Fairness and Welfare in Division of Goods When Utility is Transferable

Ruggiero Cavallo
Yahoo! Research
111 West 40th Street
New York, NY 10018
cavallo@yahoo-inc.com

Abstract

We join the goals of two giant and related fields of research in group decision-making whose connection has historically been underdeveloped: fair division, and efficient mechanism design with monetary payments. To do this we assume a context where utility is quasilinear and thus transferable across agents. We generalize the traditional binary criteria of envy-freeness, proportionality, and efficiency to measures of degree that range between 0 and 1. We observe the impossibility of achieving optimal social welfare with strategic agents in allocation of divisible or indivisible goods. We then set as the goal a strategyproof mechanism that achieves *high* welfare, *low* envy, and *low* disproportionality. We demonstrate that for the canonical fair division settings the VCG mechanism is typically not a satisfactory candidate, but the redistribution mechanism of [Bailey, 1997; Cavallo, 2006] is.

1 Introduction

The starting point in designing or evaluating any prospective group decision-making procedure is the question: what goals do we want to achieve? The answer of course will depend on the setting and who you ask. If individuals are selfish, then each will answer "maximize the value I get from the procedure". But this is usually a non-starter because, very often, what is optimal for one individual will be suboptimal for another. A goal that has a much more plausible chance of being endorsed by individuals in a group, selfish though they may be, is to achieve some notion of *fairness*. In settings that have a certain symmetric separability in the description of each outcome, we can consider notions such as *envy* and *proportionality*. Would any agent prefer the outcome obtained by another agent? Does each agent get at least a certain proportion of the value they would obtain if they could make the decision themselves, as a dictator?

These are exactly the fairness goals that have been taken up and formally studied by researchers in mathematics, economics, political science, and, most recently, computer science. The prototypical decision setting addressed in such work is that of *fair division*, where either a divisible good must be split up—typically analogized as a cake to be cut—or a set of indivisible goods is to be allocated amongst a set of stakeholders.

Perhaps the most basic and well-known example of a fair division procedure is the "you cut I choose" method for two agents: one agent determines a bisection (cuts the cake), and the other decides who gets which piece. This simple approach achieves the desirable properties of envy-freeness and proportionality: neither agent would prefer to swap pieces with the other, and both agents—in their own estimation—obtain at least half of the cake. Indeed, if we make no further assumptions about the agents it is difficult to see any way of improving on this approach. Yet, from a broader perspective we can see that a crucial aspect of the problem has been ignored: how *much* does each agent like cake? What if one of the agents' enjoyment (call her Alice) is only marginally improved from obtaining anything more than a small sliver, while the other (call him Bob) obtains only marginally increasing enjoyment until he obtains a very large portion? In such a situation, intuitively we feel it would be more just to "tip the scale" in favor of Bob, since his gain could be enormous while Alice's loss would be negligible for a skewed division.

We can formalize this intuition as a concern for *social welfare*. However, as intuitively basic as the concept is, the way we've described the setting so far does not allow us to consider it—there is a problem of comparing one agent's welfare to another's.[1] When Bob claims to have lower value for the same size piece of cake as Alice, how do we interpret that? The comparison becomes possible if we assume a quasilinear structure to agent utilities, as an agent's value for an allocation can then be interpreted as their "willingness to pay" for it. We can then also bring to bear the powerful tool of monetary payments: besides receiving a piece of cake, each agent can either be given money or have money taken away. The social welfare can then neatly and legitimately be defined as the sum of the agent utilities.

As we will see, even granting this quasilinear context, in general there will exist no mechanism that perfectly satisfies all three of our criteria: efficiency (i.e., full social welfare, defined as the social utility of the allocation that maximizes the

[1] And so the best we could aim for is a Pareto optimal allocation where no agent could benefit from a change that doesn't cause another to lose.

sum of agent values), envy-freeness, and proportionality. In fact, there can be no mechanism that yields full social welfare alone, because any unsubsidized efficient allocation requires the agents to make payments outside of the group. At the same time, although previous work in cake-cutting demonstrates the existence of perfectly envy-free and proportional allocations for arbitrary size groups [Neyman, 1946], feasible methods for determining such allocations are currently known only for groups of size less than 5. But the fact that procedures that perfectly satisfy our criteria don't exist is little reason to abandon hope. Instead, in this paper we pursue methods that, in expectation, obtain "good" performance along each metric—*high* social welfare, *low* envy, and *low* disproportionality for each agent.

1.1 Related work

We build on two significant bodies of literature: the fair division literature which typically assumes little to nothing about the nature of agent utility functions, and the mechanism design literature which, with few exceptions, has at its foundation the assumption of transferable, quasilinear utility.

Work in fair division, at least in a modern research context, seems to have been initiated by Steinhaus [1948] and Banach & Knaster (whom Steinhaus credits as discovering one of the foundational constructive approaches), who addressed the question of proportionality for groups of size greater than two. More recently Brams has been a key figure, providing, with coauthors, a series of procedures for obtaining envy-free allocations for 3 or 4 players that involve a limited number of "cuts" to the cake (see the text [Brams and Taylor, 1996]).

Also recently, the question of *truthfulness* has been introduced in this context—can an agent gain from misrepresenting his preferences about pieces of cake? Brams et al. [2006] consider a very limited kind of truthfulness, requiring for each agent only that there exist a case (i.e., preferences of other agents) where lying would not be beneficial. Chen et al. [2010] consider a much stronger and more compelling notion, the standard concept of *strategyproofness* from the social choice literature, wherein lying can never be beneficial regardless of the behavior of others; they propose a procedure that is strategyproof and proportional for restricted classes of value functions; Mossel and Tamuz [2010] address essentially the same problem.

Fairness has also been studied in a context of allocating *indivisible* goods (the "assignment problem"); the canonical example is "room assignment, rent division", where a group of housemates must divvy up the rooms in the house and decide what share of the rent is paid by whom. Brams and Kilgour [2001], Haake et al. [2002], and Abdulkadiroglu et al. [2004] all introduce efficient procedures that (in some cases) also achieve envy-freeness; however all simply assume truthful participation and break down in a context of strategic agents. This is perhaps unsurprising, as Alkan et al. [1991] earlier showed that there exists no envy-free and strategyproof mechanism (that is, without allowing for "extra" payments that diminish social welfare). In a similar spirit to the evaluation methodology we propose in the current paper, Lipton et al. [2004] consider measures of envy, and seek allocation procedures that are approximately envy-free.

Mechanism design (initiated Hurwicz [1960]) introduces payments as a way to obtain good outcomes in equilibrium when agents are self-interested and strategic. The hallmark positive result is the class of Groves mechanisms, wherein each agent reports a value function over outcomes, the socially optimal one is chosen, and each agent is paid the reported value of the others minus a constant. Green and Laffont [1977] and Holmstrom [1979] showed that this class exactly characterizes the efficient and strategyproof mechanisms for most practical problem domains. In settings where no outcome yields anyone negative value, the Vickrey–Clarke–Groves (VCG) mechanism [Vickrey, 1961; Clarke, 1971; Groves, 1973]—an instance of the Groves class where agents make payments commensurate with the negative externality they impose on others—additionally has the properties of *ex post individual rationality* and *no-deficit*: no agent is ever worse off from participating and aggregate payment to the agents is never positive.

Despite these attributes, in a group decision-making problem where the goal is welfare of the group, the VCG mechanism is unsatisfactory because it generates high *revenue*, payments that must be transferred outside the group and thus detract from social welfare. *Redistribution mechanisms*, introduced by Bailey [1997] and Cavallo [2006],[2] address this issue by returning large portions of VCG revenue back to the agents in a way that does not violate strategyproofness. Subsequently Guo and Conitzer [2007] and Moulin [2009] provided a mechanism for the special case of multi-unit auctions that maximizes the *worst-case* social welfare in that context.

Studies of the fairness properties of strategyproof mechanisms has mainly been confined to VCG. Exceptions are [Papai, 2003], which characterizes the set of all envy-free Groves mechanisms (i.e., all strategyproof, efficient, and envy-free mechanisms); and [Moulin, 2010], which examines an efficiency/fairness tradeoff in single-item allocation. In the assignment problem setting, Leonard [1983] showed that VCG is envy-free; Cohen et al. [2010] recently extended this result to a generalization of the assignment problem where individuals have additive value for obtaining more than one good.

Finally, like the current paper, [Porter *et al.*, 2004] also straddles the fair division and mechanism design literatures, there seeking to equitably allocate costly tasks throughout a population (see also [Moulin, 2010]). Interestingly, for the case of single-item allocation the mechanism earlier introduced in [Bailey, 1997] and later generalized in [Cavallo, 2006] is proposed.

1.2 Summary of contributions

Our first step in this paper will be to generalize the notions of efficiency (welfare), envy-freeness, and proportionality from the strict "yes or no" conception to *degrees*. So, for instance, given a probability distribution over types a mechanism may yield social welfare that is close to opti-

[2] Bailey was the first, to my knowledge, to derive a redistribution mechanism; his approach applies to single-item auctions as well as some other settings. The mechanism of Cavallo [2006] coincides with Bailey's in those cases but is applicable to all decision scenarios, including important allocation domains to which Bailey's is not.

mal, be close to envy-free, and close to proportional for every agent in expectation. Next we will motivate our relaxation of a hard efficiency constraint by observing that no efficient mechanisms exist for canonical fair division settings, independent of fairness criteria. Finally we will demonstrate that the redistribution mechanism of [Bailey, 1997; Cavallo, 2006] performs exceedingly well on all three metrics in cake-cutting and assignment problems; this is in opposition to the simpler VCG mechanism, which, generally speaking, performs well on envy but not well with respect to welfare and proportionality.

1.3 Preliminaries

There is a set of agents $I = \{1, \ldots, n\}$ and a compact set of outcomes A (potentially infinite), where each $a \in A$ is an n-tuple $(a_1, a_2, \ldots, a_n)$ representing an allocation for each agent $i \in I$.There is a typespace Θ which represents the set of possible valuations for allocations. The joint typespace is Θ^n, and for any $\theta = (\theta_1, \ldots, \theta_n) \in \Theta^n$ and $a = (a_1, \ldots, a_n) \in A$, each agent i's *value* is $v_i(\theta_i, a_i)$. A mechanism is a tuple (f, T) where $f : \Theta^n \rightarrow A$ is a choice function and $T = (T_1, \ldots, T_n)$ defines a transfer function $T_i : \Theta^n \rightarrow \Re$ for each agent $i \in I$. In a mechanism agents report types, and then allocations and transfer payments are made according to f and T, respectively. We use notation $f_i(\theta)$ to denote a_i for the outcome a chosen by f given type profile θ (i.e., $f(\theta) = a = (f_1(\theta), \ldots, f_n(\theta)) = (a_1, \ldots, a_n)$ for some $a \in A$). We assume, for each $i \in I$, that i is self-interested and acts to maximize a *quasilinear* utility function u_i. Given mechanism (f, T), true joint type θ, and reported type $\hat{\theta}$, i then obtains utility: $v_i(\theta_i, f_i(\hat{\theta})) + T_i(\hat{\theta})$.We will specifically consider two classes of decision problems: cake-cutting and assignment.

Cake-cutting: There is a single infinitely divisible good to be allocated. The good may be heterogeneous, so values may depend not just on "how much" but also "which part" of the cake is received. Though our formal approach is completely general, in the evaluation section we will consider the following special classes of valuation functions:

- *Linear satiation*: value is homogeneous over all sections of the cake, and increases linearly with quantity, at slope determined by the agent's type, until plateauing at 1. If agent i with type θ_i receives $x\%$ of the cake, he obtains value: $v_i(\theta_i, x) = \min\{1, x\theta_i\}$. This captures different "satiation rates".
- *Exponential*: value is homogeneous over all sections of the cake; if allocated $x\%$ of the cake, an agent i with type θ_i obtains value $v_i(\theta_i, x) = 1 - e^{-x\theta_i}$.
- *Piecewise constant*: if K is the set of "kinds" of cake, each agent i's type has a component $\theta_{i,k}$ for every distinct kind $k \in K$. If, for each $k \in K$, agent i is allocated $x_k\%$ of the cake of kind k, he obtains value: $\sum_{k \in K} x_k \theta_{i,k}$.

Assignment: There are n agents and a heterogenous set of m items. Each agent's type determines a value for each item, and each agent can be allocated no more than one item.[3]

[3]Equivalently one can imagine that each agent's value for a bundle is restricted to equal the max of its values for any single item in the bundle, in which case an efficient allocation would not allocate multiple items to one agent.

2 Fairness metrics when utility is transferable

We generalize the either/or notions of efficiency, envy-freeness, and proportionality to "rates" that can be computed for any problem instance (defined by a joint type θ). Throughout the paper we assume a context of strategyproofness—we will only discuss the rates with respect to strategyproof mechanisms—so the measures are computed with respect to the truthful outcome.

Definition 1 (Welfare rate). *The ratio of the aggregate social welfare to the agents including payments, to the social value of the efficient allocation without payments. I.e., for mechanism (f, T) and joint type $\theta \in \Theta^n$:*

$$\frac{\sum_{i \in I}(v_i(\theta_i, f_i(\theta)) + T_i(\theta))}{\sum_{i \in I} v_i(\theta_i, f_i^*(\theta))} \tag{1}$$

For a *no-deficit* mechanism (one in which aggregate payments never exceed 0), the welfare rate is bounded above by 1. A mechanism that achieves *full social welfare* is one with a welfare rate of 1 for all $\theta \in \Theta^n$.

We now generalize the notions of envy-freeness and proportionality to "envy rate" and "disproportionality rate" representing the average extent throughout the population to which, respectively, an agent prefers the outcome for another agent, and an agent fails to obtain a "fair share" $1/n$ fraction of the utility he could obtain as a dictator. Both measures range between 0 and 1. In the spirit of fairness, the measures give equal weight to each agent's envy or disproportionality, in the sense that, e.g., the disproportionality measure for an agent who obtains only $\frac{\epsilon}{n} < \frac{1}{n}$ of his maximum possible utility u is the same whether u is minuscule or enormous.

Definition 2 (Envy rate). *Let u_{max} denote the utility an agent would have experienced if he received, maximizing over all agents j, j's allocation and j's payment. The envy rate equals, averaging over all agents, the difference between u_{max} and the agent's utility, divided by u_{max}. I.e., for mechanism (f, T) and joint type $\theta \in \Theta^n$:*

$$\frac{1}{n} \sum_{i \in I} \frac{\max_{j \in I}\{v_i(\theta_i, f_j(\theta)) + T_j(\theta)\} - \{v_i(\theta_i, f_i(\theta)) + T_i(\theta)\}}{\max_{j \in I}\{v_i(\theta_i, f_j(\theta)) + T_j(\theta)\}} \tag{2}$$

The envy rate never goes below 0 since each agent's actual allocation is included in the maximization. Envy-freeness is equivalent to the requirement that the envy rate be 0 for every problem instance.

Definition 3 (Disproportionality rate). *Averaging over all agents, the maximum of 0 and $1/n$ minus the ratio of an agent's allocation value plus payment to the value the agent would experience from obtaining his optimal allocation and no payment, divided by $1/n$. I.e., for mechanism (f, T) and joint type $\theta \in \Theta^n$:*

$$\frac{1}{n} \sum_{i \in I} \max\left\{0,\ \left(\frac{1}{n} - \frac{v_i(\theta_i, f_i(\theta)) + T_i(\theta)}{\max\limits_{a \in A} v_i(\theta_i, a_i)}\right) \Big/ \frac{1}{n}\right\} \tag{3}$$

The disproportionality rate is fixed to never be below 0 for any agent so that it penalizes the failure to meet traditional proportionality but does not reward a mechanism for going "above and beyond" proportionality for some agents; this is in the spirit of fairness. Traditional proportionality[4] is equivalent to the requirement that the disproportionality rate be 0 for every problem instance.

3 On the impossibility of full social welfare

In this section we consider the question of whether, even disregarding envy and proportionality considerations, a worst-case welfare rate of 1 ("full social welfare") can be achieved. In a setting where subsidies are not available, this is equivalent to the question of whether implementing a dominant strategy efficient choice function with a mechanism that is strongly budget-balanced (0 revenue, 0 deficit) is possible. To answer the question we must specify something about the problem setting, i.e., the typespace. Green and Laffont [1979] showed that for unrestricted values settings, no mechanism achieves full social welfare in dominant strategies. In the case of multi-unit auctions,[5] we can also deduce that no strategyproof mechanism achieves full social welfare by the results of Guo and Conitzer [2007] and Moulin [2009]: they (independently) derived the mechanism for that setting that has the worst-case welfare rate when values are positive but otherwise unrestricted, and that rate is lower than 1.

In an extended version of this paper we complement those results with a proof technique that allows us to consider arbitrary restricted settings, and apply a sufficient condition for the non-existence of mechanisms that achieve full social welfare. This theorem and proofs are omitted here due to space constraints. The result establishes that in any anonymous,[6] dominant strategy efficient, and strongly budget-balanced mechanism, for any two possible types θ, θ' in the typespace, letting SW_k be the social welfare that results when k agents have type θ and $n-k$ agents have type θ', a specific linear combination of $SW_0, SW_1, \ldots, SW_n$ must equal 0. This is only a necessary condition for the possibility of full welfare and far from a sufficient one, yet alone it is an extremely restrictive condition and can be applied to very directly show that full social welfare is impossible in settings including assignment and cake-cutting, even with highly restricted values.

Theorem 1. *For the assignment problem with any number of goods, if the agent value spaces are symmetric, smoothly connected, and include values 0 and x for each item, for some $x > 0$, there exists no anonymous, dominant strategy efficient, and strongly budget-balanced mechanism.*

Theorem 2. *For cake-cutting, if the typespace is symmetric, smoothly connected and admits linear satiation values with types in the range $[0, n-1]$ (where n is the number of agents), there exists no anonymous, dominant strategy efficient, and strongly budget-balanced mechanism.*

4 The redistribution mechanism

While full social welfare may be impossible, this of course does not preclude the existence of solutions that obtain very good social welfare, i.e., achieve a high welfare rate in expectation. The most well-known general social choice mechanism is VCG; but though VCG always achieves an outcome in dominant strategies that maximizes the sum of agent values, it requires that much of this value be transferred away from the group (high "revenue"). In fact, amongst all mechanisms that choose outcomes that maximize aggregate value, VCG requires the *maximum* transfer of that value outside of the group (see Theorem 2.10 of [Cavallo, 2008]).

In settings that are extremely lacking of structure, such as settings where each agent's value function over outcomes is completely unrestricted, no improvement over VCG is possible. However, in practically all allocation settings values have significant structure—for instance, in single-item allocation an agent obtains 0 value for any outcome in which he does not receive the item. Exploiting this structure to improve social welfare is the idea introduced, for restricted settings, by Bailey [1997], and for general settings, by Cavallo [2006].[7] The general *redistribution mechanism* (RM) proposed in [Cavallo, 2006] is as follows: implement VCG, then pay each agent i a quantity equal to $1/n$ times the minimum VCG revenue that would result independent of the agent's mode of participation. In the versions of the cake-cutting and assignment problems we examine here, the redistribution payment reduces to $1/n$ times the revenue that would result if the agent were not present.

To illustrate the mechanism, consider the 3-agent (i, j, k), 3-item (A, B, C) assignment problem depicted in Table 1, which one can think of as room assignment, rent division for the purpose of narrative.

	v_i	v_j	v_k
A	500	600	800
B	900	1000	900
C	600	900	600

Table 1: 3-agent, 3-item assignment problem example.

The optimal allocation is A to agent k, B to i, and C to j. Omitting the details of computation, under VCG i pays $\$100$, and neither j or k pay anything. Under RM i pays $\$66.67$, and j and k are each paid $\$33.33$. On this instance the welfare rate under VCG is $\frac{2500}{2600}$ and under RM it is 1. The envy and disproportionality rates for both mechanisms are 0 here. If this were a room assignment, rent division problem where the rent for the house is $\$1500$, starting with the equal-share payments of $\$500$ each to ensure no-deficit, under VCG agent i ends up paying $\$600$ and the other two agents

[4] The more basic idea of extending proportionality to a transferable utility context is not new; see, e.g., [Cramton *et al.*, 1987].

[5] The multi-unit auction setting is different from the assignment problem in that the goods are identical and so the problem can be described as simply choosing "who to serve" with an item.

[6] Anonymity requires that the expected utility obtained by two agents with the same type is the same, which is natural in the spirit of fairness.

[7] Unlike Cavallo's proposal, Bailey's mechanism is not feasible for cake-cutting unless we assume the type "no value for any amount of cake" is included in the typespace.

pay $\$500$ each—the surplus $\$100$ must be transferred outside of the group (e.g., to a charity that no agent obtains utility from giving to). Under RM i pays $\$566.67$, and the other two agents each pay $\$466.67$. In this fortuitous example there is no surplus; in general there may be a surplus, but under RM it is never greater (and is typically far less) than under VCG.

We will see in the next section that in both cake-cutting and assignment, VCG does well with respect to minimizing envy, but very poorly with respect to welfare and, typically, proportionality. RM typically does well in all three metrics. Though in some cases VCG achieves a lower envy rate, it is always dominated by RM in terms of welfare and proportionality.

Theorem 3. *On any problem instance, in any domain, RM has a weakly higher welfare rate and weakly lower disproportionality rate than VCG.*

In the case of assignment with a single good, it is particularly easy to compare the traditional binary fairness properties of VCG (which reduces to a Vickrey auction) and RM. RM reduces to the following simple form: the high bidder is allocated the good and pays the second highest bid, and every agent is paid $1/n$ times the second highest bid amongst the other agents.

Theorem 4. *In any single-item allocation problem instance, RM yields an outcome that is envy-free and proportional for at least $n-2$ agents. VCG yields an outcome that is envy-free for all agents but proportional for a maximum of 1 agent that has non-zero value for the item.*

5 Evaluation

In this section we evaluate VCG and RM along the metrics of welfare, envy, and disproportionality rates introduced in Section 2. We do an average case analysis, measuring the *expected value* of each rate given a probability distribution over agent values.[8] In cake-cutting,[9] we examine values drawn from the linear satiation class (with typespace $[0, n]$), the exponential class (with typespace $[0, 9]$), and the piecewise constant class (with 3 kinds of cake[10] and value space $[0, 1]$ for each kind). The results are given in Table 2. We report results for a type distribution that is uniform over the typespace (we also considered Gaussian type distributions, but the results were very similar and are thus omitted); in the case of piecewise constant values the typespace is multidimensional, and we considered values that are uniformly distributed and independent across different kinds of cake. In all three cases VCG performs poorly with respect to welfare and proportionality, but has a low envy rate. RM performs well along all three measures, notably with welfare going to 1 and envy and disproportionality to 0 as the population size grows.

[8]Expected values were computed by a Monte Carlo sampling method, with each data point averaged over 2000–10000 (depending on the setting) randomly drawn joint type instances.

[9]When utilities are a concave function of quantity allocated (as we consider here), optimal allocations can be computed with a greedy algorithm that allocates each incremental crumb to the agent whose marginal utility per crumb is currently highest.

[10]Variants with more or less kinds (heterogeneity) of cake were considered; results were very similar.

metric	n	VCG (a)	RM (a)	VCG (b)	RM (b)	VCG (c)	RM (c)
WR	3	0.566	0.728	0.719	0.825	0.333	0.778
	5	0.505	0.852	0.569	0.898	0.200	0.920
	10	0.459	0.936	0.417	0.956	0.100	0.980
	15	0.442	0.959	0.347	0.974	0.067	0.991
ER	3	0.032	0.116	0.041	0.041	0	0.011
	5	0.029	0.076	0.021	0.012	0	0.011
	10	0.018	0.026	0.006	0.002	0	0.007
	15	0.015	0.013	0.003	0.001	0	0.004
DR	3	0.361	0.171	0.126	0.041	0.532	0.050
	5	0.376	0.027	0.224	0.000	0.693	0.002
	10	0.373	0.000	0.355	0.000	0.835	0.000
	15	0.375	0.000	0.431	0.000	0.887	0.000

Table 2: **Cake-cutting**. Expected welfare (WR), envy (ER), and disproportionality (DR) rates under VCG and RM in three cake-cutting settings: (a) homogeneous, with values that rise linearly in quantity with slope equal to the agent's type, until reaching 1; (b) homogeneous, with values that equal $1 - e^{-x\theta_i}$ for an agent with type θ_i that receives $x\%$ of the cake; and (c) heterogeneous, with values linear in quantity of each kind of cake, with distinct slope for each kind.

In the assignment problem, each agent's type is represented as a vector of m values, one for each item. In our evaluation we take values drawn independently and uniformly over $[0, 1]$ for each item. We examined the following cases, with n the number of agents: n items; $n - 1$ items; and $n - 2$ items. The results are depicted in Table 3. Somewhat surprisingly, in the classical linear assignment problem (n agents, n items; Table 3 (a)) we find that VCG is a serviceable solution, obtaining a reasonably high welfare rate, zero envy, and a low disproportionality rate. Moving to RM improves the welfare rate at the cost of a marginal increase in the envy rate. In the case of $n-1$ items (Table 3 (b)), neither VCG nor RM achieve near-optimal performance: although RM's welfare rate is significantly better than VCG's, both are poor. When there are $n - 2$ goods (Table 3 (c)), VCG is poor while RM shines.

Finally we consider the case of assignment with one good, i.e., single-item allocation. In this case alone, there is another strategyproof mechanism in the literature to which we can compare VCG and RM: the worst-case optimal mechanism proposed by Guo and Conitzer [2007] and Moulin [2009] (we'll call it GCM). The mechanism has no concise form, and is instead specified by a system of equations that depends on the number of agents, so we refer the reader to the source papers for its description. As illustrated in Table 4, both RM and GCM perform superbly with respect to welfare and proportionality; VCG's welfare and disproportionality rates are abysmal, but it achieves no-envy, as in all assignment problems. The differences in performance between RM and GCM on welfare and disproportionality are negligible, but RM's expected envy rate is only about $1/3$ of GCM's.

6 Conclusion

In many group decision-making settings approaches that excel at meeting welfare or fairness criteria, but not both, will be unsatisfactory; broader evaluation metrics and different so-

metric	n	VCG (a)	RM (a)	VCG (b)	RM (b)	VCG (c)	RM (c)
WR	3	0.882	0.907	0.457	0.528	0.337	0.781
	5	0.864	0.915	0.372	0.491	0.281	0.833
	10	0.878	0.94	0.269	0.389	0.2	0.901
	15	0.895	0.955	0.211	0.318	0.16	0.932
ER	3	0	0.02	0	0.233	0	0.195
	5	0	0.021	0	0.171	0	0.1
	10	0	0.013	0	0.109	0	0.044
	15	0	0.009	0	0.082	0	0.026
DR	3	0.015	0.013	0.463	0.391	0.765	0.202
	5	0.001	0.001	0.31	0.183	0.532	0.007
	10	0.000	0.000	0.162	0.041	0.301	0.000
	15	0.000	0.000	0.111	0.012	0.208	0.000

(a) (b) (c)

Table 3: **Assignment**. Welfare (WR), envy (ER), and disproportionality (DR) rates under VCG and RM in the assignment problem with n agents and different numbers of items: (a) n items; (b) $n-1$ items; and (c) $n-2$ items.

metric	n	VCG	RM	GCM
welfare	3	0.334	0.774	0.774
	5	0.196	0.921	0.893
	10	0.1	0.98	0.991
	15	0.067	0.991	~ 1.0
envy	3	0	0.199	0.199
	5	0	0.056	0.126
	10	0	0.012	0.037
	15	0	0.005	0.015
disproportionality	3	0.764	0.207	0.207
	5	0.867	0.057	0.069
	10	0.935	0.012	0.011
	15	0.957	0.005	0.005

Table 4: **Single-item**. Welfare, envy, and disproportionality rates under VCG, RM, and GCM in single-item assignment.

lutions are called for. When utility is quasilinear in money, mechanisms using payments can be considered, allowing us to elicit truthful participation, formulate meaningful measures of both welfare and fairness, and even "redistribute" utility. If agents are strategic it is impossible to achieve *full* social welfare (efficient allocation with no aggregate payments outside the group), but the redistribution mechanism—pre-existing in the literature—comes close in the canonical fair division settings, particularly for larger groups of agents. At the same time, the redistribution mechanism approximates the traditional fairness criteria of envy-freeness and proportionality. This makes it a compelling solution for division of goods when utility is transferable and the objective is fairness, welfare, or achieving both simultaneously.

References

[Abdulkadiroglu *et al.*, 2004] Atila Abdulkadiroglu, Tayfun Snmez, and M. Utku nver. Room assignment-rent division: A market approach. *Social Choice and Welfare*, 22(3):515–530, 2004.

[Alkan *et al.*, 1991] Ahmet Alkan, Gabrielle Demange, and David Gale. Fair allocation of indivisible goods and criteria of justice. *Econometrica*, 59(4):1023–1039, 1991.

[Bailey, 1997] Martin J. Bailey. The demand revealing process: To distribute the surplus. *Public Choice*, 91:107–126, 1997.

[Brams and Kilgour, 2001] Steven J. Brams and D. M. Kilgour. Competitive fair division. *Journal of Political Economy*, 109:418–443, 2001.

[Brams and Taylor, 1996] Steven J. Brams and Alan D. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press, UK, 1996.

[Brams *et al.*, 2006] Steven J. Brams, Michael A. Jones, and Christian Klamler. Better ways to cut a cake. *Notices of the AMS*, 53(11):1314–1321, 2006.

[Cavallo, 2006] Ruggiero Cavallo. Optimal decision-making with minimal waste: Strategyproof redistribution of VCG payments. In *Proceedings of the 5th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'06)*, pages 882–889, 2006.

[Cavallo, 2008] Ruggiero Cavallo. *Social Welfare Maximization in Dynamic Strategic Decision Problems*. Ph.D. Thesis, Harvard University, 2008.

[Chen *et al.*, 2010] Yiling Chen, John Lai, David C. Parkes, and Ariel D. Procaccia. Truth, justice, and cake cutting. In *Proceedings of the 24th Annual Conference on Artificial Intelligence (AAAI-10)*, pages 56–61, 2010.

[Clarke, 1971] Edward Clarke. Multipart pricing of public goods. *Public Choice*, 8:19–33, 1971.

[Cohen *et al.*, 2010] Edith Cohen, Michal Feldman, Amos Fiat, Haim Kaplan, and Svetlana Olonetsky. Truth and envy in capacitated allocation games. unpublished, 2010.

[Cramton *et al.*, 1987] P. Cramton, R. Gibbons, and P. Klemperer. Dissolving a partnership efficiently. *Econometrica*, 55(3):615–632, 1987.

[Green and Laffont, 1977] Jerry Green and Jean-Jacques Laffont. Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica*, 45:427–438, 1977.

[Green and Laffont, 1979] Jerry R. Green and Jean-Jacques Laffont. *Incentives in public decision-making*. North Holland, New York, 1979.

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

[Guo and Conitzer, 2007] Mingyu Guo and Vincent Conitzer. Worst-case optimal redistribution of VCG payments. In *Proceedings of the 8th ACM Conference on Electronic Commerce (EC-07), San Diego, CA, USA*, pages 30–39, 2007.

[Haake *et al.*, 2002] Claus-Jochen Haake, Matthias G. Raith, and Francis Edward Su. Bidding for envy-freeness: A procedural approach to n-player fair-division problems. *Social Choice and Welfare*, 19(4):723–749, 2002.

[Holmstrom, 1979] Bengt Holmstrom. Groves' scheme on restricted domains. *Econometrica*, 47(5):1137–1144, 1979.

[Hurwicz, 1960] Leonid Hurwicz. Optimality and informational efficiency in resource allocation processes. In Karlin Arrow and Suppes, editors, *Mathematical Methods in the Social Sciences*. Stanford University Press, 1960.

[Leonard, 1983] Herman B. Leonard. Elicitation of honest preferences for the assignment of individuals to positions. *Journal of Political Economy*, 91(3):461–479, 1983.

[Lipton *et al.*, 2004] Richard Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM conference on Electronic commerce (EC-04)*, 2004.

[Mossel and Tamuz, 2010] Elchanan Mossel and Omer Tamuz. Truthful fair division. In *Proceedings of the 3rd International Symposium on Algorithmic Game Theory (SAGT-10)*, pages 288–299, 2010.

[Moulin, 2009] Hervé Moulin. Almost budget-balanced VCG mechanisms to assign multiple objects. *Journal of Economic Theory*, 144:96–119, 2009.

[Moulin, 2010] Hervé Moulin. Auctioning or assigning an object: some remarkable VCG mechanisms. *Social Choice and Welfare*, 34:193–216, 2010.

[Neyman, 1946] J. Neyman. Un théorème d'existence. *Centre Recherche Academie de Science Paris 222*, pages 843–845, 1946.

[Papai, 2003] S. Papai. Groves sealed bid auctions of heterogenous objects with fair prices. *Social Choice and Welfare*, 20(3):371–386, 2003.

[Porter *et al.*, 2004] R. Porter, Y. Shoham, and M. Tennenholtz. Fair imposition. *Journal of Economic Theory*, 118(2):209–228, Oct 2004.

[Steinhaus, 1948] H. Steinhaus. The problem of fair division. *Econometrica*, 16:101–104, 1948.

[Vickrey, 1961] William Vickrey. Counterspeculations, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.