Sufficient Conditions for Weak Group-Strategy-Proofness

T.C.A. Madhav Raghavan*

31 July, 2014

Abstract

In this note we study group-strategy-proofness, which is the extension of strategy-proofness to groups of agents. This property comes in a standard form and a weak form. The distinction between the two forms is non-trivial as important rules in the literature fail the standard form but satisfy the weak form. It is well-known that in allocation models such as ours, a strategy-proof rule that is also non-bossy is (standard) group-strategy-proof. But the link between strategy-proofness and weak group-strategy-proofness is not as well established. We make steps towards this in this paper. We identify conditions (which we call ultra-weak Maskin monotonicity and weak non-bossiness) that are sufficient to ensure that a strategy-proof rule is weakly group-strategy-proof. These conditions are natural weaker forms of commonly used axioms in the literature. We also demonstrate that the conditions are 'weak enough', in that a rule satisfying them may not be (standard) group-strategy-proof.

1 Introduction

In this chapter we continue to consider a situation where heterogenous objects are to be distributed among a set of claimants. Same as before, the objects in question are indivisible, so they cannot be split or shared. There is no money in this economy so no compensation is possible.

Objects are to be assigned to individuals based on their preferences. Preferences in this context are rankings over the set of objects, one for each individual, such that an object preferred to another is also ranked higher than it. We assume that all objects can be ranked, that each object is acceptable to each individual, and that any ranking is possible in principle.

Moreover, the key feature of preferences is that this information is private to each individual. The designer or implementer of the solution to the allocation problems must elicit this information from individuals before making assignments. Individuals may reveal any preferences at all.

It is assumed that individuals may seek to game the system if it is to their advantage. If falsely revealing preferences gives an agent an object she prefers to what she might get if she instead truthfully revealed her preferences, then there is no reason to believe that a rational agent would not do so. A desirable property that a designer would like the allocation rule to satisfy is immunity from such undue gain for deviating agents. In particular, a strategy-proof rule ensures that it is a dominant strategy for every individual to truthfully reveal her preferences.[1]

There are many allocation problems where even individually strategy-proof rules do not exist. Demanding more from such rules is futile. However, in problems where strategy-proof rules exist, it is natural to ask whether the immunity from manipulation can be extended to coordinated actions by groups of agents as well. This property is called group-strategy-proofness.

*Senior Research Fellow, Indian Statistical Institute, New Delhi. Thesis advisor: Arunava Sen.

[1]This property has also been motivated by its informational simplicity. In a strategy-proof environment, an agent has to only consider her own preferences and not worry about what other agents might or might not do.

Group-strategy-proofness comes in two forms. Both look at the consequences for a group of agents who deviate by reporting different preferences. The standard form requires that it should not be the case that all deviating agents are as well off as before in terms of their original preferences, and some agent is strictly better off. A weaker form requires that it should not be the case that all deviating agents are strictly better off as a result. It is clear to see that the standard form implies the weak form while the converse is not true in general.

There is a clear connection between individual and group-strategy-proofness. In a wide class of allocation problems, group-strategy-proofness can be shown to be equivalent to a combination of strategy-proofness and a property called non-bossiness. Non-bossiness preempts situations where one agent can be bossy with another by affecting her assignment without changing her own.

Several important strategy-proof rules in the literature are also group-strategy-proof. The list includes inheritance rules (Pápai (2000)), other generalisations of top trading cycles rules (Abdulkadiroğlu and Sönmez (1999), Pycia and Ünver (2013)) and sequential and serial dictatorships (Svensson (1999), Pápai (2001), Hatfield (2009)). These rules are non-bossy and are also weakly group-strategy-proof by default.

However, there are also important strategy-proof rules that are weakly group-strategy-proof but not group-strategy-proof. The famous Gale-Shapley Deferred Acceptance (DA) rule (Gale and Shapley (1962)) is a case in point. Kojima (2010) shows that it is impossible for a stable rule to be non-bossy. Since the DA rule always produces a stable outcome, it is bossy. Thus it cannot be group-strategy-proof. Yet Hatfield and Kojima (2009) show that under general conditions (including the ones that apply in our model) the DA rule is weakly group-strategy-proof.

There is a non-trivial distinction between the two properties. It is useful therefore to ask the question: what makes a strategy-proof rule weakly group-strategy-proof but not group-strategy-proof? [2]

Non-bossiness is too strong a condition and is not necessary for weak group-strategy-proofness. What we seek in this paper is a condition (or set of conditions) that is weaker than non-bossiness that is nevertheless enough to guarantee weak group-strategy-proofness. We identify two fairly weak properties, which we call partial weak Maskin Monotonicity and weak non-bossiness respectively, that are sufficient to guarantee that a strategy-proof rule is also weakly group-strategy-proof.

We show the robustness of these conditions. That is, we can find examples of rules that violate one or other of these properties in turn. We also show that these properties are 'weak enough', in that there are rules that satisfy these properties (and are thus weakly group-strategy-proof) but are not group-strategy-proof.

A similar exercise is conducted in a recent working paper by Barberà et al. (2014) where they draw connections between strategy-proofness and weak group-strategy-proofness in different allocation environments. In particular, they show that it is the features of the models of house allocation, matching and division that result in the relationship between the two properties. Our work is independent of theirs. Their motivation is similar to ours but the formal results are different.

The paper is organised as follows. In Section 2 we provide the formal notation that we use throughout the paper. In Section 3 we formally describe the various forms of strategy-proofness and highlight some well-known connections between them. In Section 4 we discuss some of the important strategy-proof rules in the literature and why they satisfy one or other variant of group-strategy-proofness. In Section 5 we present the partial weak Maskin Monotonicity and weak non-bossiness properties and relate them to other conditions in the literature. Section 6 presents our main result and proof. In this section we also show that our conditions are independent. Section 7 concludes.

2 Notation and Definitions

The details of the model are given below:

- There is a finite set of *agents* $\mathcal{N} = \{1, ..., i, j, , k, ..., N\}$ and a finite set of *objects* $\mathcal{Z} = \{a, b, c, d, ...\}$.

[2] Barberà et al. (2010) consider a related question in the context of domains of preferences. In particular, they provide conditions on domains guaranteeing that for all rules defined on them, individual and weak group-strategy-proofness become equivalent.

- An *allocation* $x \in \mathcal{Z}^N$ with $x = (x_1, ..., x_N)$ is a vector that associates an object with each agent. For any agent $i \in \mathcal{N}$, $x_i \in \mathcal{Z}$ is the *assignment* of agent i in x.

- Depending on the context of the model, an allocation may be *feasible* or otherwise. The set of all feasible allocations in an allocation problem is given by $\mathcal{A}$.

- Preferences over assignments are strict. Formally, agent $i \in \mathcal{N}$ has *preferences*, denoted R_i, that are given by a binary relation over $\mathcal{Z}$. For any a, b, aR_ib is interpreted as 'object a is at least as good as object b for agent i under preferences R_i'. The binary relation is reflexive (for all a, aR_ia), complete (for all a, b, aR_ib or bR_ia), transitive (for all a, b, c, aR_ib and bR_ic imply aR_ic) and antisymmetric (for any a, b, aR_ib and bR_ia imply $a = b$). The associated strict relation is given by P_i, such that aP_ib if aR_ib and $a \neq b$. For any a, b, aP_ib means 'a is preferred by i to b under preferences R_i'.

- Agent preferences over allocations are selfish, in that they care only about the assignment they receive. Agents are indifferent between all allocations that give them the same assignment. An agent's preferences between two allocations that give her different assignments are governed by her preferences over the respective assignment she receives.

- A collection of preferences for all agents is called a *preference profile*, or simply a *profile*, and is denoted by $R = (R_1, ..., R_N)$. The set of all preference profiles is $\mathcal{R}$. In this model we shall usually suppress reference to $\mathcal{R}$, with the understanding that we operate on the full domain of preferences everywhere. As is the convention, we write R_{-i} for a sub-profile of preferences of all agents other than i. Similarly, for a subset of agents M, we write R_M and R_{-M} to denote the sub-profile of preferences of agents in subsets M and $\mathcal{N} \setminus M$, respectively.

- An *allocation rule* $f : \mathcal{R} \to \mathcal{A}$ takes as input a preference profile R and prescribes an associated feasible allocation $f(R)$. For any agent i, $f_i(R)$ is the assignment she receives at preference profile R according to the rule f. Similarly, for any subset of agents M, $f_M(R)$ is the M-dimensional vector of assignments of M at R, according to f.

3 Standard Strategy-Proofness Properties

In what follows we set up the properties of strategy-proofness and the two versions of group-strategy-proofness and draw some important connections between them.

Strategy-proofness is a condition which requires truth-telling to be a dominant strategy for all agents. In other words, given the reports of all other agents, an agent must be as well off reporting her true preferences as any other preferences. When this is true for all agents and all preferences, the mechanism is said to be strategy-proof. Formally:

Definition **1.** A rule is *strategy-proof* if there does not exist a profile R, an agent i, and preferences R'_i such that:

$$f_i(R'_i, R_{-i})P_if_i(R)$$

Group-strategy-proofness is a stronger condition than strategy-proofness. In its standard form, it ensures that groups of agents do not have profitable deviations, i.e., if a group of agents deviates by reporting different preferences, then a group-strategy-proof rule ensures that it is not the case that all agents in the deviating group are at least as well off as before, and some agent strictly better off. Formally:

Definition **2.** A rule f is *group-strategy-proof* if there does not exist a profile R, a subset of agents M, and a sub-profile R'_M such that

$$f_i(R'_M, R_{-M})R_if_i(R) \text{ for all } i \in M \text{ and } f_j(R'_M, R_{-M})P_jf_j(R) \text{ for some } j \in M$$

Weak group-strategy-proofness ensures that groups of agents do not have strictly profitable deviations, i.e., if a group of agents deviates by reporting different preferences, then a group-strategy-proof rule ensures that it is not the case that all agents in the deviating group are strictly better off than before. Formally:

Definition 3. A rule f is *weakly group-strategy-proof* if there does not exist a profile R, a subset of agents M, and a sub-profile R'_M such that:

$$f_i(R'_M, R_{-M})P_i f_i(R) \text{ for all } i \in M$$

It is clear that group-strategy-proofness implies weak group-strategy-proofness, which in turn implies strategy-proofness.

In the presence of non-bossiness, group-strategy-proofness is equivalent to strategy-proofness. Non-bossiness is an axiom that is pervasive in the literature on assignment rules. The condition was introduced by Satterthwaite and Sonnenschein (1981) and requires that an agent not be able to affect other agents' outcomes without affecting her own.

Definition 4. A rule f is *non-bossy* if, for all preference profiles R, for all agents i, and all reports R'_i, we have:

$$[f_i(R'_i, R_{-i}) = f_i(R)] \implies [f(R'_i, R_{-i}) = f(R)]$$

Non-bossiness negates any effect that an agent can have on other agents' assignments in cases where she does not change her own assignment. Its main justification is that it keeps the distribution of influence in the allocation process from unduly depending on any one agent. Another justification has to do with its strategic effects. Also, its original use by Satterthwaite and Sonnenschein (1981) is on the basis of considerations of informational simplicity. Non-bossiness disqualifies rules in exchange economies that "assign all the resources to one or the other of two agents depending upon some arbitrary feature of some third agent's preferences."[3]. However, Thomson (2014) also notes that its main value is in providing technical support for characterisation results.

As mentioned above, non-bossiness in conjunction with strategy-proofness is equivalent to group-strategy-proofness. This has been demonstrated in several contexts[4], so we state it here without proof.

Proposition 1. *A rule f is strongly group-strategy-proof if and only if it is strategy-proof and non-bossy.*

4 Strategy-Proof Rules and their Group-Strategy-Proofness

In this section we present some frequently used and important strategy-proof rules in the literature on allocation problems. We show by example how these rules satisfy (or fail to satisfy) the variants of group-strategy-proofness presented above.

4.1 Constant rules

A constant rule is one that prescribes the same allocation for every preference profile. It is easy to see that such rules are trivially strategy-proof and group-strategy-proof. A partially constant rule is one that is constant for some agents, i.e., assigns the same object to them at every preference profile. Such agents can never be part of a strictly profitable deviating group. If every group contains one such agent, this rule will be weakly group-strategy-proof.

However, these rules can still be bossy. Consider the following example: There are two agents $\{1, 2\}$ and three objects $\{a, b, c\}$. For any profile R, the rule assigns object a to agent 1 ($f_1(R) = a$) and to agent 2 assigns the top-ranked object in R_1 that is distinct from a ($f_2(R) = b$ if bP_1c and $f_2(R) = c$ if cP_1b). This is a partially constant rule as agent 1 gets the same object at all profiles. Yet agent 1 can be bossy with agent 2. Thus this rule is not group-strategy-proof. Yet it is weakly group-strategy-proof.

[3]See Thomson (2014)

[4]See, for example, Pápai (2000) and Chapter 2 of this thesis.

4.2 Dictatorships and their variants

The serial dictatorship (or serial priority) rule works as follows: There is an exogenous and fixed ordering of agents σ such that agents sequentially select objects in that order ($\sigma(1)$ selects first, $\sigma(2)$ goes next, and so on). Each agent selects her top-ranked objects from the ones that are available, given the choices of earlier agents in the sequence. It is easy to see that for any preference profile, the first agent always gets her top-ranked object, while the second agent always gets her top-ranked object whenever it is distinct from the selection of the first agent, and so on.

The sequential dictatorship (or sequential priority) rule differs from this rule only by making the identity of subsequent agents dependent on the assignment of earlier agents in the sequence. Thus for instance depending on what object $\sigma(1)$ selects, the identity of $\sigma(2)$ could differ.

Both these classes of rules are strategy-proof and group-strategy-proof in a wide variety of contexts (see Svensson (1999), Pápai (2000), Pápai (2001), Rhee (2011), Hatfield (2009) and Chapter 2 of this thesis for more details).

Consider, however, a variant of the serial dictatorship that we call a 'pure' dictatorship. There is an exogenous and fixed ordering of agents σ such that for any preference profile, the first agent in the sequence $\sigma(1)$ gets her top-ranked object. The next agent $\sigma(2)$ gets the object that $\sigma(1)$ ranks second, and in general agent $\sigma(k)$ gets the k^{th}-ranked object according to preferences of $\sigma(1)$.

This rule is strategy-proof. Since it is bossy, it is not group-strategy-proof. Yet it is easy to check that this rule is weakly group-strategy-proof.

4.3 The Top Trading Cycles Rule

The famous top trading cycles (TTC) rule is attributed to David Gale (see Shapley and Scarf (1974)). In brief, the TTC works as follows:

Each object is initially owned by one agent, who brings it to the market for trade[5]. Some agents may initially own more than one object, while others may own none at all. The procedure works in stages. In any stage, each agent who is yet to receive an assignment points to the owner of the object she most prefers from the ones that are available. A top trading cycle is made up of agents who successively point to the next agent, with the last agent pointing to the first. A cycle can be a singleton, such that an agent points to herself (she owns the object she most prefers.) Since there is a finite number of agents, at every stage there must always be a cycle. Agents in a cycle trade their objects along the cycle until they receive the object they desire. This becomes their assignment and such agents leave the market along with those objects. If there are still agents and objects left unassigned, the procedure repeats in the reduced market. If preferences are strict, then given an initial ownership, the resulting allocation is unique.

The TTC rule is illustrated by an example. Suppose there are three agents $(1, 2, 3)$ and three objects (a, b, c). Suppose agent 1 initially owns a, b and agent 3 initially owns c. Agent 1 desires c, while agents 2 and 3 desire b. The TTC procedure would look as follows:

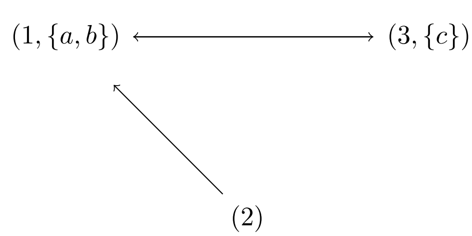

[5]Objects that an agent initially owns form a part of his or her 'endowment'

Agent 1 'points to' agent 3 who owns c, and agents 2 and 3 in turn point to agent 1 who owns b. The cycle in this stage is between agents 1 and 3, who consequently trade those objects. The TTC would assign b to agent 3 and c to agent 1.

TTC rules and their generalisations to inheritance rules (Pápai (2000)) are group-strategy-proof. They are thus also weakly group-strategy-proof. (An inheritance rule in the above example would also specify how agent 2 'inherits' the remaining object a. The TTC procedure in the second stage would just be agent 2 pointing to herself, and a would become her assignment.)

4.4 The Deferred Acceptance Rule

In addition to agent preferences, some allocation models assume that each object in turn has a 'priority', which is a ranking over agents or subsets of agents. A priority essentially captures the relative eligibility of an agent for an object vis-a-vis another agent.

The Deferred Acceptance (DA) rule (Gale and Shapley (1962)) uses preference and priority information and works in a series of rounds, as follows.

In the first round, all agents apply to their most preferred object. Any object that receives more applications than its maximum capacity is forced to reject the excess agents, provisionally accepting the rest. The agents that are rejected are those that are the lowest in that object's priority among its pool of applicants. In the next round, all rejected agents apply to their next preferred object that has not rejected them already. A object considers its existing applications plus any fresh ones it might receive, and provisionally accepts the top agents according to its priority, rejecting the lowest ones that are excess to capacity. If any agent is rejected, we go to the next round. The rule terminates in any round in which no agent is rejected. All provisional acceptances become final.

Consider the following example. There are three agents $\{1, 2, 3\}$ and three objects $\{a, b, c\}$. Preference and priority information are given in the table below:

Priorities			Preferences		
a	b	c	1	2	3
1	2	3	b	a	a
2	3	1	a	c	b
3	1	2	c	b	c

In the first round, agent 1 applies to her top-ranked object b while agents 2 and 3 apply to a. Since a has two applicants, it rejects the lower-ranked one according to its priority, which is agent 3. In the next round, agent 3 applies to her next preferred object, which is b. Now b has two applicants, so it rejects the lower-ranked one according to its priority, which is agent 1. In the next round, agent 1 applies to his next preferred object, which is a. Again, object a must reject one application, and so agent 2 is rejected. In the final round, agent 2 applies to c. There are no more rejections and the rule terminates here, giving us the final matching $((1, a), (2, c), (3, b))$.

The DA rule is strategy-proof. But it is bossy. To see this, consider a profile where only agent 2 changes her preferences such that she now ranks c above a and b. Evaluating the DA rule for this profile, we get that agent 2 continues to get c, as in the first profile, but now agents 1 and 3 swap a and b with each other. Thus despite not changing her own assignment, agent 2 affects the assignments of the other agents.

Thus the DA rule is not group-strategy-proof. It is, however, weakly group-strategy-proof. We will show that it satisfies our sufficient conditions for weak group-strategy-proofness.

5 Weaker Conditions on Allocation Rules

In this section we define our notions of partial weak Maskin monotonicity and weak non-bossiness which we will show are sufficient for a strategy-proof rule to be weakly group-strategy-proof. We first define some useful concepts.

The strict upper contour set of an object at a preference is the set of all objects that are preferred to it. Consider an agent i, a preference R_i and an object a. Formally, the *strict upper contour set of a at R_i* is given by $U(R_i, a) = \{b \in \mathcal{Z} : bP_i a\}$.

Let R be a preference profile and $f(R)$ the corresponding allocation produced by a rule f. For an agent i, we say that a preference R'_i is a *strict monotonic transformation of R_i at $f_i(R)$* if $U(R'_i, f_i(R)) \subset U(R_i, f_i(R))$. A preference is a strict monotonic transformation of another if the strict upper contour set of the assignment at the old preference is a strict superset of the strict upper contour set of that object in the new preference. We say that a preference R'_i is a *monotonic transformation of R_i at $f_i(R)$* if the subset relation in the above condition is weak. A particular kind of monotonic transformation is an upper-contour-set preserving transformation where, as the name suggests, the upper contour set remains the same in the new preference as well. In particular, a preference R'_i is an *upper-contour-set preserving transformation of R_i at $f_i(R)$* if $U(R'_i, f_i(R)) = U(R_i, f_i(R))$.

We can extend these concepts to profiles in a natural way. A profile R' is a *monotonic transformation of R at $f(R)$* if R'_i is a monotonic transformation of R_i at $f_i(R)$ for all $i \in \mathcal{N}$. We say that R' is a *strict monotonic transformation of R at $f(R)$* if in addition we have that R'_j is a strict monotonic transformation of R_j at $f_j(R)$ for some $j \in \mathcal{N}$. Similarly, a profile R' is an upper-contour-set preserving transformation of R at $f(R)$ if $U(R'_i, f_i(R)) = U(R_i, f_i(R))$ for all $i \in \mathcal{N}$.

We will use these concepts to define our two conditions: partial weak Maskin monotonicity and weak non-bossiness.

5.1 Partial Weak Maskin Monotonicity

A rule is Maskin monotonic if the allocation at any profile that is a monotonic transformation of another remains the same. Weak Maskin monotonicity relaxes the requirement that the allocations remain the same. A rule is weakly Maskin monotonic if every agent weakly prefers her assignment at a monotonic transformation of a profile to her assignment at that profile.

Formally, a rule satisfies *Maskin monotonicity* (Maskin (1999)) if for all R, R' such that R' is a monotonic transformation of R at $f(R)$, we have that $f(R') = f(R)$. A rule satisfies *weak Maskin monotonicity* if for all R, R' such that R' is a monotonic transformation of R at $f(R)$, we have that $f_i(R')R'_i f(R)$ for all $i \in \mathcal{N}$.

Weak Maskin monotonicity plays an important role in the characterisation of DA rules (see Kojima and Manea (2010)). It is also shown in that paper that the DA rule is the only stable rule at an exogenously specified priority profile that satisfies weak Maskin monotonicity.

In this paper we weaken the notion of Maskin monotonicity further. Partial weak Maskin monotonicity weakens the condition by requiring that only at least one agent weakly prefer her assignment at a strict monotonic transformation. Formally:

Definition **5.** A rule satisfies *partial weak Maskin monotonicity* if for all R, R' such that R' is a strict monotonic transformation of R at $f(R)$, we have that $f_i(R')R'_i f(R)$ for some $i \in \mathcal{N}$.

It is clear that a Maskin monotonic rule is also weakly Maskin monotonic and that a weakly Maskin monotonic rule is also partially weak Maskin monotonic. But the reverse implications do not hold in general.

5.2 Weak Non-Bossiness

Our other main property is weak non-bossiness. As its name suggests, this is a relaxation of the requirements of non-bossiness. Non-bossiness requires that assignments for all agents remain fixed for any deviation by any agent at any profile that does not change her assignment. Weak non-bossiness restricts the type of preference for which this invariance is true. In particular, if an agent's assignment does not change when she changes her preference via a upper-contour-set preserving transformation of her original preferences, then no other agent's assignment should change. Weak non-bossiness places no restrictions on bossy behaviour at other kinds of preference changes. Formally:

Definition **6.** A rule f is weakly non-bossy if, for any preference profile R, agent i and preferences R'_i such that R'_i is an upper-contour-set preserving transformation of R_i at $f_i(R)$, we have that:

$$[f_i(R'_i, R_{-i}) = f_i(R)] \implies [f(R'_i, R_{-i}) = f(R)]$$

Non-bossiness implies weak non-bossiness but the reverse implication does not hold in general.

6 Results

We are now ready to state our main theorem. Theorem 1 shows that the combination of strategy-proofness, weak non-bossiness and partial weak Maskin monotonicity is sufficient to give us weak group-strategy-proofness.

Theorem **1.** *If a strategy-proof rule f is weakly non-bossy and partially weak Maskin monotonic, then it is weakly group-strategy-proof.*

Proof: Let f be strategy-proof, weakly non-bossy and partially weak Maskin monotonic. Let R be a profile. Let M be a subset of agents, and R'_M be preferences for M such that $f_i(R'_M, R_{-M})R_i f_i(R)$ for all $i \in M$. Define the profile R' as (R'_M, R_{-M}).

For f to be weakly group-strategy-proof, we must show that $f_j(R') = f_j(R)$ for at least one agent $j \in M$. For contradiction, suppose that $f_i(R')P_i f_i(R)$ for all $i \in M$. Construct $\hat{R}$ such that $top(\hat{R}_i) = f_i(R')$ for all $i \in M$ and other objects are ranked the same as in R_i.

By assumption, $f_i(R')P_j f_i(R)$ for all $i \in M$. Thus by construction $\hat{R}$ is an upper-contour-set preserving transformation of R at $f(R)$. Consider the sequence of profiles $R^0 = R, R^1 = (\hat{R}_1, R_{-1}), R^2 = (\hat{R}_{\{1,2\}}, R_{-\{1,2\}}), ..., R^N = \hat{R}$. By strategy-proofness, $f_1(R^1) = f_1(R^0)$ since $\hat{R}_1$ is an upper-contour-set preserving transformation of R_1 for $f_1(R^0)$ and by weak non-bossiness, $f(R^1) = f(R^0)$. Repeating the argument for other agents and noting that for every i, R^i_i is an upper-contour-preserving transformation of R_i at $f_i(R^{i-1})$, we have by strategy-proofness that $f_i(R^i) = f_i(R^{i-1})$ and by weak non-bossiness that $f(R^i) = f(R^{i-1})$. Thus $f(\hat{R}) = f(R)$.

It is easy to see that $\hat{R}$ is also a monotonic transformation of R' at $f(R')$. If it is not a strict monotonic transformation, then it must be an upper-contour-set preserving transformation. By the arguments above, we have that $f(\hat{R}) = f(R')$. So let $\hat{R}$ be a strict monotonic transformation of R' at $f(R')$. Then by partial weak Maskin monotonicity, we have that $f_J(\hat{R})\hat{R}_J f_J(R')$ for some $J \in M$. In particular, we can find a $J \in M$ such that $f_J(\hat{R}) = top(\hat{R}_J)$.

But then $f_J(\hat{R}) = f_J(R')$, $f_J(\hat{R}) = f_J(R)$, but $f_J(R')P_j f_J(R)$ by assumption. Since preferences are strict, this is a contradiction. Thus our initial supposition was false, and there is at least one $i \in M$ such that $f_i(R') = f_i(R)$. Thus f is weakly group-strategy-proof. ■

6.1 Independence of Conditions

We will now show that our two properties are independent.

6.1.1 Partial Weak Maskin Monotonicity

Consider a situation with three agents $\mathcal{N} = \{1, 2, 3\}$ and four objects $\mathcal{Z} = \{a, b, c, d, e, h\}$. For any preference R_i and subset of objects X, let $bottom(R_i, X)$ be the last-ranked object in X according to R_i. Let the rule work as follows.

Let R be a preference profile. Then the rule is defined as follows: $f_2(R) = bottom(R_1, \mathcal{Z})$, $f_3(R) = bottom(R_2, \mathcal{Z} \setminus \{f_2(R)\})$, and $f_1(R) = bottom(R_3, \mathcal{Z} \setminus \{f_2(R), f_3(R)\})$.

Let R be a profile as given below. The allocation is given in boxes.

Preferences

R_1	R_2	R_3
d	h	a
a	d	b
c	b	e
e	c	h
h	e	d
$\boxed{b}$	$\boxed{a}$	$\boxed{c}$

This rule is strategy-proof. No agent can alter her own assignment via any change in preferences. Moreover, for any alternative report that keeps the same last-ranked object, the entire allocation remains the same. Thus the rule also satisfies weak non-bossiness. But the rule is not partially weak Maskin monotonic. To see this, consider the following profile R', where each preference R'_i is a strict monotonic transformation of R_i at $f_i(R)$.

Preferences

R'_1	R'_2	R'_3
d	h	a
a	d	b
c	b	e
e	c	h
b	a	c
$\boxed{h}$	$\boxed{e}$	$\boxed{d}$

Partial weak Maskin monotonicity requires that at least one of the agents gets an object weakly preferred to the original assignment at the new profile. But this is not the case here. Moreover, the rule is not weakly group-strategy-proof. In particular, agents $1, 2, 3$ can manipulate at R' via profile R making them all strictly better off according to R'.

6.1.2 Weak Non-bossiness

Let $\mathcal{N} = \{1, 2, 3\}$ be two agents and $\mathcal{Z} = \{a, b, c, d, e, h\}$ be a set of objects. For any preference R_i and set of objects X, let $bottom(R_i, X)$ denote the last-ranked object in X according to R_i. Let R be a profile and let the rule work as follows: $f_3(R) = h$. If bP_1a then $f_2(R) = d$. If aP_1b then $f_2(R) = e$. Also, $f_1(R) = bottom(R_2, \{a, b, c\})$.

Consider a preference profile R as in the table below. The allocation is marked in boxes.

Preferences

R_1	R_2	R_3
b	h	a
a	a	b
$\boxed{c}$	b	e
e	$\boxed{d}$	$\boxed{h}$
d	c	c
h	e	d

It is easy to see that this rule is strategy-proof. No agent can affect her own assignment via any other preference report. It is also partially weak Maskin monotonic as agent 3 gets object h for all preference profiles. However, it is not weakly non-bossy. In particular, let a profile R' be given as follows, where R'_1 is an upper-contour-set preserving transformation of R_1 at $f_1(R)$ and other preferences remain the same.

Preferences

R_1'	R_2'	R_3'
a	h	a
b	a	b
[c]	b	e
e	d	[h]
d	c	c
h	[e]	d

Even though agent 1 gets the same object as before, the allocation is no longer the same as agent 2 now gets e. This rule is also not weakly group-strategy-proof. Consider the following profile R'':

Preferences

R_1''	R_2''	R_3''
[b]	h	a
a	a	b
c	c	e
e	[d]	[h]
d	b	c
h	e	d

It is easy to see that agents $1, 2$ will manipulate at R' via (R_1'', R_2'').

6.2 The Rules

In what follows we will show that the rules we described earlier that are weakly group-strategy-proof but not group-strategy-proof do indeed satisfy our properties.

6.2.1 'Pure' Dictatorship

Recall the 'pure' dictatorship rule in this context. There is an exogenous and fixed ordering of agents σ such that for any preference profile, the first agent in the sequence $\sigma(1)$ gets her top-ranked object. The next agent $\sigma(2)$ gets the object that $\sigma(1)$ ranks second, and in general agent $\sigma(k)$ gets the k^{th}-ranked object according to preferences of $\sigma(1)$.

It is clear that this rule is strategy-proof. Note that the upper-contour-set for the first agent is empty. Also, that any reshuffling of the upper-contour-set for other agents does not affect any agent's assignment. Thus the rule is weakly non-bossy. Also, the first agent always gets her top-ranked object for any profile, and so the rule is partially weak Maskin monotonic.

6.2.2 Deferred Acceptance

The Deferred Acceptance rule is strategy-proof (Roth (1982), Dubins and Freedman (1984)). For any agent, and two objects she prefers to her assignment, reversing their order in the upper-contour-set will not affect the allocation. To see this, note that each of the two objects is assigned to an agent who is higher in that object's priority. Thus this agent will be rejected from both objects in either case. Thus the DA rule is weakly non-bossy. It is also partially weak Maskin monotonic. In particular, if an agent strictly raises her assignment in her preference and affects some other agent's assignment, it is via an improvement. The latter agent will be better off as a result. Thus at least one agent is at least as well off in the new profile as the old one.

7 Conclusion

The distinction between group-strategy-proofness and weak group-strategy-proofness is non-trivial as there are significant classes of rules that satisfy the weak property only. The connection between strategy-proofness and group-strategy-proofness have been well studied. In this chapter we contribute to this literature by identifying conditions that along with strategy-proofness are sufficient for weak group-strategy-proofness. These conditions, partial weak Maskin monotonicity and weak non-bossiness, are both weaker forms of well-known conditions in the literature. We also demonstrate that a rule satisfying all these properties need not be group-strategy-proof. Thus our conditions are 'weak enough' to fill the gap between the two versions of group-strategy-proofness.

A comment on extensions is in order. The conditions that we have identified are merely sufficient for weak group-strategy-proofness. Showing the necessity of these conditions is a difficult exercise. Firstly, note that weak group-strategy-proofness implies strategy-proofness. It is also possible to show that weak group-strategy-proofness implies partial weak Maskin monotonicity, but only when the number of agents that we consider is 2. Furthermore, weak non-bossiness is not a direct implication of weak group-strategy-proofness.

Finding suitable variants of these properties that are both necessary and sufficient to ensure that a strategy-proof rule is also weakly group-strategy-proof is therefore an open question.

References

Abdulkadiroğlu, A. and T. Sönmez (1999): "House allocation with existing tenants," *Journal of Economic Theory*, 88, 233–260.

Barberà, S., D. Berga, and B. Moreno (2010): "Individual versus group-strategy-proofness: When do they coincide?" *Journal of Economic Theory*, 145, 1648–1674.

——— (2014): "Group-strategy-proofness in private good economies without money: matching, division and house allocation," Mimeo.

Dubins, L. and D. Freedman (1984): "Machiavelli and the Gale-Shapley algorithm," *American Mathematical Monthly*, 88, 485–494.

Gale, D. and L. S. Shapley (1962): "College admissions and the stability of marriage," *American Mathematical Monthly*, 69, 9–15.

Hatfield, J. W. (2009): "Strategy-proof, efficient, and non-bossy quota allocations," *Social Choice and Welfare*, 33, 505–515.

Hatfield, J. W. and F. Kojima (2009): "Group incentive compatibility for matching with contracts," *Games and Economic Behavior*, 67, 745–749.

Kojima, F. (2010): "Impossibility of stable and non-bossy matching mechanisms," *Economics Letters*, 107, 69–70.

Kojima, F. and M. Manea (2010): "Axioms for Deferred Acceptance," *Econometrica*, 78, 633–653.

Maskin, E. (1999): "Nash Equilibrium and Welfare Optimality," *Review of Economic Studies*, 66, 23–38.

Pápai, S. (2000): "Strategyproof assignment by hierarchical exchange," *Econometrica*, 68, 1403–1433.

——— (2001): "Strategyproof and non-bossy multiple assignments," *Journal of Public Economic Theory*, 3, 257–271.

Pycia, M. and M. U. Ünver (2013): "Incentive-compatible allocation and exchange of discrete resources," Mimeo.

Rhee, S. (2011): "Strategy-proof allocation of individual goods among couples," *Japanese Economic Review*, 62, 289–303.

Roth, A. (1982): "Incentive compatibility in a market with indivisible goods," *Economics Letters*, 9, 127–132.

Satterthwaite, M. and H. Sonnenschein (1981): "Strategy-proof allocation mechanisms at differentiable points," *Review of Economic Studies*, 48, 587–597.

Shapley, L. and H. Scarf (1974): "On cores and indivisibility," *Journal of Mathematical Economics*, 1, 23–37.

Svensson, L.-G. (1999): "Strategy-proof allocation of indivisible goods," *Social Choice and Welfare*, 16, 557–567.

Thomson, W. (2014): "On the axiomatics of resource allocation: Interpreting nonbossiness," Mimeo, University of Rochester.