathbf{w}_{\bar{\mathbf{t}}}) - u_i(A(\mathbf{w}_{\mathbf{t}}), \mathbf{w}_{\mathbf{t}})| = |\xi_i(\bar{\pi}, \mathbf{w}_{\bar{\mathbf{t}}}) - \xi_i(\pi, \mathbf{w}_{\mathbf{t}})| \leq 3 \times \text{dist}(\mathbf{t}, \bar{\mathbf{t}})$. \square

In fact, we next show that (A, \mathbf{p}^ξ) , with A being an optimal allocation algorithm, satisfies all the remaining properties discussed in Section 3.1. To this end, we first discuss an interpretation of \mathbf{p}^ξ in the context of coalitional games.

5. A Coalitional Game Theory Viewpoint

A *coalitional game* can be modeled as a pair $\mathcal{G} = \langle N, \varphi \rangle$, where $N = \{1, \dots, n\}$ is a finite set of agents, and φ is a function associating with each *coalition* $R \subseteq N$ a real-value $\varphi(R) \in \mathbb{R}$, with $\varphi(\{\}) = 0$, which is meant to encode the worth that agents in \mathcal{C} obtain by collaborating with each other. The function φ is *supermodular* (resp., *submodular*) if $\varphi(R \cup T) + \varphi(R \cap T) \geq \varphi(R) + \varphi(T)$ (resp., $\varphi(R \cup T) + \varphi(R \cap T) \leq \varphi(R) + \varphi(T)$) holds, for each pair of coalitions $R, T \subseteq N$.

A fundamental problem for coalitional games is to single out the most desirable outcomes, usually called *solution concepts*, in terms of appropriate notions of worth distributions, i.e., of vectors of payoffs $\mathbf{x} = (x_1, \dots, x_n) \in \mathbb{R}^n$ such that $\sum_{i \in N} x_i = \varphi(N)$. This question was studied in economics and game theory with the aim of providing arguments and counterarguments about why such proposals are reasonable mathematical renderings of the intuitive concepts of fairness and stability. For further background on coalitional games, the reader is referred to, e.g., the work of Osborne and Rubinstein (1994).

Here, we consider the *Shapley value* of $\mathcal{G} = \langle N, \varphi \rangle$, which is a well-known solution concept such that:

$$\phi_i(\mathcal{G}) = \sum_{R \subseteq N} \frac{(|N| - |R|)! (|R| - 1)!}{|N|!} (\varphi(R) - \varphi(R \setminus \{i\})), \text{ for each } i \in N.$$

Indeed, we shall show that the mechanism defined in Section 4 has a nice interpretation in terms of the Shapley value of some suitable-defined coalitional games. The correspondence will be exploited to prove further properties of our mechanism, at the equilibrium \mathbf{t} where agents truthfully report their types.

5.1 The Shapley Value of Allocation Games

We consider two coalitional games defined on top of an allocation problem.

Definition 5.1. Given the valuation vector \mathbf{w} , we define $\mathcal{G}_{\mathbf{w}}^{marg} = \langle \mathcal{A}, \text{marg}_{\mathbf{w}} \rangle$ and $\mathcal{G}_{\mathbf{w}}^{best} = \langle \mathcal{A}, \text{best}_{\mathbf{w}} \rangle$ as the coalitional games such, that for each set $\mathcal{C} \subseteq \mathcal{A}$ of agents,

- $\text{marg}_{\mathbf{w}}(\mathcal{C}) = \text{opt}(\langle \mathcal{A}, G, \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}}) - \text{opt}(\langle \mathcal{A} \setminus \mathcal{C}, G, \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}})$; and,
- $\text{best}_{\mathbf{w}}(\mathcal{C}) = \text{opt}(\langle \mathcal{C}, G, \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}})$. □

Recall that in Section 3.1 we have defined the concept of marginal contribution $\text{marg}_{\pi, \mathbf{w}}(\mathcal{C})$ of a coalition \mathcal{C} with respect to a given allocation π (because it is sensible to the set of goods allocated by π). In the above definition, with a slight abuse of notation, we have defined a similar concept $\text{marg}_{\mathbf{w}}(\mathcal{C})$, which does not depend on any goods allocation. It turns out that these two concepts actually coincide over optimal allocations.

Theorem 5.2. *Let π be any optimal allocation for $S_{\mathbf{t}} = \langle \mathcal{A}, G, \zeta_{\mathbf{t}} \rangle$ w.r.t. $\mathbf{w}_{\mathbf{t}}$, and let $\mathcal{C} \subseteq \mathcal{A}$ be an arbitrary set of agents. Then, $\text{marg}_{\pi, \mathbf{w}}(\mathcal{C}) = \text{marg}_{\mathbf{w}}(\mathcal{C})$.*

Proof. Let π be an optimal allocation for $S_{\mathbf{t}} = \langle \mathcal{A}, G, \zeta_{\mathbf{t}} \rangle$ w.r.t. $\mathbf{w}_{\mathbf{t}}$, hence optimal for $\langle \mathcal{A}, \text{img}(\pi), \zeta_{\mathbf{t}} \rangle$ w.r.t. $\mathbf{w}_{\mathbf{t}}$. That is, $\text{opt}(S_{\mathbf{t}}, \mathbf{w}_{\mathbf{t}}) = \text{opt}(\langle \mathcal{A}, \text{img}(\pi), \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}})$. Moreover, for each set $\mathcal{C} \subseteq \mathcal{A}$ of agents, by Corollary 4.5, $\text{opt}(\langle \mathcal{A} \setminus \mathcal{C}, \text{img}(\pi), \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}}) = \text{opt}(\langle \mathcal{A} \setminus \mathcal{C}, G, \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}})$ holds. Therefore, $\text{marg}_{\mathbf{w}}(\mathcal{C}) = \text{opt}(S_{\mathbf{t}}, \mathbf{w}_{\mathbf{t}}) - \text{opt}(\langle \mathcal{A} \setminus \mathcal{C}, G, \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}}) = \text{opt}(\langle \mathcal{A}, \text{img}(\pi), \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}}) - \text{opt}(\langle \mathcal{A} \setminus \mathcal{C}, \text{img}(\pi), \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}})$, and the result follows as this value is equivalent to the definition of $\text{marg}_{\pi, \mathbf{w}}(\mathcal{C})$ in Equation 1 (on page 416). □

Note also that $\text{best}_{\mathbf{w}}(\mathcal{C})$ is the *best contribution* of \mathcal{C} , computed assuming that agents in \mathcal{C} were the only agents in the allocation problem. In particular, the game $\mathcal{G}_{\mathbf{w}}^{best}$ has already been considered by Moulin (1992), precisely in the setting of fair division for allocation problems. There, it is shown that the cost function associated with $\mathcal{G}_{\mathbf{w}}^{best}$ is *submodular*.

Proposition 5.3. *The function $\text{best}_{\mathbf{w}}$ is submodular.*

Since $\text{opt}(\langle \mathcal{C}, G, \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}}) = \text{best}_{\mathbf{w}}(\mathcal{C})$, it turns out that $\mathcal{G}_{\mathbf{w}}^{\text{marg}}$ is what is called in the literature the *dual game* of $\mathcal{G}_{\mathbf{w}}^{\text{best}}$, and the following result is known to hold. Nevertheless, we give a direct proof, for completeness.

Corollary 5.4. *The function $\text{marg}_{\mathbf{w}}$ is supermodular.*

Proof. Let $S_{\mathbf{t}} = \langle \mathcal{A}, G, \zeta_{\mathbf{t}} \rangle$ be the given scenario. The result just follows by noticing that $\text{marg}_{\mathbf{w}}(\mathcal{C}) = \text{opt}(\langle \mathcal{A}, G, \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}}) - \text{opt}(\langle \mathcal{A} \setminus \mathcal{C}, G, \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}}) = \text{opt}(\langle \mathcal{A}, G, \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}}) - \text{best}_{\mathbf{w}}(\mathcal{A} \setminus \mathcal{C})$, for each set of agents $\mathcal{C} \subseteq \mathcal{A}$. Therefore, $\text{best}_{\mathbf{w}}(\mathcal{C}) = \text{opt}(\langle \mathcal{A}, G, \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}}) - \text{marg}_{\mathbf{w}}(\mathcal{A} \setminus \mathcal{C})$. Thus, if $\text{best}_{\mathbf{w}}(R \cup T) + \text{best}_{\mathbf{w}}(R \cap T) \leq \text{best}_{\mathbf{w}}(R) + \text{best}_{\mathbf{w}}(T)$ holds $\forall R, T \subseteq \mathcal{A}$, then we have that $\text{marg}_{\mathbf{w}}(\mathcal{A} \setminus (R \cup T)) + \text{marg}_{\mathbf{w}}(\mathcal{A} \setminus (R \cap T)) \geq \text{marg}_{\mathbf{w}}(\mathcal{A} \setminus R) + \text{marg}_{\mathbf{w}}(\mathcal{A} \setminus T)$ holds as well, $\forall R, T \subseteq \mathcal{A}$. Eventually, by letting $R' = \mathcal{A} \setminus R$ and $T' = \mathcal{A} \setminus T$, we get $\text{marg}_{\mathbf{w}}(R' \cap T') + \text{marg}_{\mathbf{w}}(R' \cup T') \geq \text{marg}_{\mathbf{w}}(R') + \text{marg}_{\mathbf{w}}(T')$, for each $\forall R', T' \subseteq \mathcal{A}$. That is, $\text{marg}_{\mathbf{w}}$ is supermodular. \square

The second relevant property is that the payment rule in Section 4 coincides with the Shapley value of the game $\mathcal{G}_{\mathbf{w}}^{\text{best}}$ associated with \mathbf{w} . The result follows by comparing the utility function as in Theorem 4.10 with the expression for the Shapley value of the coalitional game $\mathcal{G}_{\mathbf{w}}^{\text{best}}$. Moreover, we show that the same result can be established for the “dual” game $\mathcal{G}_{\mathbf{w}}^{\text{marg}}$, so that the Shapley values of the two games are identical—for similar correspondences between Shapley values of different games, see also the works by Maniquet (2003) and Kalai and Samet (1983).

Theorem 5.5. *For each optimal allocation π for $S_{\mathbf{t}}$ w.r.t. $\mathbf{w}_{\mathbf{t}}$, and for each agent $i \in \mathcal{A}$, it holds that $u_i(\pi, \mathbf{w}_{\mathbf{t}}) = \xi_i(\pi, \mathbf{w}_{\mathbf{t}}) = \phi_i(\mathcal{G}_{\mathbf{w}}^{\text{best}}) = \phi_i(\mathcal{G}_{\mathbf{w}}^{\text{marg}})$.*

Proof. By comparing the utility function as in Theorem 4.10 with the expression for the Shapley value of the coalitional game $\mathcal{G}_{\mathbf{w}}^{\text{best}}$ associating with each coalition \mathcal{C} of agents the worth $\text{opt}(\langle \mathcal{C}, G, \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}})$, we immediately get that, for each optimal allocation π for $S_{\mathbf{t}}$ w.r.t. $\mathbf{w}_{\mathbf{t}}$, and for each agent $i \in \mathcal{A}$, it holds that $u_i(\pi, \mathbf{w}_{\mathbf{t}}) = \xi_i(\pi, \mathbf{w}_{\mathbf{t}}) = \phi_i(\mathcal{G}_{\mathbf{w}}^{\text{best}})$.

In order to conclude the proof, we show that for each agent $i \in \mathcal{A}$, $\phi_i(\mathcal{G}_{\mathbf{w}}^{\text{marg}}) = \phi_i(\mathcal{G}_{\mathbf{w}}^{\text{best}})$ holds. To this end, first note that these Shapley values can be written as follows:

$$\begin{aligned} - \phi_i(\mathcal{G}_{\mathbf{w}}^{\text{marg}}) &= \sum_{\mathcal{C} \subseteq \mathcal{A}, i \in \mathcal{C}} \frac{(|\mathcal{A}| - |\mathcal{C}|)! (|\mathcal{C}| - 1)!}{|\mathcal{A}|!} T'_{\mathcal{C}}, \text{ and} \\ - \phi_i(\mathcal{G}_{\mathbf{w}}^{\text{best}}) &= \sum_{\mathcal{C} \subseteq \mathcal{A}, i \in \mathcal{C}} \frac{(|\mathcal{A}| - |\mathcal{C}|)! (|\mathcal{C}| - 1)!}{|\mathcal{A}|!} T_{\mathcal{C}}, \end{aligned}$$

where $T'_{\mathcal{C}} = \text{marg}_{\mathbf{w}}(\mathcal{C}) - \text{marg}_{\mathbf{w}}(\mathcal{C} \setminus \{i\})$ and $T_{\mathcal{C}} = \text{opt}(\langle \mathcal{C}, G, \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}}) - \text{opt}(\langle \mathcal{C} \setminus \{i\}, G, \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}})$.

Then, we claim that:

- (1) for each set $\mathcal{C} \subseteq \mathcal{A}$ of agents with $i \in \mathcal{C}$, the set $\bar{\mathcal{C}} = (\mathcal{A} \setminus \mathcal{C}) \cup \{i\}$ is such that $T'_{\mathcal{C}} = T_{\bar{\mathcal{C}}}$, and
- (2) for each set $\bar{\mathcal{C}} \subseteq \mathcal{A}$ of agents with $i \in \bar{\mathcal{C}}$, the set $\mathcal{C} = (\mathcal{A} \setminus \bar{\mathcal{C}}) \cup \{i\}$ is such that $T'_{\mathcal{C}} = T_{\bar{\mathcal{C}}}$.

(1) Let $\mathcal{C} \subseteq \mathcal{A}$ such that $i \in \mathcal{C}$, and observe that $T'_\mathcal{C} = \text{marg}_\mathbf{w}(\mathcal{C}) - \text{marg}_\mathbf{w}(\mathcal{C} \setminus \{i\}) = (\text{opt}(S_\mathbf{t}, \mathbf{w}_\mathbf{t}) - \text{opt}(\langle \mathcal{A} \setminus \mathcal{C}, G, \zeta_\mathbf{t} \rangle, \mathbf{w}_\mathbf{t})) - (\text{opt}(S_\mathbf{t}, \mathbf{w}_\mathbf{t}) - \text{opt}(\langle \mathcal{A} \setminus (\mathcal{C} \setminus \{i\}), G, \zeta_\mathbf{t} \rangle, \mathbf{w}_\mathbf{t})) = \text{opt}(\langle \mathcal{A} \setminus (\mathcal{C} \setminus \{i\}), G, \zeta_\mathbf{t} \rangle, \mathbf{w}_\mathbf{t}) - \text{opt}(\langle \mathcal{A} \setminus \mathcal{C}, G, \zeta_\mathbf{t} \rangle, \mathbf{w}_\mathbf{t}) = \text{opt}(\langle (\mathcal{A} \setminus \mathcal{C}) \cup \{i\}, G, \zeta_\mathbf{t} \rangle, \mathbf{w}_\mathbf{t}) - \text{opt}(\langle \mathcal{A} \setminus \mathcal{C}, G, \zeta_\mathbf{t} \rangle, \mathbf{w}_\mathbf{t})$. Thus, let $\bar{\mathcal{C}} = (\mathcal{A} \setminus \mathcal{C}) \cup \{i\}$, and note that $T'_\mathcal{C} = T_{\bar{\mathcal{C}}}$.

(2) Let $\bar{\mathcal{C}} \subseteq \mathcal{A}$ such that $i \in \bar{\mathcal{C}}$, and observe that $T_{\bar{\mathcal{C}}} = \text{opt}(\langle \bar{\mathcal{C}}, G, \zeta_\mathbf{t} \rangle, \mathbf{w}_\mathbf{t}) - \text{opt}(\langle \bar{\mathcal{C}} \setminus \{i\}, G, \zeta_\mathbf{t} \rangle, \mathbf{w}_\mathbf{t}) = (\text{opt}(S_\mathbf{t}, \mathbf{w}_\mathbf{t}) - \text{opt}(\langle \bar{\mathcal{C}} \setminus \{i\}, G, \zeta_\mathbf{t} \rangle, \mathbf{w}_\mathbf{t})) - (\text{opt}(S_\mathbf{t}, \mathbf{w}_\mathbf{t}) - \text{opt}(\langle \bar{\mathcal{C}}, G, \zeta_\mathbf{t} \rangle, \mathbf{w}_\mathbf{t})) = \text{marg}_\mathbf{w}((\mathcal{A} \setminus \bar{\mathcal{C}}) \cup \{i\}) - \text{marg}_\mathbf{w}(\mathcal{A} \setminus \bar{\mathcal{C}})$. Thus, let $\mathcal{C} = (\mathcal{A} \setminus \bar{\mathcal{C}}) \cup \{i\}$ and note that $T'_\mathcal{C} = T_{\bar{\mathcal{C}}}$.

As (1) and (2) hold, and given the two expressions for $\phi_i(\mathcal{G}_\mathbf{w}^{\text{marg}})$ and $\phi_i(\mathcal{G}_\mathbf{w}^{\text{best}})$, we conclude that the two values coincide. \square

5.2 Marginality, Budget-Balancedness, and Individual Optimality

Now that we have established a precise correspondence between our mechanism and the Shapley value of its associated allocation games, we can show further desirable properties of \mathbf{p}^ξ . In fact, we exploit the following well-known properties (see, e.g., Osborne & Rubinstein, 1994; Young, 1985) of the Shapley value of any game $\mathcal{G} = \langle N, \varphi \rangle$:

- (I) $\sum_{i \in N} \phi_i(\mathcal{G}) = \varphi(N)$;
- (II) If φ is *supermodular* (resp., *submodular*), then $\sum_{i \in R} \phi_i(\mathcal{G}) \geq \varphi(R)$ (resp., $\sum_{i \in R} \phi_i(\mathcal{G}) \leq \varphi(R)$), for each coalition $R \subseteq N$.
- (III) If $\mathcal{G}' = \langle N, \varphi' \rangle$ is a game such that $\varphi'(R) \geq \varphi(R)$, for each $R \subseteq N$, then $\phi_i(\mathcal{G}') \geq \phi_i(\mathcal{G})$, for each agent $i \in N$.

In particular, Property (II) entails that the Shapley value of the supermodular game $\mathcal{G}_\mathbf{w}^{\text{marg}}$ is such that $\sum_{i \in R} \phi_i(\mathcal{G}_\mathbf{w}^{\text{marg}}) \geq \text{marg}_\mathbf{w}(R)$. This means that the Shapley value of this game belongs to another important solution concept for coalitional games called *core* (Osborne & Rubinstein, 1994), which is highly desirable for its “stability”, as every coalition of agents in $\mathcal{G}_\mathbf{w}^{\text{marg}}$ gets at least the worth it deserves according to the function $\text{marg}_\mathbf{w}$.

In our context, this will easily entail that the utility of every group (coalition) of agents \mathcal{C} (i.e., the sum of the utilities of agents in \mathcal{C}), is not less than its actual marginal contribution to the best possible allocations. Note that considering such a collective utility is particularly useful whenever we reason in terms of fairness for groups of agents, rather than just singletons. We first pinpoint that $\text{best}_\mathbf{w}(\mathcal{C})$ and $\text{marg}_\mathbf{w}(\mathcal{C})$ provide an upper bound and a lower bound, respectively, to the utility of \mathcal{C} .

Theorem 5.6. *Let π be an optimal allocation for $S_\mathbf{t}$ w.r.t. $\mathbf{w}_\mathbf{t}$. Then, for each set $\mathcal{C} \subseteq \mathcal{A}$ of agents, $\text{best}_\mathbf{w}(\mathcal{C}) \geq \sum_{i \in \mathcal{C}} u_i(\pi, \mathbf{w}_\mathbf{t}) \geq \text{marg}_\mathbf{w}(\mathcal{C})$.*

Proof. By Theorem 5.5, we know that $u_i(\pi, \mathbf{w}_\mathbf{t}) = \phi_i(\mathcal{G}_\mathbf{w}^{\text{best}}) = \phi_i(\mathcal{G}_\mathbf{w}^{\text{marg}})$, for each agent $i \in \mathcal{A}$ and optimal allocation π . Then, we can simply recall that the function $\text{marg}_\mathbf{w}$ (resp., $\text{best}_\mathbf{w}$) associated with the game $\mathcal{G}_\mathbf{w}^{\text{marg}}$ (resp., $\mathcal{G}_\mathbf{w}^{\text{best}}$) is supermodular (resp., submodular) by Corollary 5.4 (resp., Proposition 5.3). Hence, the result follows as $\sum_{i \in \mathcal{C}} u_i(\pi, \mathbf{w}_\mathbf{t}) = \sum_{i \in \mathcal{C}} \phi_i(\mathcal{G}_\mathbf{w}^{\text{best}}) = \sum_{i \in \mathcal{C}} \phi_i(\mathcal{G}_\mathbf{w}^{\text{marg}})$ and by property (II). \square

By combining the above result and Theorem 5.2, we immediately get the following.

Corollary 5.7 (marginality). *Let A be any optimal allocation algorithm. Then, the mechanism (A, \mathbf{p}^ξ) is such that $\sum_{i \in \mathcal{C}} u_i(A(\mathbf{w}_t), \mathbf{w}_t) \geq \text{marg}_{A(\mathbf{w}_t), \mathbf{w}}(\mathcal{C})$, for each set $\mathcal{C} \subseteq \mathcal{A}$.*

Hence, under the payment rule \mathbf{p}^ξ and optimal allocation algorithms, we get mechanisms that accurately take care of the marginality property defined in Section 3.1. Our next result pertains the budget-balance property of the mechanisms. Again, the correspondence with the Shapley value is crucial to establish the result.

Theorem 5.8. *Let π be an optimal allocation for S_t w.r.t. \mathbf{w}_t . Then, it holds that $\sum_{i \in \mathcal{A}} p_i^\xi(\pi, \mathbf{w}_t) = 0$.*

Proof. By Theorem 5.5, we know that $u_i(\pi, \mathbf{w}_t) = \phi_i(\mathcal{G}_{\mathbf{w}}^{\text{best}})$, for each agent $i \in \mathcal{A}$ and optimal allocation π , where $\phi_i(\mathcal{G}_{\mathbf{w}}^{\text{best}})$ is the Shapley value of $\mathcal{G}_{\mathbf{w}}^{\text{best}}$. By property (I) of the Shapley value, we know that $\sum_{i \in \mathcal{A}} \phi_i(\mathcal{G}_{\mathbf{w}}^{\text{best}}) = \text{best}_{\mathbf{w}}(\mathcal{A})$. Thus, $\sum_{i \in \mathcal{A}} u_i(\pi, \mathbf{w}_t) = \sum_{i \in \mathcal{A}} \phi_i(\mathcal{G}_{\mathbf{w}}^{\text{best}}) = \text{opt}(\langle \mathcal{A}, G, \zeta_t \rangle, \mathbf{w}_t)$. By definition of the utility and letting $\mathbf{w}_t = (w_1, \dots, w_n)$, it then follows that $\text{opt}(\langle \mathcal{A}, G, \zeta_t \rangle, \mathbf{w}_t) = \sum_{i \in \mathcal{A}} w_i(\pi) - \sum_{i \in \mathcal{A}} p_i^\xi(\pi, \mathbf{w})$. Hence, $\sum_{i \in \mathcal{A}} p_i^\xi(\pi, \mathbf{w}_t) = \text{opt}(\langle \mathcal{A}, G, \zeta_t \rangle, \mathbf{w}_t) - \text{val}(\pi, \mathbf{w}_t) = 0$, as π is indeed an optimal allocation w.r.t. \mathbf{w}_t . \square

Corollary 5.9 (budget-balancedness). *Let A be any optimal allocation algorithm. Then, the mechanism (A, \mathbf{p}^ξ) is budget-balanced.*

Finally, we complete the picture of our analysis by proving a strong fairness property of the proposed payment rule \mathbf{p}^ξ : In words, the best outcome for every agent is always determined by a (global) optimal allocation. Moreover, from Corollary 4.11, any agent is indifferent about the specific optimal allocation being considered. That is, *any* chosen optimal allocation leads to the best results for all agents.

Lemma 5.10. *Let π and π' be two feasible allocations for S_t such that π is optimal w.r.t. \mathbf{w}_t , and hence $\text{val}(\pi, \mathbf{w}_t) \geq \text{val}(\pi', \mathbf{w}_t)$. Then, $u_i(\pi, \mathbf{w}_t) \geq u_i(\pi', \mathbf{w}_t)$, for each $i \in \mathcal{A}$. Moreover, if π' is not optimal, there exists some agent $i \in \mathcal{A}$ such that $u_i(\pi, \mathbf{w}_t) > u_i(\pi', \mathbf{w}_t)$.*

Proof. For any allocation $\bar{\pi}$, consider the coalitional game $\mathcal{G}^{\bar{\pi}} = \langle \mathcal{A}, v^{\bar{\pi}} \rangle$ such that $v^{\bar{\pi}}(\mathcal{C}) = \text{opt}(\langle \mathcal{C}, \text{img}(\bar{\pi}), \zeta_t \rangle, \mathbf{w}_t)$, for each $\mathcal{C} \subseteq \mathcal{A}$. By looking at the expression of the Shapley value for $\mathcal{G}^{\bar{\pi}}$, it is easy to check that $u_i(\bar{\pi}, \mathbf{w}_t) = \phi_i(\mathcal{G}^{\bar{\pi}})$ (just use the same reasoning in the proof of Theorem 4.10). Assume now that π' is an allocation with $\text{val}(\pi, \mathbf{w}_t) \geq \text{val}(\pi', \mathbf{w}_t)$, and consider the value $v^{\pi'}(\mathcal{C}) = \text{opt}(\langle \mathcal{C}, \text{img}(\pi'), \zeta_t \rangle, \mathbf{w}_t)$, for each $\mathcal{C} \subseteq \mathcal{A}$. By Corollary 4.6, we have that $v^\pi(\mathcal{C}) \geq v^{\pi'}(\mathcal{C})$, for each $\mathcal{C} \subseteq \mathcal{A}$. Then, we derive that $u_i(\pi, \mathbf{w}_t) = \phi_i(\mathcal{G}^\pi) \geq \phi_i(\mathcal{G}^{\pi'}) = u_i(\pi', \mathbf{w}_t)$ for every $i \in \mathcal{A}$, because of property (III).

Now assume that π' is not optimal, and thus $\text{val}(\pi, \mathbf{w}_t) > \text{val}(\pi', \mathbf{w}_t)$. Therefore, for the grand-coalition \mathcal{A} , we have $v^\pi(\mathcal{A}) > v^{\pi'}(\mathcal{A})$. Because of property (I) of the Shapley value, only (and all) the total value $v^{\pi'}(\mathcal{A})$ is distributed to agents. It follows that there exists some agent $i \in \mathcal{A}$ such that $u_i(\pi, \mathbf{w}_t) = \phi_i(\mathcal{G}^\pi) > \phi_i(\mathcal{G}^{\pi'}) = u_i(\pi', \mathbf{w}_t)$. \square

By focusing on optimal allocation algorithms, the above lemma immediately entails the following two fairness properties.

Theorem 5.11 (individual optimality). *Let A be any optimal allocation algorithm. Then, for any agent $i \in \mathcal{A}$ and any feasible allocation π for $S_{\mathbf{t}}$, $u_i(A(\mathbf{w}_{\mathbf{t}}), \mathbf{w}_{\mathbf{t}}) \geq u_i(\pi, \mathbf{w}_{\mathbf{t}})$.*

Corollary 5.12 (Pareto-efficiency and envy-freeness). *Let A be any optimal allocation algorithm. Then, the mechanism (A, \mathbf{p}^{ξ}) is Pareto efficient and envy-free.*

Note that individual optimality guarantees much more than classical Pareto efficiency and envy-freeness, because it entails that the mechanism leads to a unique evaluation, independently of the chosen optimal allocation. In particular, the Pareto set is a singleton.

6. Complexity Issues

In this section, we shall reconsider our mechanism with verification from a computational perspective. Note first that computing an optimal allocation on the basis of the reported types is an easy task, which can be carried out via adaptations of classical matching algorithms. Indeed, in the light of Fact 4.2, computing an optimal allocation for an arbitrary scenario (with agents in \mathcal{A} and goods in G) reduces to computing an optimal allocation for a scenario where each agent (in \mathcal{A}^1) gets one good. This is equivalent to finding a matching of maximum weight over a complete bipartite graph over the set of disjoint nodes \mathcal{A}^1 and G , and where edge weights are encoded via the function \mathbf{w}^1 . Such a combinatorial problem is known to be feasible in polynomial time (e.g., Schrijver, 2003).

6.1 Hardness Result

Despite the fact that optimal allocations can be computed in polynomial time, our mechanism is not computationally-efficient, since payments are unlikely to be computable in polynomial time. Indeed, we next show that this computation problem is complete for the complexity class $\#P$ (see Papadimitriou, 1993).

Let us recall that a *counting Turing machine* is a standard nondeterministic Turing machine with an auxiliary output device that prints in binary notation the number of accepting computations induced by the input. It has (worst-case) time complexity $f(n)$ if the longest accepting computation induced by the set of all inputs of size n takes $f(n)$ steps. Then, $\#P$ is the class of all functions that can be computed by counting Turing machines of polynomial time complexity. A prototypical $\#P$ -complete problem is to count the number of truth variable assignments that satisfy a Boolean formula. Of course, $\text{NP} \subseteq \#P$, and a polynomial-time algorithm for solving a $\#P$ -complete problem would imply $\text{P} = \text{NP}$.

Theorem 6.1. *Computing the Shapley value of coalitional games associated with allocation problems (as in Definition 5.1) is $\#P$ -complete.*

Proof. The problem belongs to $\#P$, because computing the Shapley value is known to be feasible in $\#P$ for any class of coalitional games with polynomial-time value/cost functions (c.f. Deng & Papadimitriou, 1994). To show that it is $\#P$ -hard, we exhibit a reduction from the following problem: Let $G = (A \cup B, E)$ be a bipartite graph with $|A| = |B| = n$, $E \subseteq A \times B$, and $|E| = m \geq n$. Recall that a matching is a set $E' \subseteq E$ of edges such that for each pair of distinct edges (a, b) and (a', b') in E' , $a \neq a'$ and $b \neq b'$ hold. The matching E' is *perfect* if $|E'| = n$. The problem of counting the number of perfect matchings in such bipartite graphs is $\#P$ -complete (Valiant, 1979a).

Given a graph $\mathbf{G} = (A \cup B, E)$ as above and a constant $k \geq 1$ (which we shall fix below), we build in polynomial-time a tuple $S(\mathbf{G}) = \langle \mathcal{A}, G, \zeta_{\mathbf{t}} \rangle$ and a vector $\mathbf{w}_{\mathbf{t}}$ such that:

- (1) $\mathcal{A} = \{\alpha\} \cup \bigcup_{(a,b) \in E} \{(a,b)^1, \dots, (a,b)^k\}$, i.e., agents are one-to-one associated with k distinct clones of each edge $(a,b) \in E$, plus a distinguished agent α . Note that $|\mathcal{A}| > n$, because in the considered bipartite graphs $m \geq n$ holds.
- (2) $G = \{g_{\alpha}\} \cup A \cup B$, i.e., goods correspond to nodes, plus a distinguished good g_{α} .
- (3) $\zeta_{\alpha} = 1$ and $\zeta_{(a,b)^i} = 2$, for each $(a,b) \in A$ and $i \in \{1, \dots, k\}$, are the components of $\zeta_{\mathbf{t}}$ associated with agents α and $(a,b)^i$, respectively.
- (4) For each agent $c \in \mathcal{A}$, the associated component w_c in the vector $\mathbf{w}_{\mathbf{t}}$ is defined as follows. For each $(a,b)^i \in \mathcal{A}$, $w_{(a,b)^i}(a) = 2$, $w_{(a,b)^i}(b) = 2$, $w_{(a,b)^i}(g_{\alpha}) = 1$, $w_{(a,b)^i}(x) = 0$, $\forall x \in (A \cup B) \setminus \{a,b\}$. Moreover, $w_{\alpha}(g_{\alpha}) = 1$, $w_{\alpha}(x') = 0$, $\forall x' \in (A \cup B)$.

Let us now fix some notations. For any set $E' \subseteq E$ of edges, let $\text{match}(E')$ denote the size of the largest set $E'' \subseteq E'$ of edges that is a matching. For any set $\mathcal{C} \subseteq \mathcal{A} \setminus \{\alpha\}$ of agents, let $\mathbf{A}(\mathcal{C}) = \{a \mid (a,b)^i \in \mathcal{C}\}$ and $\mathbf{B}(\mathcal{C}) = \{b \mid (a,b)^i \in \mathcal{C}\}$. Finally, we say that $\mathcal{C} \subseteq \mathcal{A} \setminus \{\alpha\}$ is *tight* if it does not contain two agents of the form $(a,b)^i$ and $(a,b)^j$, with $i \neq j$, i.e., associated with the same edge of \mathbf{G} .

Observe that, for each set $\mathcal{C} \subseteq \mathcal{A} \setminus \{\alpha\}$ of agents,

$$\text{opt}(\langle \mathcal{C} \cup \{\alpha\}, G, \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}}) - \text{opt}(\langle \mathcal{C}, G, \zeta_{\mathbf{t}} \rangle, \mathbf{w}_{\mathbf{t}}) = \begin{cases} 1 & \text{if } \mathcal{C} \text{ is tight, and } |\mathcal{C}| = \mathbf{A}(\mathcal{C}) = \mathbf{B}(\mathcal{C}) \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

Indeed, if $\pi'_{\mathcal{C}}$ is an optimal allocation for $\langle \mathcal{C} \cup \{\alpha\}, G, \zeta_{\mathbf{t}} \rangle$ w.r.t. $\mathbf{w}_{\mathbf{t}}$, then we always have that $\text{val}(\pi'_{\mathcal{C}}, \mathbf{t}) = 2 \times |\mathbf{A}(\mathcal{C})| + 2 \times |\mathbf{B}(\mathcal{C})| + 1$. Instead, if $\pi_{\mathcal{C}}$ is an optimal allocation for $\langle \mathcal{C}, G, \zeta_{\mathbf{t}} \rangle$ w.r.t. $\mathbf{w}_{\mathbf{t}}$, then we have

$$\text{val}(\pi_{\mathcal{C}}) = \begin{cases} 2 \times |\mathbf{A}(\mathcal{C})| + 2 \times |\mathbf{B}(\mathcal{C})| & \text{if } \mathcal{C} \text{ is tight, and } |\mathcal{C}| = \mathbf{A}(\mathcal{C}) = \mathbf{B}(\mathcal{C}) \\ 2 \times |\mathbf{A}(\mathcal{C})| + 2 \times |\mathbf{B}(\mathcal{C})| + 1 & \text{otherwise} \end{cases}$$

By exploiting Equation 5, we can now express the Shapley value of the game $\mathcal{G}_{S(\mathbf{G}), \mathbf{t}}^{\text{best}}$ for agent α in a convenient way. Let X_h denote the number of sets $\mathcal{C} \subseteq \mathcal{A} \setminus \{\alpha\}$ of agents which are tight and such that $|\mathcal{C}| = |\mathbf{A}(\mathcal{C})| = |\mathbf{B}(\mathcal{C})| = h$, and let $X_0 = 1$. Then,

$$\phi_{\alpha}(\mathcal{G}_{\mathbf{w}}^{\text{best}}) = \sum_{h=0}^{|\mathcal{A}|-1} \frac{(|\mathcal{A}| - h - 1)!(h)!}{|\mathcal{A}|!} X_h. \quad (6)$$

In particular, let us now focus on the coefficient X_h . Denote by Y_h the number of matchings in \mathbf{G} whose cardinality is h . By construction of $S(\mathbf{G})$ it is immediate to check that for each matching of cardinality h in \mathbf{G} , there are precisely k^h sets of agents $\mathcal{C} \subseteq \mathcal{A} \setminus \{\alpha\}$ that are tight and such that $|\mathcal{C}| = |\mathbf{A}(\mathcal{C})| = |\mathbf{B}(\mathcal{C})| = h$. Thus, we can rewrite the above expression:

$$\phi_{\alpha}(\mathcal{G}_{\mathbf{w}}^{\text{best}}) = \sum_{h=0}^{|\mathcal{A}|-1} (Z_h \times Y_h) \times k^h, \text{ with } Z_h = \frac{(|\mathcal{A}|-h-1)!(h)!}{|\mathcal{A}|!}. \quad (7)$$

For an expression as the one above, given the value of $\phi_\alpha(\mathcal{G}_{\mathbf{w}}^{best})$, it is known that under certain circumstances we can reconstruct in polynomial time the value of each single term of the form $Z_h \times Y_h$ (see Fact 6 in the work by Valiant (1979b)): We need an integer constant $\mathbf{A} > 2$ such that, for each $h \in \{0, \dots, |\mathcal{A}| - 1\}$, $Z_h \times Y_h \leq \mathbf{A}$, and $k \geq \mathbf{A}^2$. In our case, it can be noticed that, for each $h \in \{0, \dots, |\mathcal{A}| - 1\}$, $Z_h \times Y_h \leq 1$ holds, as $Y_h \leq |\mathcal{A}|! / ((h)! (|\mathcal{A}| - h)!)$. Thus, for $k = 9$, we have that, given the value of $\phi_\alpha(\mathcal{G}_{\mathbf{w}}^{best})$, we can compute in polynomial time all such terms. In particular, we can compute in polynomial time the term associated to $h = |A| = |B| = n$, where recall that $|\mathcal{A}| > n$. This term has the form $Z_n \times Y_n$, with Y_n being the number of perfect matchings in \mathbf{G} . Thus, by putting it all together and since Z_n can be computed in polynomial time (as the size of the numbers n and $|\mathcal{A}|$ are logarithmic w.r.t. the size of \mathbf{G}), the number of perfect matchings in bipartite graphs can be counted in polynomial time too, which concludes the proof. \square

By Lemma 4.8 and Theorem 5.5, the following is immediate.

Corollary 6.2. *Computing the payments as given by the rule \mathbf{p}^ξ is $\#P$ -complete.*

6.2 A Fully Polynomial-Time Randomized Approximation Scheme

An approach to circumvent the intractability of the Shapley value is based on approximation: For a game $\mathcal{G} = \langle N, \varphi \rangle$, a vector $\hat{\phi}$ is an ε -approximation of the Shapley value if $|\hat{\phi}_i - \phi_i(\mathcal{G})| \leq \varepsilon \times \phi_i(\mathcal{G})$ holds, for each $i \in N$.

Recently, a sampling method conceived by Bachrach, Markakis, Resnick, Procaccia, Rosenschein, and Saberi (2010) for the special class of *simple* coalitional games has been extended to deal with arbitrary games that are supermodular and *monotone*⁶ (Liben-Nowell, Sharp, Wexler, & Woods, 2012), under the assumption that the value $\varphi(R)$ can be computed by an oracle having unitary cost, for each $R \subseteq N$. The result is that, for any $\varepsilon > 0$ and $\delta > 0$, it is possible to compute in time $\text{poly}(N, 1/\varepsilon, \log(1/\delta))$ a vector $\hat{\phi}$ that is an ε -approximation of the Shapley value with probability of failure at most δ . A method with these properties is called a fully polynomial-time randomized approximation scheme.

Next, we propose a payment rule $\hat{\mathbf{p}}^\xi$ that is founded on the sampling strategy described in the work by Liben-Nowell et al. (2012). The payment rule, reported in Figure 5, samples m subsets of \mathcal{A} storing them in $\hat{\mathbf{C}}$ (together with other subsets functionally determined by the samples), and then computes the value $\xi(\pi, \mathbf{w}_\theta)$ as in Figure 4, but with $\hat{\mathbf{C}}$ playing the role of the power-set \mathbf{C} . The process is repeated $\Theta(\log(1/\delta))$ times, and the component-wise median vector of all such payments is computed. The resulting payment is eventually defined at step 9.

The new rule $\hat{\mathbf{p}}^\xi$ gives rise to a *randomized* mechanism that is of course still verifiable and uses no punishment. Moreover, the mechanism is truthful even when the realization of the set $\hat{\mathbf{C}}$, sampled in step 1, is known beforehand.⁷ Formally, the mechanism turns out to be *universally truthful*, i.e., it is a probability distribution over deterministic truthful mechanisms (see, e.g., Dobzinski & Dughmi, 2009).

6. Monotonicity of $\mathcal{G} = \langle N, \varphi \rangle$ means that $\varphi(R) \geq \varphi(T)$, $\forall T \subseteq R \subseteq N$.

7. This evidences that if truthfulness is our only desideratum, then there is no need to implement a payment rule computing the Shapley value of the coalitional game associated with the allocation problem (as in Definition 5.1), which in fact is a concept defined over all possible subsets of \mathcal{A} and which we have shown to be computationally hard in Theorem 6.1.

<p>Input: A type vector $\theta \in \Theta$, a feasible allocation π for S_θ, and an integer $m > 0$;</p> <p>Assumption: The verifier \mathbf{v} (for \mathbf{t}) is available, with $\mathbf{v}(\pi) = (v_1, \dots, v_n)$;</p> <p>Notation: $S_\theta = \langle \mathcal{A}, G, \zeta_\theta \rangle$, $\mathbf{w}_\theta = (w_1, \dots, w_n)$, $\mathbf{w}_{\mathbf{v}(\pi)} = (w_{v_1}, \dots, w_{v_n})$;</p>
<ol style="list-style-type: none"> 1. Generate a set $\overline{\mathcal{C}}$ of m subsets of \mathcal{A}, and let $\hat{\mathcal{C}} := \overline{\mathcal{C}} \cup \{(\mathcal{A} \setminus \mathcal{C}) \cup \{i\} \mid \mathcal{C} \in \overline{\mathcal{C}}, i \in \mathcal{C}\} \cup \{\mathcal{A}\}$; 2. For each $i \in \mathcal{A}$ and $\mathcal{C} \in \hat{\mathcal{C}}$, 3. Compute an optimal allocation $\pi_{\mathcal{C},i}$ for $\langle \mathcal{C}, \text{img}(\pi), \zeta_{(v_i, \theta_{-i})} \rangle$ w.r.t. $\mathbf{w}_{(v_i, \theta_{-i})}$; 4. Compute an optimal allocation for $\pi_{\mathcal{C} \setminus \{i\},i}$ for $\langle \mathcal{C} \setminus \{i\}, \text{img}(\pi), \zeta_{(v_i, \theta_{-i})} \rangle$ w.r.t. $\mathbf{w}_{(v_i, \theta_{-i})}$; 5. For each agent $i \in \mathcal{A}$, 6. Compute $\xi_i(\pi, \mathbf{w}_\theta)$ as in Figure 4 (steps 4–8), with $\mathcal{C} := \hat{\mathcal{C}}$; 7. Repeat $\Theta(\log(1/\delta))$ times steps 1, 2, and 5, and 8. Let $\hat{\xi}(\pi, \mathbf{w}_\theta)$ be the component-wise median vector of these vectors $\xi(\pi, \mathbf{w}_\theta)$; 9. Define $\hat{p}_i^\xi(\pi, \mathbf{w}_\theta) := \hat{\xi}_i(\pi, \mathbf{w}_\theta) - w_{v_i}(\pi)$;

 Figure 5: Payment rule $\hat{\mathbf{p}}^\xi$.

Theorem 6.3. *Let A be any optimal allocation algorithm. Then, the mechanism $(A, \hat{\mathbf{p}}^\xi)$ is universally truthful.*

Proof. The result follows by inspecting the proof of Theorem 4.9 for rule \mathbf{p}^ξ in Section 4. Indeed, it can be immediately checked that the proof does not depend on the specific subset of coalitions \mathcal{C} , and thus it smoothly applies if any set of coalitions $\hat{\mathcal{C}}$ is used as in Figure 5, instead of all possible subsets of \mathcal{A} . Note in particular that, in the proof of Theorem 4.9, properties (A) and (B) are precisely those guaranteeing truthfulness, and that these properties hold for each given coalition \mathcal{C} . Therefore, they still hold for any subset of coalitions randomly chosen by the mechanism. \square

Similarly, by inspecting the proofs in Section 4, in this universally truthful mechanism, the following properties are seen to hold for every realization of the random coin tosses.

Theorem 6.4 (basic properties). *Let A be any optimal allocation algorithm. Then, the mechanism $(A, \hat{\mathbf{p}}^\xi)$ is efficient and guarantees an equal treatment of the equals. Moreover, if all valuations are non-negative, then $(A, \hat{\mathbf{p}}^\xi)$ is individually rational.*

For a deeper analysis of the payment rule $\hat{\mathbf{p}}^\xi$, we next point out a relationship between utility values and approximations of the Shapley value, at the equilibrium when all agents report their true types.

Lemma 6.5. *Let $\mathcal{A} = \{1, \dots, |\mathcal{A}|\}$, and let $m = \Theta(|\mathcal{A}|^2/\varepsilon^2)$. For each optimal allocation π for $S_{\mathbf{t}}$ w.r.t. $\mathbf{w}_{\mathbf{t}}$, the vector $(u_{1, \hat{\mathbf{p}}^\xi}(\pi, \mathbf{w}_{\mathbf{t}}), \dots, u_{|\mathcal{A}|, \hat{\mathbf{p}}^\xi}(\pi, \mathbf{w}_{\mathbf{t}}))$ is an ε -approximation of the Shapley value of $\mathcal{G}_{\mathbf{w}}^{\text{marg}}$ (and $\mathcal{G}_{\mathbf{w}}^{\text{best}}$) with probability $1 - \delta$, and coincides with it in expectation.*

Proof. By exploiting the same line of reasoning as in the proofs in Section 4 for \mathbf{p}^ξ , we can see that $\xi_i(\pi, \mathbf{w}_{\mathbf{t}})$ (at step 6 of the algorithm in Figure 5) can be rewritten as follows:

$$\xi_i(\pi, \mathbf{w}_{\mathbf{t}}) = \sum_{\mathcal{C} \in \hat{\mathcal{C}}} \frac{(|\mathcal{A}| - |\mathcal{C}|)! (|\mathcal{C}| - 1)!}{|\mathcal{A}|!} (\text{best}_{\mathbf{w}}(\mathcal{C}) - \text{best}_{\mathbf{w}}(\mathcal{C} \setminus \{i\})).$$

Then, we observe that by construction of $\hat{\mathcal{C}}$ in step 1, for each set $\mathcal{C} \in \hat{\mathcal{C}}$ and each agent $i \in \mathcal{A}$, the set $(\mathcal{A} \setminus \mathcal{C}) \cup \{i\}$ is in $\hat{\mathcal{C}}$, too. Thus, we can apply the same line of reasoning as in the proof of Theorem 5.5, in order to conclude that:

$$\xi_i(\pi, \mathbf{w}_t) = \sum_{\mathcal{C} \in \hat{\mathcal{C}}} \frac{(|\mathcal{A}| - |\mathcal{C}|)! (|\mathcal{C}| - 1)!}{|\mathcal{A}|!} (\text{marg}_{\mathbf{w}}(\mathcal{C}) - \text{marg}_{\mathbf{w}}(\mathcal{C} \setminus \{i\})).$$

Now, recall that the game $\mathcal{G}_{\mathbf{w}}^{\text{marg}} = \langle \mathcal{A}, \text{marg}_{\mathbf{w}} \rangle$ is supermodular by Corollary 5.4. Moreover, $\mathcal{G}_{\mathbf{w}}^{\text{marg}}$ is clearly monotone. Thus, by Liben-Nowell et al. (2012), the above expression coincides in expectation with the Shapley value of $\mathcal{G}_{\mathbf{w}}^{\text{marg}}$, and it approximates this value with constant probability. Steps 7 and 8 serve to amplify the probability (c.f. Liben-Nowell et al., 2012), and to get a fully polynomial-time randomized approximation scheme. The result follows as step 9 implements the usual bonus and compensation approach, so that $u_{i, \hat{\mathbf{p}}^\xi}(\pi, \mathbf{w}_t)$ coincides with the above expression, for each $i \in \mathcal{A}$ (cf. Lemma 4.8). \square

Note that in the approach by Liben-Nowell et al. (2012), a final normalization step is carried out that would guarantee that $\sum_{i \in \mathcal{A}} u_{i, \hat{\mathbf{p}}^\xi}(\pi, \mathbf{w}_t) = \text{marg}_{\mathbf{w}}(\mathcal{A}) = \text{best}_{\mathbf{w}}(\mathcal{A})$. Unfortunately, this way truthfulness might be lost, hence we did not include such a normalization procedure in the above payment rule. As a consequence, the mechanism $\hat{\mathbf{p}}^\xi$ does not guarantee (for instance) budget-balancedness and Pareto-efficiency. However, because of Lemma 6.5, since the expected utility profile coincides with the Shapley value, $\hat{\mathbf{p}}^\xi$ enjoys in expectation all the properties of \mathbf{p}^ξ including the above two ones. In particular, we can still have approximate counterparts for Theorem 5.6 and Theorem 5.8.

Theorem 6.6. *Let π be an optimal allocation for S_t w.r.t. \mathbf{w}_t . Let $m = \Theta(|\mathcal{A}|^2/\varepsilon^2)$. Then, with probability $1 - \delta$,*

- $(1 + \varepsilon) \times \text{best}_{\mathbf{w}}(\mathcal{C}) \geq \sum_{i \in \mathcal{C}} u_{i, \hat{\mathbf{p}}^\xi}(\pi, \mathbf{w}_t) \geq (1 - \varepsilon) \times \text{marg}_{\mathbf{w}}(\mathcal{C})$, for each $\mathcal{C} \subseteq \mathcal{A}$;
- $\varepsilon \times \text{val}(\pi, \mathbf{w}_t) \geq \sum_{i \in \mathcal{A}} \hat{p}_i^\xi(\pi, \mathbf{w}_t) \geq -\varepsilon \times \text{val}(\pi, \mathbf{w}_t)$.

Proof. Here, just observe that, in the light of Lemma 6.5 and Theorem 5.5, for each set $\mathcal{C} \subseteq \mathcal{A}$, we have $(1 - \varepsilon) \times \sum_{i \in \mathcal{C}} u_{i, \hat{\mathbf{p}}^\xi}(\pi, \mathbf{w}_t) \leq \sum_{i \in \mathcal{C}} u_{i, \mathbf{p}^\xi}(\pi, \mathbf{w}_t) \leq (1 + \varepsilon) \times \sum_{i \in \mathcal{C}} u_{i, \hat{\mathbf{p}}^\xi}(\pi, \mathbf{w}_t)$. The result then follows by substituting such bounds in Theorem 5.6 and Theorem 5.8, respectively, with simple algebraic manipulations. \square

Finally, we propose a further randomized mechanism that is able to guarantee both economic efficiency and budget-balancedness. The price to be paid is however that truthfulness holds in expectation only. The mechanism is based on a payment rule that we call $\bar{\mathbf{p}}^\xi$.

Theorem 6.7. *Let A be any optimal allocation algorithm. Then, there is a (randomized) mechanism with verification $(A, \bar{\mathbf{p}}^\xi)$ that is truthful in expectation, and that (at the truthful equilibrium) is efficient and budget-balanced. Moreover, if all valuations are non-negative, then $(A, \bar{\mathbf{p}}^\xi)$ is individually rational.*

Proof. The payment rule $\bar{\mathbf{p}}^\xi$ follows the steps in Figure 5, with minor modifications in step 8 and step 9: First, in step 8, whenever we compute the median value $\hat{\xi}_i(\pi, \mathbf{w}_\theta)$ for agent i , we also compute the corresponding value $\hat{\xi}_i(\pi, \mathbf{w}_{\mathbf{v}(\pi)})$ (evaluated on the revealed

types rather than on the reported ones). Then, we define a normalization factor $R = \text{opt}(\langle \mathcal{A}, \text{img}(\pi), \zeta_{\mathbf{v}(\pi)} \rangle, \mathbf{w}_{\mathbf{v}(\pi)}) / (\sum_{i \in \mathcal{A}} \hat{\xi}_i(\pi, \mathbf{w}_{\mathbf{v}(\pi)}))$, so that, in step 9, $\bar{p}_i^\xi(\pi, \mathbf{w}_\theta)$ is eventually returned as $w_{v_i}(\pi) - \hat{\xi}_i(\pi, \mathbf{w}_\theta) \times R$.

Concerning truthfulness, we can just note that the expected value of R is 1. Indeed, by Lemma 6.5, the expected value of $\hat{\xi}_i(\pi, \mathbf{w}_{\mathbf{v}(\pi)})$ is $\xi_i(\pi, \mathbf{w}_{\mathbf{v}(\pi)})$; hence, the sum of all these values coincides with $\text{opt}(\langle \mathcal{A}, \text{img}(\pi), \zeta_{\mathbf{v}(\pi)} \rangle, \mathbf{w}_{\mathbf{v}(\pi)})$ by the efficiency of the Shapley value (as in the proof of Theorem 5.8). Thus, the expected utility of an agent i under the payment rule $\bar{\mathbf{p}}^\xi$ coincides with the (actual, i.e., not in expectation) utility of i under the rule $\hat{\mathbf{p}}^\xi$. Hence, truthfulness in expectation follows by Theorem 6.3. Now, we can just check that, at the truthful equilibrium, the maximum social welfare is achieved (equilibrium efficiency) and $\sum_{i \in \mathcal{A}} \bar{p}_i^\xi(\pi, \mathbf{w}_\mathbf{t}) = \text{opt}(\langle \mathcal{A}, \text{img}(\pi), \zeta_\mathbf{t} \rangle, \mathbf{w}_{\mathbf{v}(\pi)}) - \sum_{i \in \mathcal{A}} \hat{\xi}_i(\pi, \mathbf{w}_\mathbf{t}) \times R = 0$. That is, the mechanism is budget-balanced, too. Finally, the mechanism $\bar{\mathbf{p}}^\xi$ is seen to be individually-rational, by exploiting the same line of reasoning as the one used for the mechanism based on $\hat{\mathbf{p}}^\xi$, since the corresponding proof in Section 4 is not affected by the sampling strategy. \square

Again, by Lemma 6.5, all the remaining properties hold in expectation.

7. Related approaches to Mechanisms with Verification

We next review the main approaches in the literature for mechanisms with verification, and point out the differences w.r.t. our proposal.

First we observe that, differently from our general setting where the types of the agents can determine the vector of the upper bounds for the allocation problem, in earlier approaches it is assumed that agent types do not have any impact on the underlying combinatorial problem, so that the type vector precisely coincides with the valuation vector. Instead, we deal explicitly with both allocation constraints and goods valuations. For instance, in a scheduling problem formalized in our setting, the private type of an agent/machine can be its speed, while its valuation function is fully determined by this speed and the size of the job to be processed. Then, the verifier is just asked to measure the speed of such agents/machines—see also Appendix A. In fact, it is immediate to encode the valuations as types as well, at the price however of hiding the true complexity (or simplicity) of the setting. For instance, in the above example, if the type of an agent is the vector of its valuations, then we would miss the information that the agent is ultimately characterized by one (observable) parameter only.

More substantial differences between earlier approaches and our proposal come in the very definition of the utility of an agent. In the works by Auletta et al. (2009), Penna and Ventre (2012a, 2012b), Krysta and Ventre (2010), Auletta et al. (2006) and Ferrante et al. (2009), the individual welfare of an agent i , given the outcome π and the vector \mathbf{d} of reported types, is assumed to be of the following form:

$$u_{i,\mathbf{p}}(\pi, \mathbf{d}) = t_i(\pi) - \begin{cases} 0 & \text{if } i \text{ is caught lying} \\ p_i(\pi, \mathbf{d}) & \text{otherwise} \end{cases}$$

where $p_i(\cdot, \cdot)$ is a payment that does not depend on the vector \mathbf{t} of the true types.

In these papers, the only information that is assumed to be available at payment time is whether the reported type d_i of agent i differs or not from its actual true type t_i . This is the

role played by the verifier, and in fact $p_i(\pi, \mathbf{d})$ does not exploit the possibility of partially revealing t_i . Moreover, the payment scheme adopted punishes those agents that are caught lying. Hence, while the verification process provides a smaller amount of information than the verification process in our approach, the rules used to discourage strategic behaviors are stronger than ours and based on punishing agents. Also, the above works assume that agents' misreporting is restricted to certain kinds of lies (e.g., values lower than the corresponding true ones), so that a form of "one-sided" verification suffices.

Recently, the above model of (partial) verification has been extended by Caragiannis et al. (2012) to a setting where an agent cheating on her/his type will be identified with some probability that may depend on her/his true type, the reported type, or both. In fact, Caragiannis et al. also showed that there are cases in which verification does not help if it is not one-sided. The payment scheme is exactly the same as the one discussed above and, hence, verification does not exploit any (possibly partial) knowledge of the actual true type and a punishment approach is still used. The main novelty, in addition to the probabilistic verification, is that there is no constraint on the type that an agent can report while cheating.

Finally, a different kind of verification model goes back to the seminal paper by Nisan and Ronen (2001), and it is actually closer to our "no-punishment" perspective, because an agent i can in principle be paid by the mechanism even if i has been caught lying. Given n agents, Nisan and Ronen consider a vector $\mathbf{e} = (e_1, \dots, e_n)$ of "observed" agent types, which are completely known after the verification process. Moreover, the individual utility of any agent i has the form $u_{i,\mathbf{p}}(\pi, \mathbf{d}) = e_i(\pi) - p_i(\pi, \mathbf{d})$, so that the vector \mathbf{e} in such a framework plays the same role as the verifier's output in our approach. A first difference between the work by Nisan and Ronen and our approach is that, in the above model, agents' misreporting is again restricted only to certain kinds of lies. On the other hand, we consider a restriction on the form of valuation functions, since our verification model assumes that valuations are determined by some objective and observable properties of goods and agents (encoded by functions of the form λ_g and χ_i , as defined in Section 2). As a consequence, at payment time, the valuation of any agent i is known for every good in $\text{img}(\pi)$. Instead, in the setting by Nisan and Ronen, the valuation of each agent i for goods in $\text{img}(\pi) \setminus \pi(i)$ remains unknown even after the verification is performed. From the results presented in this paper, it turns out that this difference between the two framework is crucial to overcome classical impossibility results, and to meet all desirable properties at once (without using any punishing power).

8. Conclusion

In this paper, we have proposed and analyzed mechanisms for fair allocation problems in a setting where agents' declarations regard objective properties of goods or agents, and thus can be (partially) verified before payments are made. In particular, we have considered a model of verification that is able to disclose the true values of allocated goods, in contrast to previous approaches in the literature where partial and probabilistic verification have been considered. However, the use of this verification power is in fact quite limited, because payment rules have been designed so that no punishment is meted out to agents whose declarations do not match the output of the verification process. This requirement is crucial

for practical applications of the framework, such as the Italian research evaluation described in (Greco & Scarcello, 2013) (and summarized in the Appendix), because any discrepancy between declared and verified values may be due to sensing errors or subjective issues that cannot be interpreted as agents' lies to be punished.

The challenge was to show that, in this framework, truthfulness, efficiency, budget-balancedness, and fairness (as well as further desirable properties discussed in Section 2) can be achieved simultaneously, unlike the classical setting. It is worthwhile noting that, if one is guaranteed that all agents truthfully report their true types (no strategic behaviors) or, equivalently, if these types are given as public knowledge, then the problem is quite easy. For instance, one can just use payments in a way that each agent gets as her/his utility the Shapley value according to either coalitional game defined in Section 5. Moreover, if fairness is not an issue, then the other properties can be obtained by even using the simple "uniform" payment rule. However, whenever agents may behave strategically, reporting the true type is no longer a dominant strategy, if such approaches are used. In fact, truthfulness might be enforced by equipping the mechanism with suitable punishment rules. However, we already argued that this would be not appropriate in our context, and typically the resulting mechanism would not be error tolerant.

By looking at the proposed framework from an abstract perspective, one may notice that it is based on two fundamental ingredients: a base combinatorial problem that determines feasible and optimal allocations, and a game-theoretic notion that describes what is considered fair, with respect to agents' contributions and expectations. Namely, in the application domain of allocation problems addressed in the paper, the weighted matching is the basic combinatorial problem and the Shapley value is the most natural game-theoretic solution concept. Under this perspective, a natural research direction is to study different instances of such an abstract framework for mechanisms with verification, where other combinatorial problems (colorings, coverings, etc.) and different solution concepts (Nucleolus, Banzhaf index, etc.) may be more appropriate and best describe the problem at hands.

Another interesting avenue for further research would be to study a modification of the framework where agents' preferences are not directly expressed in terms of real-valued functions, but they are rather formalized in terms of orderings/ranks over available goods. This is particularly interesting given that, in the strategic setting we have defined, agents' declarations only contribute to the definition of which goods will be selected, but they do not determine their values.

Acknowledgments

We thank the anonymous referees and the Associate Editor for their very useful comments.

Appendix A. Application Scenarios

Fair allocation with monetary compensation has been intensively studied in the literature, very often in settings where all parameters of interest are public knowledge (or, simply, by getting rid of strategic issues). One example application discussed in the literature is parking space and benefit allocation at a workplace, where each employee gets a parking space and a share from a fixed benefits package. House allocation problems are another

classical example, where agents collectively own a set of houses, and we look for a systematic way of exclusively assigning a house to each agent, possibly with monetary compensations. A third example is room assignment-rent division, where a group of agents rent a house, with each of them getting a room and paying a share of the rent.

In the following, we illustrate further applications of fair allocation problems that fit our general framework discussed in Section 2. In the applications we shall discuss, part of the relevant information is private knowledge of the agents, and verification can be adopted to measure observable properties before the payment phase. Nonetheless, we stress here that, even if all the relevant information were available as public knowledge or could be measured in advance, such applications would still remain of interest, as they show the need of defining allocation policies that guarantee fair and/or error tolerant solutions.

A.1 The Italian Research Assessment Program (VQR) 2004–2010

We start our overview of possible application scenarios by focusing on a real-world case study that first motivated our investigation on allocation problems. This application is best described in a companion paper (Greco & Scarcello, 2013), which applies the results presented here to the specific case of the Italian research assessment program.

A.1.1 THE SETTING

In 2012, the National Agency for the Evaluation of Universities and Research Institutes (ANVUR) has promoted the ‘VQR’ assessment program devoted to evaluate the quality of the whole Italian research production. In its first application, the program focuses on the period 2004-2010, while the evaluation will be repeated on a regular basis (the next one should cover years 2011–2014).

In the first phase of the program, every structure R (a university or a research institute) selects and submits to ANVUR a set \mathcal{P}_π of its products, under the constraints that (i) each product has to be univocally associated with an author (even if the product is co-authored) and that (ii) at most three products can be associated with each author affiliated with R .⁸ In abstract terms, R computes in this phase an allocation π , with $\text{img}(\pi) = \mathcal{P}_\pi$, for the scenario where \mathcal{R} is the set of the agents/researchers affiliated with R , where \mathcal{P} is the set of all the goods/products they have co-authored, and where for each $r \in \mathcal{R}$, $\zeta_r = 3$ is the associated constraint on the number of goods/products that can be allocated to r . The way this allocation is performed is described below.

In the second phase, ANVUR evaluates the products in \mathcal{P}_π by equipping each of them with a quality score, expressed as a number.⁹ Hence, this phase defines an official valuation such that, for each product $p \in \pi(r)$, $w(p)$ is the quality score assigned by ANVUR to p . The overall score of R will be the sum of the values of all products in \mathcal{P}_π , and it will be used to proportionally transfer to R the funds allocated by the Ministry to support research activities in the next years, until data from a new evaluation for the subsequent period will be available.

8. We simplify here. Actually, the number of publications may be less than three, for some authors.

9. The set of the possible scores is defined in the VQR guidelines. To our ends, this detail is immaterial and scores are just viewed as (arbitrary) real numbers.

A.1.2 THE NEED OF A FAIR DIVISION

The VQR program actually assigns a score not only to the structure R , but also to all its substructures (e.g., to departments, if R is a university). Of course, this is expected to have an impact on the funds redistribution inside every research structure, and therefore a crucial problem is to define a fair rule for funds redistribution, that is, a rule that is capable to assign funds by clearly reflecting the true contribution of each researcher/substructure to the performances of the structure as a whole. Moreover, the recruitment policies of universities are going to be evaluated as well, by looking at the (VQR) performances of researchers that were hired recently. However, no redistribution rule has been specified in the program, and most researchers believe that the evaluation of (the contribution of) each $r \in \mathcal{R}$ will just coincide with the overall value $w_r(\pi) = \sum_{p \in \pi(r)} w(p)$ of those products allocated to r in the submission phase. Of course this is not fair in general, because if a product $p \in \mathcal{P}_\pi$ is co-authored by two researchers of R , then they both (and, in turn, their structures) might claim this contribution even if the product has been formally allocated just to one of them.¹⁰ Therefore, more sophisticated approaches have to be defined for a proper evaluation of individuals and substructures.

A.1.3 STRATEGIC ISSUES AND VERIFICATION IN THE VQR PROGRAM

Recall that the goal of structure R is to submit to ANVUR (in the first phase of the program) the products that will likely get the highest possible scores (in the second phase). To this end, there are publicly available evaluation criteria that should allow one to perform a ranking of products, ideally by equipping each of them with the quality score that will be assigned to it by ANVUR in the subsequent phase. Both for time and economic reasons, this first evaluation phase was based only on self-evaluations performed by the authors of the products, which are clearly assumed to be the best experts on their research subjects. Moreover, in most structures, to decide the precise allocation of products, hence to deal with the conflicts related to co-authored products, researchers performed choices in a decentralized way; only a few structures set up an optimization framework to compute an optimal allocation, based on agents' declarations. In any case, by abstracting from the specific method adopted to end up with a feasible allocation to be submitted for the VQR, strategic issues emerged in this phase. In particular, there are some researchers $r \in R$ that "guided" the allocation of the products (e.g., cheating on their quality) for the supposed personal interest of maximizing the value $w_r(\pi)$. Clearly enough, this way optimal product selections can hardly be achieved by the structure.

Note that the VQR case perfectly fits our framework of allocation problems: here the research products are the indivisible goods to be allocated to researchers/agents, who initially declare their goods' valuations. Moreover, the social welfare is the total "ANVUR score" of the research structure, that should be distributed in a budget-balanced way to researchers, and hence to their substructures, in a fair way. Note that the upper bound constraint for each researcher (3 products, with some exceptions) does not depend on the type in this case,

10. This problem does not occur for products co-authored by researchers belonging to different structures. Indeed, the same product can be submitted by several structures (assigned to a different co-author in each of them). Thus, for a given research structure R , the co-authors of interest are only those affiliated with R .

because it is a parameter fixed by ANVUR. Therefore, the private information of agents is limited to good valuations. Moreover, observe that the specific valuation functions occurring in the VQR fit the verification model proposed in the paper. In particular, every research product has an objective value, which can be “measured” by a verifier for each allocated and hence submitted product (while nothing can be said for products that are not allocated to any researcher). Thus, ANVUR precisely acts as the verifier of the model.

A.1.4 MECHANISMS FOR THE VQR PROGRAM

The basic approach of assigning to each researcher $r \in \mathcal{R}$ the value $w_r(\pi)$ is precisely the (trivial no-payment) rule \mathbf{p}° in Example 2.6. There, we have observed that mechanisms based on this rule are not truthful in general. Moreover, from Example 3.1 is also emerged that the rule is not fair (in particular) when combined with an optimal allocation algorithm (in order to to guarantee efficiency). Therefore, this rule is highly undesirable, and different kinds of mechanisms have to be defined for the VQR setting.

In fact, we would like to end up with a mechanism that is able as able to collect the correct self-evaluations (truthfulness) and to obtain the maximum possible performance (efficiency) for the research structures. These two goals, however, have to be accomplished together with fairness. Indeed, we would like to guarantee to every researcher or group of researchers to get at least their marginal contributions, and that the resulting mechanism is also individually optimal. The latter property is very important for this application, since it guarantees that, for any researcher, her/his score is the maximum one over all possible alternative allocations, including allocations using products that were not submitted by the structure and hence not verified. Moreover, it is also relevant that these properties are guaranteed via mechanisms enjoying the verifiability property, in that the payment rule only uses “certified” valuations (it would be unacceptable to use product values that are not verified by ANVUR).

Finally, it is worthwhile noting that ANVUR guidelines defines a range of possible product values determined, e.g., by the publishing venues and by citation indices, by pointing out that peer-reviews can also be used to override such a basic classification (e.g., for papers published on a good journal but without many citations, or after an explicit author’s request). Clearly enough, this entails that discrepancies between authors’ declarations and ANVUR evaluations might well emerge even in absence of any malicious behavior. Hence, the no-punishment (and error-tolerant) approach seems to be unavoidable for a concrete and politically acceptable application of any mechanism.

A.2 (Cooperative) Scheduling and Task Allocation

Assume that h jobs in the set $\mathcal{J} = \{j_1, \dots, j_h\}$ have to be executed within a deadline d , and that n agents/machines, m_1, \dots, m_n , with $n \leq h$, are available to execute in parallel these jobs. For each machine m_i , with $i \in \{1, \dots, n\}$, let $\gamma_i \subseteq \mathcal{J}$ denote the set of jobs that m_i can in principle execute. These sets are determined, for instance, by physical, technological, or accessibility constraints and are known to the scheduler, while the precise speed that the machines guarantee for this process is not known to the scheduler, and should be declared by agents. It is assumed that each machine m_i can always work at such a speed independently of the workload actually assigned to it (for the considered jobs γ_i at hand). For the sake of

simplicity, assume also that each job consists of the same workload wl , and that for each executed job, a fixed profit pr is earned. Moreover, a further profit pr_{add} will be earned by the scheduler if all jobs are correctly executed, and (part of it) can be distributed to agents. Every machine participating in the process will execute at least one job (i.e., we are assuming we have many jobs that can be executed by any machine and machines have comparable speeds so that it is convenient to use all of them). Furthermore, every machine is dedicated to this process and it aims at executing as many jobs from \mathcal{J} as possible, because it cannot earn any profit from external jobs (not in \mathcal{J}).

Therefore, we have an allocation problem where the goods/jobs in the set $\{j_1, \dots, j_h\}$ have to be allocated/scheduled to the agents/machines in the set $\{m_1, \dots, m_n\}$. Here, the private type of agent m_i is its speed s_i (devoted to this process), for each $i \in \{1, \dots, n\}$. Moreover, the vector ζ of the upper bounds to the number of goods that can be allocated to the agents is defined as $\zeta_i(s_i) = \lfloor \frac{d \times s_i}{wl} \rfloor$, for each $i \in \{1, \dots, n\}$. Note, in particular, that the jobs have to be completed within the deadline, so that the upper bound constraints of the allocation scenario are functions of the types/speeds of agents/machines. Finally, we can define the valuation vector $\mathbf{w} = (w_1, \dots, w_h)$ so that, for each $i \in \{1, \dots, n\}$, $w_i(j_x) = pr$, if $j_x \in \gamma_i$; and $w_i(j_x) = -1$, if $j_x \notin \gamma_i$. Note that in this case valuations are independent on the types of the agents (of course, we are assuming that any single job in γ_i can be executed within the deadline d).

Given that the same task can be carried out by different agents, it is sensible that allocations are perceived to be as fair ones (see, e.g., Porter et al., 2004). To this end, the mechanism with monetary compensations described in this paper can be adopted. In fact, note that strategic issues come into play, and the egoistic behavior of some agents may lead to allocations that are not optimal and possibly can miss the extra-reward pr_{add} . Observe that the scheduler may act as the verifier in our model, at the end of the process. Indeed, after the allocation is computed and jobs are executed based on it, it can immediately verify the truthfulness of agents' declarations by looking at the amount of time used by each machine to execute the jobs assigned to it.

Furthermore, in practice, speed values are given with a finite precision, as well as their physical verification is subject to measurement errors. For this reason, even this mechanism should be tolerant of errors, and not based on punishments.

As a further example, consider the following *task allocation problem*, which can be viewed as a variation of the previous one: Assume that a company select some agents e_1, \dots, e_n to perform a given set of tasks $\{t_1, \dots, t_m\}$ within a certain deadline, and assume that the company does not know precisely whether or not such agents have the necessary skills (experience, strength, speed, competence, etc.). For instance, this may happen if the company is starting some new line of production with new tasks, or if such tasks should be executed by means of crowdsourcing, where agents are selected through an internet call.

Therefore, in this case agents' types declarations comprise both their skills and the number of tasks they are able to execute within the required deadline. Valuation functions encode the profit to be earned by any agent executing a given task (where the value 0 means that the agent is not able to execute that task). Note that here both the upper bound constraints of the allocation scenario and the valuations functions depend on agents' types. Again, a mechanism with monetary compensations can be used, in order to provide a fair distribution of the company reward for those tasks, and to encourage agents to truthfully

declare their skills. Note that the framework with verification proposed in this paper can be used under the assumption that the agents’ skills necessary for the proposed set of tasks can be observed and evaluated by the verifier, (at least) at the end of the process, where all agents performed their work (possibly with failures).

A.3 Protocols for Wireless Communication Networks

We conclude this overview of possible applications by considering a cost problem where, moreover, there is no private information, and hence no strategic behavior. This is to remark that the proposal described in the paper can be used even when mechanism design is not necessary, but one still needs a policy for fair division that enjoys the properties described in Section 3.1 (e.g., envy freeness, individual optimality, and so on). In fact, fairness issues are currently attracting much attention in the design of scheduling protocols over wireless communication networks, where the underlying problem is bandwidth allocation; for instance, in the design of protocols for high-speed wide area wireless networks, where the role of monetary compensation is played by adjustments in the priorities of users (see, e.g., Jalali, Padovani, & Pankaj, 2000). Moreover, a number of network applications are emerging where “direct” forms of monetary compensation are considered.

As an example, let us consider *ad-hoc networks*, which are self-organizing wireless infrastructures where mobile nodes cooperatively act via multi-hop routing to transmit data even when source and destination nodes are out of their transmission ranges. Cooperation can be achieved by associating a credit balance with each node, so that nodes use credits to pay for sending their own traffic, and earn credits by forwarding traffic from other nodes to compensate bandwidth and power consumption—see, for instance, the work by Göbel, Krzesinski, and Mandjes (2009) and the references therein. Consider a setting where the nodes s_1, \dots, s_n are willing to transmit some data in an ad-hoc network. In a given configuration of the network, the nodes r_1, \dots, r_m can be used to forward these data. In particular, whenever s_i and r_j , with $i \in \{1, \dots, n\}$ and $j \in \{1, \dots, m\}$, are within their transmission ranges, we denote by $c_{i,j}$ the credit paid by s_i to transfer the data via r_j —a large enough value for $c_{i,j}$ can be used to state that the transmission is not possible. The resulting cost problem can be modeled via an allocation scenario where $\zeta_i = 1$, and where the valuation vector is such that $w_i(r_j) = -c_{i,j}$, for each $i \in \{1, \dots, n\}$ and $j \in \{1, \dots, m\}$. Thus, maximizing the social welfare amounts to minimizing the overall credits paid by source nodes.

It is easily seen that, even if there is no private information and everything is public, fairness issues emerge in this context, too. Indeed, different sources may want to use the same routing node to transfer data. For instance, in a simple setting where $m = 2$ and $n = 2$ and where r_1 is the most preferred routing node for both s_1 and s_2 (i.e., $c_{1,1} < c_{1,2}$ and $c_{2,1} < c_{2,2}$), the node that is “forced” to transfer data via r_2 might perceive the given allocation as unfair. This suggests that the credits paid by source nodes should be “adjusted” via payment rules that, in our proposal, are based on the Shapley value of the coalitional games described in Section 5.

A setting similar to the one discussed above emerges in wireless *cooperative file sharing systems*, where mobile subscribers cluster together by downloading (portions of) files of interest over long-range cellular links, and by exchanging them over short-range radio communications in a wireless local area network. These systems are cooperative environments,

whose benefits can be appreciated not only in terms of increased throughput and reduced energy consumption, but also in terms of economic advantages both for users and content providers. In order to be effective, however, fair allocation protocols have to be designed, whose goal is to encourage cooperation (see, e.g., Militano, Iera, & Scarcello, 2013).

References

- Abdulkadiroğlu, A., Sönmez, T., & Ünver, M. U. (2004). Room assignment-rent division: A market approach. *Social Choice and Welfare*, *22*, 515–538.
- Ågotnes, T., van der Hoek, W., Tennenholtz, M., & Wooldridge, M. (2009). Power in normative systems. In *Proc. of AAMAS'09*, pp. 145–152.
- Alcalde, J., & Barberà, S. (1994). Top dominance and the possibility of strategyproof stable allocations to matching problems. *Economic Theory*, *4*, 417–435.
- Alkan, A., Demange, G., & Gale, D. (1991). Fair allocation of indivisible goods and criteria of justice. *Econometrica*, *59*(4), 1023–39.
- Andersson, T. (2009). A general strategy-proof fair allocation mechanism revisited. *Economics Bulletin*, *29*(3), 1717–1722.
- Andersson, T., & Svensson, L.-G. (2008). Non-manipulable assignment of individuals to positions revisited. *Mathematical Social Sciences*, *56*(3), 350–354.
- Andersson, T., Svensson, L.-G., & Ehlers, L. (2010). Budget-balance, fairness and minimal manipulability. Working papers 2010:16, Lund University, Department of Economics.
- Aragones, E. (1995). A derivation of the money rawlsian solution. *Social Choice and Welfare*, *12*, 267–276.
- Archer, A., & Tardos, E. (2007). Frugal path mechanisms. *ACM Transactions on Algorithms*, *3*, 1–22.
- Auletta, V., De Prisco, R., Penna, P., & Persiano, G. (2009). The power of verification for one-parameter agents. *Journal of Computer and System Sciences*, *75*, 190–211.
- Auletta, V., De Prisco, R., Penna, P., Persiano, G., & Ventre, C. (2006). New constructions of mechanisms with verification. In *Proc. of ICALP'06*, pp. 596–607.
- Auletta, V., Penna, P., Persiano, G., & Ventre, C. (2011). Alternatives to truthfulness are hard to recognize. *Autonomous Agents and Multi-Agent Systems*, *22*(1), 200–216.
- Aumann, R. J., & Maschler, M. (1985). Game-theoretic analysis of a bankruptcy problem from the talmud. *Journal of Economic Theory*, *36*(2), 195–213.
- Bachrach, Y., Markakis, E., Resnick, E., Procaccia, A. D., Rosenschein, J. S., & Saberi, A. (2010). Approximating power indices: theoretical and empirical analysis. *Autonomous Agents and Multi-Agent Systems*, *20*, 105–122.
- Bachrach, Y., & Rosenschein, J. S. (2008). Coalitional skill games. In *Proc. of AAMAS'08*, pp. 1023–1030.
- Bachrach, Y., & Rosenschein, J. (2009). Power in threshold network flow games. *Autonomous Agents and Multi-Agent Systems*, *18*(1), 106–132.

- Bachrach, Y., Zuckerman, M., Wooldridge, M., & Rosenschein, J. (2013). Proof systems and transformation games. *Annals of Mathematics and Artificial Intelligence*, 67(1), 1–30.
- Beviá, C. (1998). Fair allocation in a general model with indivisible goods. *Review of Economic Design*, 3, 195–213.
- Bouveret, S., & Lang, J. (2008). Efficiency and envy-freeness in fair division of indivisible goods: Logical representation and complexity. *Journal of Artificial Intelligence Research*, 32, 525–564.
- Brams, S. J., & Kilgour, D. M. (2001). Competitive fair division. *Journal of Political Economy*, 109(2), 418–443.
- Brandt, F., Conitzer, V., & Endriss, U. (2012). *Multiagent Systems*, chap. Computational Social Choices. MIT Press.
- Caragiannis, I., Elkind, E., Szegedy, M., & Yu, L. (2012). Mechanism design: from partial to probabilistic verification. In *Proc. of EC'12*, pp. 266–283.
- Clarke, E. (1971). Multipart pricing of public goods. *Public Choice*, 8, 1933.
- Deng, X., & Papadimitriou, C. H. (1994). On the complexity of cooperative solution concepts. *Mathematics of Operations Research*, 19, 257–266.
- Dobzinski, S., & Dughmi, S. (2009). On the power of randomization in algorithmic mechanism design. In *Proc. of FOCS'09*, pp. 505–514.
- Dunne, P. E. (2005). Extremal behaviour in multiagent contract negotiation. *Journal of Artificial Intelligence Research*, 23, 41–78.
- Dunne, P. E., Wooldridge, M., & Laurence, M. (2005). The complexity of contract negotiation. *Artificial Intelligence*, 164(1-2), 23–46.
- Endriss, U., Maudet, N., Sadri, F., & Toni, F. (2006). Negotiating socially optimal allocations of resources. *Journal of Artificial Intelligence Research*, 25, 315–348.
- Feige, U., & Tennenholtz, M. (2011). Mechanism design with uncertain inputs: (to err is human, to forgive divine). In *Proc. of STOC'11*, pp. 549–558.
- Ferrante, A., Parlato, G., Sorrentino, F., & Ventre, C. (2009). Fast payment schemes for truthful mechanisms with verification. *Theoretical Computer Science*, 410, 886–899.
- Göbel, J., Krzesinski, A., & Mandjes, M. (2009). Incentive-based control of ad hoc networks: A performance study. *Computer Networks*, 53(14), 2427–2443.
- Greco, G., & Scarcello, F. (2013). Fair division rules for funds distribution: The case of the italian research assessment program (vqr 2004-2010). *Intelligenza Artificiale*, 7(1), 45–56.
- Green, J., & Laffont, J. (1977). Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica*, 45(2), 427–438.
- Groves, T. (1973). Incentives in teams. *Econometrica*, 41, 617631.
- Haake, C.-J., Raith, M. G., & Su, F. E. (2002). Bidding for envy-freeness: A procedural approach to n-player fair-division problems. *Social Choice and Welfare*, 19(4), 723–749.

- Hurwicz, L. (1975). On the existence of allocation systems whose manipulative nash equilibria are pareto optimal. *Unpublished paper, presented at the third World Congress of the Economic Society, Toronto.*
- Jain, K., & Vazirani, V. (2001). Applications of approximation algorithms to cooperative games. In *Proc. of STOC'01*, pp. 364–372.
- Jalali, A., Padovani, R., & Pankaj, R. (2000). Data throughput of cdma-hdr a high efficiency-high data rate personal communication wireless system. In *Proc. of IEEE VTC'00*, Vol. 3, pp. 1854–1858.
- Kalai, E., & Samet, D. (1983). On weighted Shapley values. Discussion papers 602, Northwestern University, Center for Mathematical Studies in Economics and Management Science.
- Klijn, F. (2000). An algorithm for envy-free allocations in an economy with indivisible objects and money. *Social Choice and Welfare*, 17(2), 201–215.
- Krysta, P., & Ventre, C. (2010). Combinatorial auctions with verification are tractable. In *Proc. of ESA'10*, pp. 39–50.
- Liben-Nowell, D., Sharp, A., Wexler, T., & Woods, K. (2012). Computing Shapley value in supermodular coalitional games. In *Proc. of COCOON'12*, pp. 568–579.
- Lindner, C. (2010). A market-affected sealed-bid auction protocol. In *Proc. of SETN'10*, pp. 193–202.
- Lipton, R. J., Markakis, E., Mossel, E., & Saberi, A. (2004). On approximately fair allocations of indivisible goods. In *Proc. of EC'04*, pp. 125–131.
- Maniquet, F. (2003). A characterization of the Shapley value in queueing problems. *Journal of Economic Theory*, 109(1), 90–103.
- Maskin, E. (1987). *On the Fair Allocation of Indivisible Goods*, pp. 341–349. MacMillan.
- Meertens, M., Potters, J., & Reijnierse, H. (2002). Envy-free and pareto efficient allocations in economies with indivisible goods and money. *Mathematical Social Sciences*, 44(3), 223–233.
- Militano, L., Iera, A., & Scarcello, F. (2013). A fair cooperative content-sharing service. *Computer Networks*, 57(9), 1955–1973.
- Mishra, D., & Rangarajan, B. (2007). Cost sharing in a job scheduling problem. *Social Choice and Welfare*, 29(3), 369–382.
- Moulin, H. (1992). An application of the Shapley value to fair division with money. *Econometrica*, 60(6), 1331–49.
- Moulin, H. (1999). Incremental cost sharing: Characterization by coalition strategy-proofness. *Social Choice and Welfare*, 16(2), 279–320.
- Moulin, H. (2003). *Fair Division and Collective Welfare*. MIT Press.
- Moulin, H., & Shenker, S. (2001). Strategyproof sharing of submodular costs: budget balance versus efficiency. *Economic Theory*, 18(3), 511–533.
- Nagamochi, H., Zeng, D.-Z., Kabutoya, N., & Ibaraki, T. (1997). Complexity of the minimum base game on matroids. *Mathematics of Operations Research*, 22, 146–164.

- Nisan, N., & Ronen, A. (2001). Algorithmic mechanism design. *Games and Economic Behavior*, 35, 166–196.
- Nisan, N., Roughgarden, T., Tardos, É., & Vazirani, V. V. (2007). *Algorithmic Game Theory*. Cambridge University Press, Cambridge, UK.
- Ohseto, S. (2004). Implementing egalitarian-equivalent allocation of indivisible goods on restricted domains. *Economic Theory*, 23, 659–670 (2004).
- Osborne, M. J., & Rubinstein, A. (1994). *A Course in Game Theory*. The MIT Press, Cambridge, MA, USA.
- Papadimitriou, C. H. (1993). *Computational Complexity*. Addison-Wesley.
- Pathak, A.P., S. T. (2013). Comparing mechanisms by their vulnerability to manipulation. *The American Economic Review*, 103(27), 80–106.
- Penna, P., & Ventre, C. (2012a). Collusion-resistant mechanisms with verification yielding optimal solutions. *ACM Transaction on Computation Theory*, 4(2), 1–17.
- Penna, P., & Ventre, C. (2012b). Optimal collusion-resistant mechanisms with verification. *Games and Economic Behavior*, in press, electronically available with doi 10.1016/j.geb.2012.09.002.
- Porter, R., Shoham, Y., & Tennenholtz, M. (2004). Fair imposition. *Journal of Economic Theory*, 118(2), 209–228.
- Potthoff, R. F. (2002). Use of linear programming to find an envy-free solution closest to the bramskilogour gap solution for the housemates problem. *Group Decision and Negotiation*, 11, 405–414.
- Quinzii, M. (1984). Core and competitive equilibria with indivisibilities. *International Journal of Game Theory*, 13, 41–60.
- Sakai, T. (2007). Fairness and implementability in allocation of indivisible objects with monetary compensations. *Journal of Mathematical Economics*, 43(5), 549–563.
- Sandholm, T. (1998). Contract types for satisficing task allocation: I theoretical results. In *AAAI Spring Symposium: Satisficing Models*.
- Schrijver, A. (2003). *Combinatorial Optimization: Polyhedra and Efficiency*. Springer-Verlag.
- Shioura, A., Sun, N., & Yang, Z. (2006). Efficient strategy proof fair allocation algorithms. *Journal of the Operations Research Society of Japan*, 49(2), 144–150.
- Shoham, Y., & Leyton-Brown, K. (2009). *Multiagent Systems*. Cambridge University Press.
- Su, F. (1999). Rental harmony: Sperners lemma in fair division. *American Mathematical Monthly*, 106, 930942.
- Svensson, L.-G. (1983). Large indivisibles: An analysis with respect to price equilibrium and fairness. *Econometrica*, 51(4), pp. 939–954.
- Svensson, L.-G. (2009). Coalitional strategy-proofness and fairness. *Economic Theory*, 40, 227–245.

- Tadenuma, K., & Thomson, W. (1991). No-envy and consistency in economies with indivisible goods. *Econometrica*, 59(6), 1755–67.
- Tadenuma, K., & Thomson, W. (1993). The fair allocation of an indivisible good when monetary compensations are possible. *Mathematical Social Sciences*, 25(2), 117–132.
- Tadenuma, K., & Thomson, W. (1995). Games of fair division. *Games and Economic Behavior*, 9(2), 191–204.
- Thomson, W. (2011). Fair allocation rules. In Kenneth J. Arrow, A. S., & Suzumura, K. (Eds.), *Handbook of Social Choice and Welfare*, Vol. 2, pp. 393–506. Elsevier.
- Valiant, L. G. (1979a). The complexity of computing the permanent. *Theoretical Computer Science*, 8(2), 189–201.
- Valiant, L. G. (1979b). The complexity of enumeration and reliability problems. *SIAM Journal on Computing*, 8(3), 410–421.
- Vickery, W. (1961). Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance*, 8–37.
- Willson, S. J. (2003). Money-egalitarian-equivalent and gain-maximin allocations of indivisible items with monetary compensation. *Social Choice and Welfare*, 20, 247–259.
- Yang, Z. (2001). An intersection theorem on an unbounded set and its application to the fair allocation problem. *Journal of Optimization Theory and Applications*, 110, 429–443.
- Yengin, D. (2012). Egalitarian-equivalent groves mechanisms in the allocation of heterogeneous objects. *Social Choice and Welfare*, 38(1), 137–160.
- Young, H. P. (1985). Monotonic solutions of cooperative games. *International Journal of Game Theory*, 14, 65–72.
- Young, H. P. (1994). *Equity in Theory and Practice*. Princeton University Press.