

Fairness and Welfare Through Redistribution 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 that have historically had little contact: fair division, and efficient mechanism design with monetary payments. To do this we adopt the standard mechanism design paradigm where utility is assumed to be quasilinear and thus transferable across agents. We generalize the traditional binary criteria of envy-freeness, proportionality, and efficiency (welfare) to measures of degree that range between 0 and 1. We demonstrate that in the canonical fair division settings under any allocatively-efficient mechanism the worst-case welfare rate is 0 and disproportionality rate is 1; in other words, the worst-case results are as bad as possible. This strongly motivates an average-case analysis. We then set as the goal identification of a mechanism that achieves *high* welfare, *low* envy, and *low* disproportionality *in expectation* across a spectrum of fair division settings. We establish that the VCG mechanism is 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 to inquire: 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 all members of the group, selfish though they may be, is to achieve some notion of *fairness*. In decision settings that have a symmetric separability in the description of each outcome—for instance, splitting up a divisible good—we can consider notions such as *envy-freeness* and *proportionality*. Does no agent prefer the outcome (allocation) obtained by another agent? Does each agent get at least a certain proportion of the value he would obtain if he could make the decision himself, as a dictator?

These are exactly the goals that have been taken up and formally studied by researchers in mathematics, economics, political science, and, most recently, computer science working in this area. The prototypical decision setting

Copyright © 2012, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

addressed 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 (“assigned”) to 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 and a divisible good: 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 it is difficult to see any way of improving on this approach. Yet taking a broader perspective it is clear 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 may 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. 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 that agent utilities are quasilinear in money, 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 (utilitarian) 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 for the agents, defined as the aggregate value of the allocation that maximizes the sum of agent values), envy-freeness, and proportionality. In fact, there can be no mech-

anism 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 obtain "good" performance in expectation along each metric—high social welfare, low envy, and low disproportionality.

Related work

In this paper we build on two very significant bodies of literature: the fair division literature which typically assumes little to nothing about the nature of agent utility functions,¹ and the mechanism design literature which takes as 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.² More recently, Brams has been a key figure, with coauthors, providing 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]). There is currently no known procedure for achieving an envy-free allocation for more than 4 agents with a bounded number of cuts, although a procedure exists for arbitrary group size that involves an unbounded number of cuts [Brams and Taylor, 1995].

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, in fact the standard concept of *strategyproofness* from game theory, 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 ("piecewise uniform" and "piecewise linear"); 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 Abdulkadiroğlu et al. [2004] all introduce efficient procedures that (in some

¹And in particular, does not assume utilities have cardinal value that can be compared across agents.

²For two agents the "divide and choose" procedure goes back at least to biblical times; see Brams and Taylor [1999, p. 53], who cite the story of Abram and Lot (Gen 13:8–9).

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 imposing payments that diminish agents' 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 (see also [Mu'alem and Schapira, 2007]). Caragiannis et al. [2009] study the efficiency loss of imposing various fairness constraints in a setting where utility is not transferable, and assuming non-strategic agents.

Mechanism design (initiated by Hurwicz [1960] in the modern context) introduces payments as a way to obtain good outcomes in equilibrium when agents are self-interested and strategic. The hallmark positive result is the strategyproof and efficient Groves class of mechanisms. 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—is of great interest because it is *ex post individually rational* 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 members 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],³ address this issue by returning portions of VCG revenue back to the agents in a way that does not violate strategyproofness or no-deficit. Subsequent work by Guo and Conitzer [2009; 2008] and Moulin [2009] addressed multi-unit auctions, and provides a mechanism that maximizes the worst-case social welfare in that context.

Study 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 for assignment with super-additive valuations; and [Moulin, 2010]. In the assignment problem setting with unit-demand, Leonard [1983] showed that VCG is envy-free; very recently, Cohen et al. [2011] extended this result to a generalized setting where individuals have additive value for obtaining multiple goods, up to some capacity that is uniform across the population.

Finally, like the current paper, [Porter, Shoham, and Tenenholtz, 2004] also straddles the fair division and mechanism design literatures, there seeking to equitably allocate costly tasks throughout a population (see also [Moulin,

³Bailey 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 some cases but is applicable to all decision scenarios, including important domains to which Bailey's is not, such as combinatorial allocation.

2010]). Interestingly, for the case of single-task allocation the mechanism earlier proposed in [Bailey, 1997] and later generalized in [Cavallo, 2006] is proposed.

Summary of contributions

Our first contribution will be to provide a useful evaluation methodology: we generalize the notions of efficiency, envy-freeness, and proportionality from the strict “yes or no” conception to *degrees*. So, for instance, given a profile of agent types a mechanism may yield social welfare that is close to optimal, be close to envy-free, and close to proportional for every agent.

Next we will motivate our relaxation of hard efficiency and fairness constraints by establishing that, even taken individually, not only is *perfect* efficiency or proportionality unachievable in the canonical fair division settings, but in fact in any allocatively-efficient mechanism in the worst-case agents will end up with *zero* welfare and the outcome will have maximum *disproportionality*. In other words no efficient mechanism can guarantee any positive performance.

This moves us to seek a mechanism that performs well in the average-case. We demonstrate that a mechanism pre-existing in the literature—the redistribution mechanism of [Bailey, 1997] and [Cavallo, 2006]—performs exceedingly well on all three metrics in cake-cutting and assignment problems; this is in stark opposition to the simpler VCG mechanism which, we show, often performs well on envy but performs very poorly with respect to welfare and proportionality.

Preliminaries

There is a set of agents $I = \{1, \dots, n\}$ and a compact set of outcomes A , where each $a \in A$ is an n -tuple (a_1, a_2, \dots, a_n) representing an allocation for each agent $i \in I$.⁴ There is a typespace Θ , the same for each agent, which represents the set of possible private information an agent may have. The joint typespace is Θ^n , and for any $\theta = (\theta_1, \dots, \theta_n) \in \Theta^n$ and $a = (a_1, \dots, a_n) \in A$, each agent i 's *value* is $v_i(\theta_i, a_i)$. The value functions are symmetric, in that $\forall i, j \in I, \forall \theta_j \in \Theta, \forall a \in A, v_i(\theta_j, a_j) = v_j(\theta_j, a_j)$.

A mechanism is a tuple (f, T) where $f : \Theta^n \rightarrow A$ is a choice function and $T = (T_1, \dots, T_n)$, with transfer function $T_i : \Theta^n \rightarrow \mathbb{R}$ 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 . We use notation $f_i(\theta)$ to denote a_i for the outcome a chosen by f given type profile θ (i.e., if $f(\theta) = a$, then $f(\theta) = (f_1(\theta), \dots, f_n(\theta)) = (a_1, \dots, a_n)$). We assume that each $i \in I$ is self-interested and acts to maximize a *quasilinear* utility function; given mechanism (f, T) , true joint type θ , and reported type $\hat{\theta}$, i obtains utility $v_i(\theta_i, f_i(\hat{\theta})) + T_i(\hat{\theta})$. $v_i(\theta_i, a_i)$ can thus be considered i 's *willingness-to-pay* for allocation a_i .

We use shorthand $v(\theta, a)$ for $\sum_{i \in I} v_i(\theta_i, a_i)$, and $v_{-i}(\theta_{-i}, a)$ for $\sum_{j \in I \setminus \{i\}} v_j(\theta_j, a_j)$. We use notation f^* for an efficient choice function (i.e., $\forall \theta \in \Theta^n, f^*(\theta) \in$

$\arg \max_{a \in A} v(\theta, a)$) and $f^*(\theta_{-i})$ for an outcome that is efficient if the preferences of i are disregarded (i.e., $f^*(\theta_{-i}) \in \arg \max_{a \in A} v(\theta_{-i}, a_{-i})$). The VCG mechanism defines $T_i(\theta) = v_{-i}(\theta_{-i}, f^*(\theta)) - v_{-i}(\theta_{-i}, f^*(\theta_{-i}))$.

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 amongst the n agents. The good may be heterogeneous, so values may depend not just on *how much* but also *which part* of the cake is received. In various places we will make reference to 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 linear:* value is heterogeneous, where 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 kind k cake, he obtains value: $\sum_{k \in K} x_k \theta_{i,k}$.

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

2 Welfare and fairness evaluation metrics

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

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

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 $\theta \in \Theta^n$.⁵

We now introduce novel generalizations of envy-freeness and proportionality, i.e., rates representing the average extent throughout the population to which, respectively, an agent prefers the outcome for another agent, and an agent

⁴Note that A may be infinite, as in the case of allocating an infinitely divisible good.

⁵A welfare metric of this nature has been used in several previous papers on redistribution mechanisms.

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 $\frac{\epsilon}{n} < \frac{1}{n}$ times his maximum possible utility u is the same whether u is minuscule or enormous.

Definition 2 (Envy rate). Let u_i^{max} denote the utility an agent i 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 i , the difference between u_i^{max} and i ’s utility, divided by u_i^{max} . I.e., for mechanism (f, T) and joint type $\theta \in \Theta^n$:

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

(We take 0 for the above ratio in the case where numerator and denominator are 0.) 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 obtain from his optimal allocation and no payment, divided by $1/n$. I.e., for mechanism (f, T) and joint type $\theta \in \Theta^n$:

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

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 is equivalent to the requirement that the disproportionality rate be 0 for every problem instance.⁶

3 Negative worst-case results

Having established a way of evaluating any proposed mechanism along the three metrics of welfare, envy, and disproportionality, the most natural goal would be to seek a mechanism that *always* performs well in each of the dimensions. However in this section any such hopes are dashed: we demonstrate that in the *worst-case*, any strategyproof mechanism for cake-cutting (with a sufficiently broad typespace) or assignment that makes efficient allocations has welfare rate 0 and disproportionality rate 1. That is, even if we seek to optimize welfare or proportionality independently, in the worst-case we can not do better than the worst possible.

While one might question our emphasis on mechanisms that always choose an efficient outcome (*allocatively-efficient* mechanisms), this can be motivated by considering

⁶The more basic idea of extending proportionality to a transferable utility context is not new; see, e.g., [Cramton, Gibbons, and Klemperer, 1987].

that a non-Pareto-efficient outcome is inherently unstable in a transferable-utility setting: for instance in single-item allocation if the good is allocated to an agent j not equal to the highest valuer i , then j and i could both gain from arranging a sale after the mechanism has been run.

We first consider the worst-case outcomes of cake-cutting. Note that the theorem places no requirement that values are of the linear satiation (or any other) class; it requires only that particular linear satiation type values are *possible*. (Proofs are omitted due to space constraints, but we provide a sketch demonstration of the worst-case welfare and disproportionality results.)

Theorem 1. For $n \geq 2$ agents and any cake-cutting typespace that is smoothly connected and admits linear satiation values with slope 0 and $n-1$, any mechanism that is efficient in dominant strategies, ex post individually rational, and no-deficit has worst-case welfare rate 0 and disproportionality rate 1.

Proof Sketch. It can be shown that for any smoothly connected typespace that admits value 0 for all outcomes, in any mechanism that is truthful and efficient in dominant strategies, ex post individually rational, and no-deficit, no agent’s payment exceeds that which VCG prescribes by more than the amount of VCG revenue that would result—fixing the reports of others—if the agent reported value 0 for all outcomes. (See also Chapter 3 of [Cavallo, 2008].)

Now, for any $n \geq 2$ a cake-cutting example with linear satiation values can be constructed where: social value equals $n-1$; VCG revenue equals $n-1$ (so social utility is 0); and *zero* VCG revenue would result if any one given agent instead reported value 0 for all outcomes. By the result described above, no mechanism meeting the required criteria yields more agent welfare than VCG here, so none yields a welfare rate exceeding 0. Moreover, in the same example each agent’s optimal allocation value is positive and under any mechanism meeting the criteria each agent’s utility is 0, yielding a disproportionality rate of 1. \square

Using the same type of proof technique, a similar result can be established for the assignment problem:

Theorem 2. For the assignment problem with $n \geq 2$ agents and n items, if each agent’s value space is smoothly connected and admits values 0 and x for each item for some $x > 0$, any mechanism that is efficient in dominant strategies, ex post individually rational, and no-deficit has worst-case welfare rate 0 and disproportionality rate 1.

These results demonstrate the impossibility of designing an efficient mechanism that satisfies any worst-case robustness criteria at all with respect to welfare or disproportionality.⁷ To complete this line of inquiry we consider envy.

⁷The result of Theorem 2 can also be achieved for the $n-1$ items case but does *not* extend all the way down to the single-item case. [Guo and Conitzer, 2009] and [Moulin, 2009] provide the mechanism with the optimal worst-case welfare rate for single-item allocation, and the rate is greater than 0; furthermore, welfare rate greater than 0 necessarily implies disproportionality rate less than 1. An alternative derivation of the worst-case welfare rate for n agents could proceed fairly directly from results in these papers.

For assignment it is known that there exists an allocatively-efficient mechanism that is *completely* robust to envy: VCG is envy-free (see [Leonard, 1983; Cohen et al., 2011]). For cake-cutting we do not provide a tight worst-case result (envy will be less than 1 in allocatively-efficient outcomes, and the best worst-case bound will be domain dependent), but we do establish the following bound:

Theorem 3. *For any cake-cutting typespace that is smoothly connected and admits linear satiation values with slope 0 and n , any mechanism that is efficient in dominant strategies, ex post individually rational, and no-deficit has a worst-case envy rate of at least $\frac{n-1}{2n}$.*

4 Fairness of the redistribution mechanism

While the previous section demonstrates that in core fair division settings the *worst-case* welfare and disproportionality rates are as bad as possible, of course this does not preclude the existence of solutions that perform very well on average, i.e., achieve good rates *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 often involves much of this value being transferred away from the group (high “revenue”). In fact amongst all allocatively-efficient and ex post individually rational mechanisms, VCG specifies the *maximum* transfer of value outside of the group (see Theorem 2.10 of [Cavallo, 2008]), and in this sense is worst.

In settings that are extremely lacking of structure, such as that in which 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, an agent typically obtains 0 value for any outcome in which he does not receive an item. Exploiting this structure to improve welfare is the idea introduced, for restricted settings, by Bailey [1997], and in the general case, by Cavallo [2006]. The general *redistribution mechanism* (RM) proposed in [Cavallo, 2006] is as follows: implement VCG, then pay each agent $1/n$ times the minimum VCG revenue that could result considering all possible types the agent could report. (In some allocation settings, including those we focus on here, the Bailey and Cavallo proposals coincide.)

To illustrate the mechanism we consider a divisible good (cake-cutting) problem with 3 agents, where values for the good are exponential and as depicted in Figure 1. The efficient outcome (implemented by both VCG and RM) allocates 6% of the good to agent 1, 48% to agent 2, and 46% to agent 3, which yields value of 0.047, 0.619, and 0.746 to the respective agents. Under the VCG mechanism agents 1, 2, and 3 are then respectively paid: -0.044 , -0.317 , and -0.307 . Thus the welfare rate of VCG is:

$$\frac{0.047 + 0.619 + 0.746 - 0.044 - 0.317 - 0.307}{0.047 + 0.619 + 0.746} = 0.53$$

To calculate redistribution payments under RM for each agent we compute the revenue that VCG would have yielded if the agent had value 0 for any size piece of cake, fixing the types of the other agents, and divide this number by 3 (this

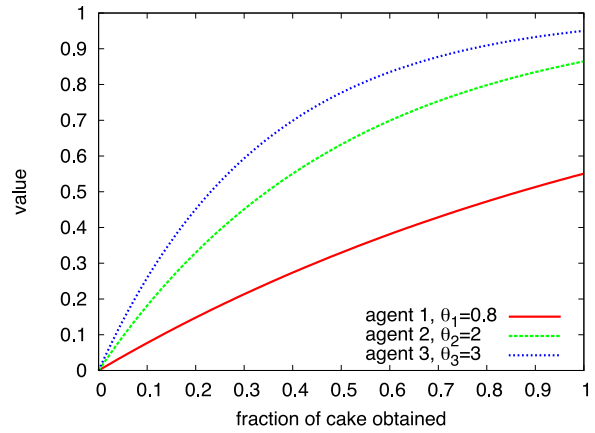


Figure 1: Example exponential value functions for three agents in a cake-cutting problem.

equals 0.135 for agent 1, 0.130 for agent 2, and 0.148 for agent 3). Thus under RM the agents are respectively paid 0.091, -0.187 , and -0.159 ; and the welfare rate is:

$$\frac{0.047 + 0.619 + 0.746 + 0.091 - 0.187 - 0.159}{0.047 + 0.619 + 0.746} = 0.82$$

The envy rate is 0.046 under VCG and 0.027 under RM. Both mechanisms are proportional in this example for agents 2 and 3, but VCG is highly disproportional for agent 1, while RM is much less so. This leads to a disproportionality rate of 0.328 for VCG and 0.083 for RM.

We will now establish some strong distribution-independent distinctions between VCG and RM. First a fact that follows trivially from the definition of RM:

Theorem 4. *On any problem instance, in any setting, RM has a weakly higher welfare rate and weakly lower disproportionality rate than VCG.*

Now in the case of assignment of a single good we can establish a particularly strong contrast between VCG (which reduces to a second-price auction) and RM with respect to the traditional binary fairness properties. RM reduces to the following simple form: the high bidder is allocated the good and pays the second highest bid, and every agent is then paid $1/n$ times the second highest bid amongst the other agents.

Theorem 5. *In any single-item assignment 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.*

Thus, at least in single-item assignment, RM is inherently fair. Moreover, the fairness performance it achieves is not simply a necessary implication of its welfare rate, which we demonstrate by example here. Consider the following alternative redistribution mechanism, which we’ll call “Mechanism X”: VCG is implemented, and then an arbitrary agent (say, agent 1) chosen a priori receives the minimum VCG revenue that could result taken over all possible types he could report, fixing the reports of others. Now consider a

single-item assignment example where agents 1, 2, 3, and 4 have values 10, 8, 8, and 6, respectively. Under RM agent 1 gets the item and pays 6, and the other agents all receive payment 2. Under Mechanism X agent 1 gets the item and pays 0, and all other agents receive nothing. Both mechanisms achieve a welfare rate of 1 on this example—in fact the *expected* welfare rate of both mechanisms is identical for any symmetric distribution over agent values—but RM’s envy rate is 0 and disproportionality rate is 0.15, whereas Mechanism X’s envy and disproportionality rates are both 0.75. The “even-handed” redistribution of RM gives it an exceptional tendency towards achieving fair outcomes.

5 Average-case evaluation

The results of Section 3 motivate the goal of finding a mechanism that performs well *in expectation* in a wide array of settings, and Section 4 provides some basic evidence that the redistribution mechanism may fit the bill. Here we verify this, evaluating VCG and RM along the metrics of welfare, envy, and disproportionality introduced in Section 2 in the average case given a probability distribution over types.

We start with cake-cutting and consider values drawn from the linear satiation class (with typespace $[0, n]$), the exponential class (with typespace $[0, 9]$), and the piecewise linear class (with 3 kinds of cake and value space $[0, 1]$ for each kind). We examined a type distribution that is uniform over the typespace; in the case of piecewise linear values the typespace is multidimensional, and we considered values that are uniformly distributed and independent across different kinds of cake. The results are given in Table 1.⁸ In all three cake-cutting settings VCG performs poorly with respect to welfare and proportionality, but has a low (but positive) envy rate. RM performs very well along all three measures, notably with welfare going to 1 and disproportionality to 0 as the population size grows.

We also evaluated the mechanisms in different versions of the assignment problem. Here each agent’s type is represented as a vector of m values, one for each item. In our evaluation we took values drawn independently and uniformly over $[0, 1]$ for each item. We present results for the following cases, with n the number of agents: n items; $n - 1$ items; and $n - 2$ items. The results are depicted in Table 2.

Somewhat surprisingly, in the n -item assignment problem (Table 2a) we find that VCG is a serviceable solution, obtaining a reasonably high welfare rate, zero envy, and a low disproportionality rate (surprising because in all other settings we considered, VCG welfare is low and disproportionality high). 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 2b), neither VCG nor RM achieve near-optimal performance: RM’s welfare rate is significantly better than VCG’s, but both are poor. When there are $n - 2$

⁸Expected values were computed via Monte Carlo sampling, with each data point averaged over 2000–10000 (depending on the setting) random joint type instances. For all cases we also examined Gaussian type distributions, and variants with more or less kinds (heterogeneity) of cake for the piecewise linear value setting; results were very similar so we only present the uniform case.

(a)				(b)		(c)	
metric	n	VCG	RM	VCG	RM	VCG	RM
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 1: **Cake-cutting.** Expected welfare (WR), envy (ER), and disproportionality (DR) rates under VCG and RM in three different cake-cutting value settings: (a) linear satiation, (b) exponential, and (c) piecewise linear.

(a)				(b)		(c)	
metric	n	VCG	RM	VCG	RM	VCG	RM
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

Table 2: **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.

goods (Table 2c), VCG is poor while RM shines.⁹

Finally we consider the case of assignment with one good, i.e., single-item allocation. In this case there exists another strategyproof mechanism in the literature to which we can compare VCG and RM: the worst-case optimal mechanism of [Guo and Conitzer, 2009] and [Moulin, 2009] (which we’ll label GCM), which we also implemented and tested. The mechanism’s description is complex, so for it we refer the reader to the source papers. 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. RM’s expected envy rate is only approximately 1/3 of GCM’s, though both are quite low.

⁹Due to space constraints we do not report results for all possible numbers of items; but informally, as the number decreases from $n - 2$ to 1, the results corroborate the pattern one would infer from the $n - 1$ (Table 2b), $n - 2$ (Table 2c), and 1-item (Table 3) data: VCG’s performance degrades while RM’s, if anything, ascends.

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 3: **Single-item allocation.** Welfare, envy, and disproportionality rates under VCG, RM, and GCM in the assignment problem with a single item.

6 Conclusion

In many group decision-making settings, solutions that excel only at achieving welfare or fairness—but not both—will be unsatisfactory; in this paper we addressed the simultaneous pursuit of both goals. We started by developing a suitable evaluation methodology that goes beyond the coarse binary efficiency and fairness concepts. Then we showed that with strategic agents it is impossible to achieve *any* welfare or proportionality in the worst-case. So we moved to seek a successful mechanism for the average-case and discovered that, at least for the settings under consideration, this did not require design of a new mechanism. We demonstrated that, unlike the VCG mechanism, for groups of more than a few agents the *redistribution mechanism* comes very close to perfect welfare and fairness across an array of canonical fair division settings. These attributes were previously unknown, and reveal it to be a compelling solution for group decision-making when utility is transferable and the objective is fairness, welfare, or achieving both simultaneously.

References

Abdulkadiroğlu, A.; Sönmez, T.; and Ünver, M. U. 2004. Room assignment-rent division: A market approach. *Social Choice and Welfare* 22(3):515–530.

Alkan, A.; Demange, G.; and Gale, D. 1991. Fair allocation of indivisible goods and criteria of justice. *Econometrica* 59(4):1023–1039.

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

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

Brams, S. J., and Taylor, A. D. 1995. An envy-free cake division protocol. *The American Mathematical Monthly* 102(1):9–18.

Brams, S. J., and Taylor, A. D. 1996. *Fair Division: From Cake-Cutting to Dispute Resolution*. UK: Cambridge University Press.

Brams, S. J., and Taylor, A. D. 1999. *The Win-Win Solution: Guaranteeing Fair Shares to Everybody*. New York: W. W. Norton.

Brams, S. J.; Jones, M. A.; and Klamler, C. 2006. Better ways to cut a cake. *Notices of the AMS* 53(11):1314–1321.

Caragiannis, I.; Kaklamanis, C.; Kanellopoulos, P.; and Kyropoulou, M. 2009. The efficiency of fair division. In *Proceedings of the 5th Workshop on Internet and Network Economics (WINE-09)*, LNCS 5929, Springer, 475–482.

Cavallo, R. 2006. 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)*, Hakodate, Japan, 882–889.

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

Chen, Y.; Lai, J.; Parkes, D. C.; and Procaccia, A. D. 2010. Truth, justice, and cake cutting. In *Proceedings of the 24th Annual Conference on Artificial Intelligence (AAAI-10)*, Atlanta, GA, 56–61.

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

Cohen, E.; Feldman, M.; Fiat, A.; Kaplan, H.; and Olonetsky, S. 2011. Truth, envy, and truthful market clearing bundle pricing. In *Proceedings of the 7th Workshop on Internet and Network Economics (WINE-11)*, Singapore, 97–108.

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

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

Guo, M., and Conitzer, V. 2008. Undominated redistribution mechanisms. In *Proceedings of the Seventh International Conference on Autonomous Agents and Multiagent Systems (AAMAS-08)*, Estoril, Portugal, 1039–1046.

Guo, M., and Conitzer, V. 2009. Worst-case optimal redistribution of VCG payments in multi-unit auctions. *Games and Economic Behavior* 67(1):69–98.

Haake, C.-J.; Raith, M. G.; and 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. 1960. Optimality and informational efficiency in resource allocation processes. In Arrow, K., and Suppes., eds., *Mathematical Methods in the Social Sciences*. Stanford University Press.

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

Lipton, R.; Markakis, E.; Mossel, E.; and Saberi, A. 2004. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC-04)*, New York, NY, 125–131.

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

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

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

Mu’alem, A., and Schapira, M. 2007. Setting lower bounds on truthfulness. In *Proceedings of the 18th ACM-SIAM Symposium on Discrete Algorithms (SODA-07)*, New Orleans, Louisiana, 1143–1152.

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

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

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

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

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