

Groupstrategyproofness of the Egalitarian Mechanism for Constrained Rationing Problems

Shyam Chandramouli* and Jay Sethuraman†

September 2011

Abstract

Motivated by applications in many economic environments, Bochet et al. [2] generalize the classic rationing model (see Sprumont [8]) as follows: there is a moneyless market, in which a non-storable, homogeneous commodity is reallocated between agents with single-peaked preferences. Agents are either suppliers or demanders. Transfers between a supplier and a demander are feasible only if they are *linked*, and the links form an arbitrary bipartite graph. Information about individual preferences is private, and so is information about feasible links: an agent may unilaterally close one of her links if it is in her interest to do so. For this problem they propose the *egalitarian transfer* solution, which equalizes the net transfers of rationed agents as much as permitted by the bilateral constraints. Furthermore, they show that the egalitarian mechanism elicits a truthful report of *both* preferences and links. In the variant where demanders are not strategic but demands need to be exactly met [1], they propose a similar mechanism for which truthfully reporting the peaks is a dominant strategy, but truthful reporting of links is not.

The key contribution of the paper is a comprehensive study of the egalitarian mechanism with respect to manipulation by a coalition of agents. Our main result is that the egalitarian mechanism is, in fact, *peak group strategyproof*: no coalition of agents can (weakly) benefit from jointly misreporting their peaks. Furthermore, we show that the egalitarian mechanism cannot be manipulated by any coalition of suppliers (or any coalition of demanders) in the model where both the suppliers and demanders are agents. Our proofs shed light on the structure of the two models and simplify some of the earlier proofs of strategyproofness. An implication of our results is that the well known algorithm of Megiddo [7] to compute a lexicographically optimal flow in a network is group strategyproof with respect to the source capacities and sink capacities.

*IEOR Department, Columbia University, New York, NY; sc3102@columbia.edu

†IEOR Department, Columbia University, New York, NY; jay@ieor.columbia.edu. Research supported by NSF grant CMMI-0916453.

1 Introduction

Motivated by applications in diverse settings, Bochet et al. [2, 1] study a model in which a homogeneous commodity is reallocated between a given set of agents with single-peaked preferences. In this environment, each agent is endowed with a certain quantity of the commodity and has an ideal consumption level (his *peak*) of that commodity. An agent who is endowed with more than his ideal consumption level can thus be thought of as a *supplier*, and an agent who is endowed with less than his ideal consumption level can be thought of as a *demand*er. Furthermore, transfers are possible only between certain pairs of agents, represented by a graph. The goal is to reallocate the commodity to balance supply and demand to the extent possible. The key difference from conventional economic models on this topic is the inability to use money: motivating applications include assigning (or reassigning) patients to hospitals, assigning students to schools, and allocating emergency aid supplies. On the other hand, it is easy to see that the resulting problem is essentially a transportation problem in a (bipartite) network. The distinguishing feature here is that the preferences of the agents (such as their peaks) and the other agents they are linked to is typically private information, so the agents must be given an incentive to report this information truthfully.

Bochet et al. [2] propose a clearinghouse mechanism (a centralized organization of the market) that prescribes an allocation that is efficient with respect to (reported) preferences and (reported) feasible links between agents. They identify a unique *egalitarian* allocation—so named because of the intimate connection with the egalitarian solution of an associated supermodular game—that Lorenz dominates all Pareto efficient allocations for this problem. Furthermore, they show that the egalitarian mechanism is strategyproof with respect to both links and peaks: no *individual* agent can strictly benefit by misreporting his peak or the set of agents he is linked to. In a companion paper, Bochet et al. [1] consider a “one-sided” model where the demanders are not strategic, and their demands have to be met exactly. For this model, they propose an egalitarian mechanism that is strategyproof with respect to peaks, but not with respect to links.

Our main result is that the egalitarian mechanism is *group* strategyproof with respect to peaks in both the one-sided and two-sided models of Bochet et al. Furthermore, we show that under the egalitarian mechanism it is a weakly dominant strategy for *any* coalition of suppliers (or any coalition of demanders) to truthfully report their links. These results thus properly generalize the corresponding (individual) strategyproofness results of Bochet et al. Our proofs result in an improved understanding of the two models and simplify some of the earlier proofs of strategyproofness.

The models of Bochet et al. [2, 1] generalize many well-known and well-understood models in the literature; we briefly discuss those that are very closely related to our work. If there is a single demander (or a single supplier), the problem reduces to a classical rationing problem of the sort considered by Sprumont [8]. The egalitarian rule then reduces to the “uniform” rule, and admits many characterizations. If the peaks are all identically 1, the problem reduces to a

matching problem with dichotomous preferences, discussed in Bogomolnaia and Moulin [3]: in this case, the flow between a supplier-demander pair can be thought of as the probability that this pair is matched. Some of the negative results related to link strategyproofness discussed later are true even in this restricted setting as has already been observed there; we mention these results in the appropriate sections for the sake of completeness. Finally, Megiddo [6, 7] considered the problem of finding an “optimal” flow in a multiple-source, multiple-sink network, and proposed an algorithm to find a lexicographically optimal flow. The egalitarian algorithm described in Bochet et al. [2, 1] is essentially Megiddo’s algorithm to compute a lexicographically optimal flow. An implication of our result is that Megiddo’s algorithm is group strategyproof with respect to the source and sink capacities, that is, if the agents are located on the edges incident to sources and sinks, and all other edge-capacities are common knowledge, then no coalition of agents have an incentive to misreport their capacities. This observation is useful in settings in which equitably sharing resources is important, such as the *sharing* problem of Brown [4].

2 Models and Notation

A single commodity is transferred from a set S of suppliers (typically indexed by i) to a set D of demanders (typically indexed by j). The commodity can only be transferred between certain supplier-demander pairs, given by a bipartite graph $G \subseteq S \times D$: $(i, j) \in G$ means that supplier i can send the commodity to demander j . A set of transfers from the suppliers to the demanders results in a vector $(x, y) \in \mathbb{R}_+^S \times \mathbb{R}_+^D$ where x_i (resp. y_j) is supplier i ’s (resp. demander j ’s) *net transfer*, with $\sum_S x_i = \sum_D y_j$. Each supplier i has *single-peaked preferences*¹ R_i (with corresponding indifference relation I_i) over her *net transfer* x_i , with peak s_i , and each demander j has single-peaked preferences R_j (I_j) over her net transfer y_j , with peak d_j . We write \mathcal{R} for the set of single peaked preferences over \mathbb{R}_+ , and $\mathcal{R}^{S \cup D}$ for the set of preference profiles.

We use the following notation. For any subset $T \subseteq S$, the set of demanders compatible with the suppliers in T is $f(T) = \{j \in D | G(T, \{j\}) \neq \emptyset\}$. Similarly, the set of suppliers compatible with the demanders in $C \subseteq D$ is $g(C) = \{i \in S | G(\{i\}, C) \neq \emptyset\}$. We abuse notation and say $f(i)$ and $g(j)$ instead of $f(\{i\})$ and $g(\{j\})$ respectively. For any subsets $T \subseteq S$, $C \subseteq D$, $x_T := \sum_{i \in T} x_i$ and $y_C := \sum_{j \in C} y_j$. A transfer of the commodity from S to D is realized by a flow φ , i.e., a vector $\varphi \in \mathbb{R}_+^G$. We write $x(\varphi), y(\varphi)$ for the transfers implemented by φ , namely:

$$\text{for all } i \in S : x_i(\varphi) = \sum_{j \in f(i)} \varphi_{ij}; \text{ for all } j \in D : y_j(\varphi) = \sum_{i \in g(j)} \varphi_{ij} \quad (1)$$

We say that the net transfers (x, y) are *feasible* if they are implemented by some flow. Given the graph G and the preferences of the agents (both suppliers and demanders), we would like to find feasible net transfers satisfying some desirable properties. An allocation rule (or a mechanism)

¹Writing P_i for agent i ’s strict preference, we have for every x_i, x_i' : $x_i < x_i' \leq s_i \Rightarrow x_i' P_i x_i$, and $s_i \leq x_i < x_i' \Rightarrow x_i P_i x_i'$.

is a function that associates feasible net transfers to any given problem. We shall be concerned exclusively with the egalitarian mechanism (described later) proposed by Bochet et al. [2, 1]. An important feature of the egalitarian mechanism is that it is *peak-only*: the egalitarian net transfers depend on the preferences of the agents *only* through the peaks. Thus it makes sense to talk of the *problem* (G, s, d) : this emphasizes the fact that the *peaks* of the agents and the *identity of their potential trading partners on the other side* can both be modeled as private information. To summarize: (i) agents (suppliers and demanders) report their peaks as well as the set of agents on the other side they are compatible with; (ii) the graph G has a link from supplier i to demander j if and only if both i and j report each other as compatible; and (iii) the egalitarian mechanism is applied to the problem (G, s, d) where s and d are the reported peaks of the suppliers and demanders respectively.

A mechanism is *link strategyproof* if for any profile of peaks (s, d) , it is a weakly dominant strategy for each agent to truthfully report their set of compatible partners. A mechanism is said to be *peak strategyproof* if for any graph G it is a weakly dominant strategy for an agent to truthfully report their peak. A mechanism is said to be *strategyproof* if it is a weakly dominant strategy for each agent to truthfully report their peak as well their set of compatible partners. It is not difficult to see that a mechanism is strategyproof if and only if it is both peak strategyproof and link strategyproof. These definitions admit a natural extension that models potential deviations by coalitions of agents: thus, a mechanism is *link group strategyproof* if for any profile of peaks (s, d) it is a weakly dominant strategy for any coalition of agents to truthfully report their set of compatible partners. Similarly, a mechanism is *peak group strategyproof* if for any graph G it is a weakly dominant strategy for any coalition of agents to truthfully report their peaks. (For truthful reporting to be a weakly dominant strategy, we require that for any group of agents, and for any possible misreports of these agents—including misreports in which only some but not all the agents in the group lie—it is not possible for every agent in the group to be weakly better off and at least one agent in the group to be strictly better off.

Finally, a mechanism is *group strategyproof* if it is a weakly dominant strategy for any coalition of agents to truthfully report both their peaks and compatible partners. Again, it is not hard to see that a mechanism is group strategyproof if it is both peak group strategyproof and link group strategyproof.

We shall focus on two related, but distinct, models. In the one-sided model, only the suppliers are modeled as strategic agents; the demanders are not strategic and their demands must be met exactly, so that some suppliers may be forced to send more than their peaks. In the two-sided model, both the suppliers and the demanders are modeled as strategic agents. Thus, in considering strategic issues in the one-sided model, we shall naturally only examine coalitions of suppliers.

3 The One-sided Model

3.1 Model

Recall that in the one-sided model, we are given a bipartite graph G with suppliers S indexed by i and demanders D indexed by j . Demander j has a demand of d_j that must be satisfied exactly, whereas supplier i has single-peaked preferences with peak s_i ; therefore, a supplier may be required to send more or less than his peak. In addition supplier i is required to send at least ℓ_i and at most u_i units of flow; we may assume without loss of generality that $\ell_i \leq s_i \leq u_i$. The peaks of the demanders, their preferences, and the ℓ_i and u_i are common knowledge; in contrast, for any supplier i , his peak s_i and the set $f(i)$ of demanders he is linked to may be private information held only by that supplier i and hence must be elicited by the mechanism.

Let $\lambda := (\lambda_i)_{i \in S}$ be non-negative. Construct the following network $G(\lambda)$: introduce a source s and a sink t ; arcs of the form (s, i) for each supplier i with capacity λ_i , arcs of the form (j, t) for each demander j with capacity d_j ; an infinite-capacity arc from supplier i to demander j if supplier i and demander j share a link. Let $\ell = (\ell_i)_{i \in S}$, $u = (u_i)_{i \in S}$, and $s = (s_i)_{i \in S}$. It is straightforward to verify that the given problem admits a feasible solution if and only if the maximum s - t flow in $G(\ell)$ and $G(u)$ are, respectively, $\sum_{i \in S} \ell_i$ and $\sum_{j \in D} d_j$. Consider now a maximum s - t flow in the network $G(s)$. By the max-flow min-cut theorem, there is a cut C (a cut is a subset of nodes that contains the source s but not the sink t) whose capacity is equal to that of the max-flow. If the set of suppliers in C is X and the set of demanders in C is Y , it is clear that $Y = f(X)$: if $Y \not\subseteq f(X)$, then C has infinite capacity, and if $Y \supset f(X)$ then C 's capacity can be improved by deleting the demanders in $Y \setminus f(X)$. Bochet et al. [1] show that in any Pareto-optimal allocation x for the suppliers, $x_i \leq s_i$ for each $i \in X$ and $x_i \geq s_i$ for each $i \in S \setminus X$.

If the min-cut is not unique, it is again well-known (see [5]) that there is a min-cut with the largest X (largest in the sense of inclusion), and a min-cut with the smallest X (again in the sense of inclusion). Call these sets \bar{X} and \underline{X} . It is easy to check that every supplier in $\bar{X} \setminus \underline{X}$ will be at his peak value in *all* Pareto optimal solutions. In the notation of Bochet et al. [1], $M_0 := \bar{X} \setminus \underline{X}$, $M_- := \underline{X}$, and $M_+ := S \setminus \bar{X}$. To keep things simple, however, we shall dispense with M_0 and use the partition $M_- = \underline{X}$, $M_+ = S \setminus \underline{X}$. In this case the partition of the demanders becomes $Q_+ = f(M_-)$ and $Q_- = D \setminus f(M_-)$. We note that our M_- is still uniquely determined for each problem. In what follows, often it will be important to talk about the set of suppliers involved in the cut, rather than the cut itself: we abuse notation and talk about the cut X when in fact the set of nodes in the cut is really $s \cup X \cup f(X)$.

3.2 Egalitarian Mechanism

Suppose $(x_i)_{i \in S}$ is a Pareto optimal allocation. From the earlier discussion it is clear that $x_i \in [s_i, u_i]$ for every supplier $i \in M_+$, and $x_i \in [\ell_i, s_i]$ for every supplier $i \in M_-$. Bochet et al. [1]

prove that the egalitarian allocation, which is defined independently for the suppliers in M_- and M_+ , Lorenz dominates all other Pareto optimal allocations.

For the suppliers in M_- , the egalitarian allocation is found by the following algorithm. Let λ be a parameter whose value is increased continuously from zero, and let $m_i(\lambda) = \text{median}(\ell_i, \lambda, s_i)$. Consider the graph $G(m(\lambda))$, where the capacity of the arc (s, i) is $m_i(\lambda)$. By the earlier discussion, we know that each supplier in M_- will send at least ℓ_i and at most s_i units of flow in a Pareto optimal solution, and that every demander j in $f(M_-)$ will receive exactly d_j units of flow. We now study the sequence of networks $G(m(\lambda))$ —specifically the maximum s - t flow in such networks—as λ is increased from zero. It is not hard to see that the maximum s - t flow in $G(m(\lambda))$ is a weakly-increasing, piecewise linear function of λ with at most $2n$ breakpoints. Moreover, each breakpoint is one of the ℓ_i , or one of the s_i (type 1), or is associated with a subset of suppliers X such that

$$\sum_{i \in X} m_i(\lambda) = \sum_{j \in f(X)} d_j \quad (2)$$

This we call a type-2 breakpoint. At a type-1 breakpoint, the associated supplier is at his peak and so will not send any more flow (recall that every supplier $i \in M_-$ will send flow at most his peak s_i); at a type-2 breakpoint, however, the group of suppliers in X are sending enough flow to satisfy the collective demand of the demanders in $f(X)$, so any further increase in flow from *any* supplier in X would cause *some* demander in $f(X)$ to accept more than his peak demand.

If the given problem does not have any type-2 breakpoint, then the egalitarian solution obtains by setting each supplier's allocation to his peak value. Otherwise, let λ^* be the first type-2 breakpoint of the max-flow function; by the max-flow min-cut theorem, for every subset X satisfying (2) at λ^* the cut $C^1 = \{s\} \cup X \cup f(X)$ is a minimal cut in $G(m(\lambda^*))$ providing a certificate of optimality for the maximum-flow in $G(m(\lambda^*))$. If there are several such cuts, we pick the one with the largest X^* (its existence is guaranteed by the usual supermodularity argument). The egalitarian solution obtains by setting

$$x_i = \text{median}(\ell_i, \lambda^*, s_i), \text{ for } i \in X^*, \quad y_j = d_j, \text{ for } j \in f(X^*),$$

and assigning to other agents their egalitarian share in the reduced problem involving the suppliers in $M_- \setminus X^*$ and the demanders in $Q_+ \setminus f(X^*)$. It is straightforward to verify that the first type-2 breakpoint λ^{**} of this reduced problem will satisfy $\lambda^{**} > \lambda^*$.

For the suppliers in M_+ , a similar algorithm is used to determine the egalitarian allocation: here, each demander $j \in Q_-$ receives exactly d_j units of flow, whereas every supplier $i \in M_+ = g(Q_-)$ sends at least s_i and at most u_i units of flow in a Pareto optimal solution. As before, we consider the graph $G(m(\lambda))$, where the capacity of the arc (s, i) is $m_i(\lambda) := \text{median}(s_i, \lambda, u_i)$. We increase λ gradually and observe that the maximum s - t flow in $G(m(\lambda))$ is a weakly-increasing, piecewise linear function of λ with at most $2n$ breakpoints. Moreover, each breakpoint is one of

the s_i , or one of the u_i (type 1), or is associated with a subset of suppliers X such that

$$\sum_{i \in X} m_i(\lambda) = \sum_{j \in f(X)} d_j$$

This we call a type-2 breakpoint. At a type-1 breakpoint, the associated supplier is at his upper-bound and so cannot send any more flow; at a type-2 breakpoint, however, the group of suppliers in X are sending enough just enough flow to satisfy the collective demand of the demanders in $f(X)$, so any decrease in flow from *any* supplier in X would cause *some* demander in $f(X)$ to receive an amount strictly below his peak demand. As before, if the given problem does not have any type-2 breakpoint, then the egalitarian solution obtains by setting each supplier's allocation to his upper bound. Otherwise, let λ^* be the first type-2 breakpoint of the max-flow function, and let X^* be the (largest) associated bottleneck set of suppliers (as before). The egalitarian solution obtains by setting

$$x_i = \text{median}(s_i, \lambda^*, u_i), \text{ for } i \in X^*, \quad y_j = d_j, \text{ for } j \in f(X^*),$$

and assigning to other agents their egalitarian share in the reduced problem involving the suppliers in $M_+ \setminus X^*$ and the demanders in $Q_- \setminus f(X^*)$. This completely defines the egalitarian solution.

3.3 Group Strategyproofness

We turn now to strategic aspects of the rationing problem with constraints. In the one-sided model, only the suppliers are modeled as “agents,” who possess potentially two pieces of information that could be modeled as private: the set of demanders they are compatible with, and their own preference over allocations². As the egalitarian mechanism is “peak-only” [1], it is sufficient for the suppliers to report only their peaks, rather than their entire preference ordering.

It is a simple matter to verify that the egalitarian mechanism is *not* link strategyproof. Consider a supplier with a peak of 1, connected to two demanders, each with a demand of 1, see Figure 1. If the supplier reveals both links, his egalitarian allocation is 2, whereas by suppressing one of the links, his egalitarian allocation improves to 1. Therefore in the rest of this section we focus only on peak strategyproofness.

Bochet et al. [1] show that the egalitarian mechanism, which finds the egalitarian allocation for any given problem, is peak strategyproof. Our main result in this section is that, in fact, the egalitarian mechanism is peak groupstrategyproof. To set the stage for this, we start with a lemma, which we first use to give an alternative proof that the egalitarian mechanism is strategyproof.

Lemma 1 *For a problem (G, s, d) , suppose the decomposition is M_+ and M_- (with Q_+ , Q_- defined as before), and the egalitarian allocation is x . Consider the problem (G, s', d) with $s'_j = s_j$ for all $j \neq i$, with the decomposition being M'_+ and M'_- .*

²The set of demanders, their individual demands, as well as the the lower and upper bounds on arc-flows are assumed to be common knowledge.




Figure 1: Counterexample for Link SP

(a) If $i \in M_-$ and $s'_i \geq s_i$, $M'_+ = M_+$ and $M'_- = M_-$.

(b) If $i \in M_+$ and $s'_i \leq s_i$, $M'_+ = M_+$ and $M'_- = M_-$.

Proof. By definition, M_- is the smallest (both in terms of cardinality and inclusion) min-cut in the graph $G(s)$ (see §3.1 for the definition). For $i \in M_-$, the arc (s, i) does not contribute to the cut-capacity. If $s'_i \geq s_i$, the capacity of any cut is weakly greater in (G, s', d) than in (G, s, d) , whereas the capacity of the cut M_- stays the same, so part (a) follows by the minimality of M_- . Similarly, for $i \in M_+$, the arc (s, i) contributes to cut-capacity, the capacity of the cut M_- is smaller in (G, s', d) than in (G, s, d) by *exactly* $s_i - s'_i$, whereas the capacity of any cut is weakly smaller in (G, s', d) than in (G, s, d) by *at most* $s_i - s'_i$. Again, part (b) follows by the minimality of M_- . ■

Theorem 1 *The egalitarian mechanism is peak strategyproof.*

Proof. For the problem (G, s, d) let x be the egalitarian allocation, and let M_+ and M_- be defined as before. Consider the problem (G, s', d) with $s_k = s'_k$ for all $k \neq i$. Suppose $i \in M_-$. If $s'_i \geq s_i$, Lemma 1 proves that the decomposition does not change; it is easy to see that the egalitarian allocation is unaffected as well, because the algorithm to compute operates identically in the problems (G, s, d) and (G, s', d) . Similarly, if $i \in M_+$ and $s'_i \leq s_i$, the decomposition does not change (by Lemma 1), and the egalitarian allocation is unaffected as well. Suppose agent i reports s'_i as his peak and the allocation changes to x'_i . To prove strategyproofness, it suffices to show that any $i \in M_-$ (weakly) prefers x_i to x'_i for all $s'_i < s_i$, and that any $i \in M_+$ (weakly) prefers x_i to x'_i for all $s'_i > s_i$.

Fix an $i \in M_-$, and suppose that i reports a peak of $s'_i < s_i$. In this case the decomposition *may* change; let M'_- and M'_+ be the new decomposition. If $i \in M'_-$, an application of Lemma 1 to the problem (G, s', d) shows that the decomposition does not change, and that $x'_i = x_i$. Suppose $i \in M'_+$. Let $D' := Q_+ \cap Q'_-$, and $X' := M_- \cap M'_+$, and note that by our supposition $X' \ni i$. Note

also that $g(D') \cap M_- \subseteq X'$, as no agent in M'_- has a link to any demander in Q'_- . Furthermore, if $i \notin g(D')$, $x'_i = 0$, and again the result follows: recall that $f(i) \subseteq Q_+$; and if $i \notin g(D')$, $f(i) \subseteq Q'_+$, and the links from M'_+ to Q'_+ do not carry any flow. So we may assume that $i \in g(D')$. We now make two simple observations about the agents in $X' \cap g(D')$ in the problem (G, s', d) : first every such agent sends flow only to the demanders in D' , and therefore $\sum_{k \in X' \cap g(D')} x'_k \leq \sum_{j \in D'} d_j$. Also, as every agent in $X' \cap g(D')$ is (weakly) above his reported peak, $x'_k \geq s_k$ for every $k \in X' \cap g(D')$, $k \neq i$, and $x'_i \geq s'_i$. This implies

$$\sum_{k \in X' \cap g(D'), k \neq i} s_k + x'_i \leq \sum_{j \in D'} d_j. \quad (3)$$

We next claim that in the problem (G, s, d) , $\sum_{k \in X' \cap g(D')} x_k = \sum_{j \in D'} d_j$. To see why, observe that the demands of D' are covered in the problem (G, s, d) by the suppliers in $M_- \cap g(D')$; but every demander in D' moves from Q_+ to Q'_- , so every supplier in $M_- \cap g(D')$ must move to M'_+ (as there cannot be an edge between a supplier in M'_- and a demander in Q'_-). This implies that any supplier supplying a positive amount to a demander in D' in the problem (G, s, d) must be in $X' \cap g(D')$. Note also that for each $k \in X' \cap g(D')$, $x_k \leq s_k$. These, along with $X' \cap g(D') \subseteq M_-$, imply

$$\sum_{k \in X' \cap g(D'), k \neq i} s_k + x_i \geq \sum_{j \in D'} d_j. \quad (4)$$

Inequalities (3) and (4) imply $x'_i \leq x_i$, as required.

Now fix an $i \in M_+$, and suppose that i reports a peak of $s'_i < s_i$. In this case the decomposition may change; let M'_- and M'_+ be the new decomposition. If $i \in M'_+$, as before, an application of Lemma 1 to the problem (G, s', d) shows that the decomposition does not change, and that $x'_i = x_i$. Suppose $i \in M'_-$. Let $D' := Q_- \cap Q'_+$, and $X' := M_+ \cap M'_-$, and note that by our supposition $X' \ni i$. Note also that $f(X') \cap Q_- \subseteq D'$, as no agent in M'_- can have a link to any demander in Q'_+ . We now make two simple observations about the demanders in $f(X') \cap D'$ in the problem (G, s', d) : first every such demander can receive flow only from the agents in X' , and therefore $\sum_{k \in X'} x'_k \geq \sum_{j \in f(X') \cap D'} d_j$. Also, as every agent in X' is (weakly) below his reported peak (in the new problem), $x'_k \leq s_k$ for every $k \in X'$, $k \neq i$, and $x'_i \leq s'_i$. This implies

$$\sum_{k \in X', k \neq i} s_k + x'_i \geq \sum_{j \in f(X') \cap D'} d_j. \quad (5)$$

We next claim that in the problem (G, s, d) , $\sum_{k \in X'} x_k = \sum_{j \in f(X') \cap D'} d_j$: in (G, s, d) the suppliers in X' send flow only to the demanders in $f(X') \cap D'$, who receive flow only from these suppliers. Furthermore, $x_k \geq s_k$ for each $k \in X'$. In particular,

$$\sum_{k \in X', k \neq i} s_k + x_i \leq \sum_{j \in f(X') \cap D'} d_j. \quad (6)$$

Inequalities (5) and (6) imply $x'_i \geq x_i$, as required. ■

In fact, the ideas in the proof of Theorem 1 can be used to prove the following result, which weakens the conditions under which the decomposition is guaranteed not to change.

Lemma 2 *For a problem (G, s, d) , suppose the decomposition is M_+ and M_- (with Q_+ , Q_- defined as before), and the egalitarian allocation is x . Consider the problem (G, s', d) with $s'_j = s_j$ for all $j \neq i$, with the decomposition being M'_+ and M'_- .*

(a) *If $i \in M_-$ and $s'_i > x_i$, $M'_+ = M_+$ and $M'_- = M_-$.*

(b) *If $i \in M_+$ and $s'_i < x_i$, $M'_+ = M_+$ and $M'_- = M_-$.*

Proof. By definition, M_- is the smallest (both in terms of cardinality and inclusion) min-cut in the graph $G(s)$ (see §3.1 for the definition). For $i \in M_-$, the arc (s, i) does not contribute to the cut-capacity. If $s'_i \geq s_i$, the capacity of any cut is weakly greater in (G, s', d) than in (G, s, d) , whereas the capacity of the cut M_- stays the same, so the result follows. Suppose now that $x_i < s'_i < s_i$, the max s - t flow in $G(s')$ is weakly below that of $G(s)$, but the egalitarian allocation x is still feasible, so x continues to be a max-flow, so M_- continues to be a min-cut in $G(s')$. We need to show that it remains the minimal min-cut. First observe that $M'_- \subseteq M_-$, as M'_- is the minimal min-cut in $G(s')$ whereas M_- is a min-cut for $G(s')$. If $i \in M'_-$, then the capacity of the cut M'_- is the same in $G(s)$ and $G(s')$, so the minimality of M_- in the problem (G, s, d) implies $M'_- = M_-$. Suppose $i \notin M'_-$. Let $X = M_- \setminus M'_-$, and note that $i \in X$. Note also that $Q_+ = f(M_-)$ and $Q'_+ = f(M'_-)$, so that the net change in the cut capacity when the suppliers in X move from M_- to M'_+ is precisely $\sum_{k \in X} s'_k - \sum_{j \in Q_+ \setminus Q'_+} d_j$. In the problem (G, s, d) , however, the demanders in $Q_+ \setminus Q'_+$ receive flow only from the suppliers in X , each of whom sends no more than his peak: thus, $\sum_{k \in X} x_k \geq \sum_{j \in Q_+ \setminus Q'_+} d_j$, and $s_k \geq x_k$ for each k . An easy implication is that $s'_k \geq x_k$ for each $k \in X$, $k \neq i$, and $s'_i > x_i$. Thus the net change in cut capacity in moving from M_- to M'_- is strictly positive, which implies M'_- cannot be a min-cut. A similar argument establishes part (b). ■

We conclude this section with a proof that the egalitarian mechanism is, in fact, group strategyproof.

Theorem 2 *The egalitarian mechanism is peak group strategyproof.*

Proof. Suppose not. Focus on a counterexample G with the *smallest* number of nodes. Suppose the true peaks of the suppliers are s and suppose they misreport their peaks to be s' . Fix a coalition A of agents: note that this coalition includes all the agents k with $s'_k \neq s_k$. Let x and x' be the respective allocations to the agents when they report s and s' respectively. As with the earlier proof, let M_+, M_- be the decomposition when the agents report s , and let M'_+, M'_- be the decomposition when the agents report s' . We shall show that when the agents report s' rather than s the only allocation in which each agent in A is (weakly) better off is one in which $x'_k = x_k$ for all $k \in A$, establishing the required contradiction.

Let $D' := Q_+ \cap Q'_-$. Note that $g(D') \subseteq M'_+$, for otherwise there will be a supplier in M'_- with a link to a demander in Q'_- . We now make two simple observations about the agents in $M_- \cap g(D')$:

- When the report is s' , every such agent can send flow only to the demanders in D' : this is because $f(M_-) \subseteq Q_+$, and each agent in $g(D')$ can send flow only to the agents in Q'_- . Therefore $\sum_{k \in M_- \cap g(D')} x'_k \leq \sum_{j \in D'} d_j$.
- When the report is s , the demanders in D' can receive flow only from such agents: the demanders in D' can receive flow only from the suppliers in M_- and they are connected only to the suppliers in $g(D')$. Therefore $\sum_{k \in M_- \cap g(D')} x_k \geq \sum_{j \in D'} d_j$.

Note also that $s'_k \leq x'_k$ and $x_k \leq s_k$ for any $k \in M_- \cap g(D')$, and that $s'_k = s_k$ for all $k \notin A$. These observations lead to

$$\sum_{\substack{k \in M_- \cap g(D') \\ k \notin A}} s_k + \sum_{\substack{k \in M_- \cap g(D') \\ k \in A}} x'_k = \sum_{\substack{k \in M_- \cap g(D') \\ k \notin A}} s'_k + \sum_{\substack{k \in M_- \cap g(D') \\ k \in A}} x'_k \leq \sum_{k \in M_- \cap g(D')} x'_k \leq \sum_{j \in D'} d_j, \quad (7)$$

and

$$\sum_{j \in D'} d_j \leq \sum_{k \in M_- \cap g(D')} x_k \leq \sum_{\substack{k \in M_- \cap g(D') \\ k \notin A}} s_k + \sum_{\substack{k \in M_- \cap g(D') \\ k \in A}} x_k. \quad (8)$$

For every agent in A to be (weakly) better off when reporting s' , we must have $x'_k \geq x_k$ for each $k \in A$. Combining this with inequalities (7) and (8), we conclude that $x'_k = x_k$ for each $k \in M_- \cap g(D') \cap A$. Moreover, these inequalities also imply that $x'_k = x_k = s_k$ for each $k \in M_- \cap g(D')$, $k \notin A$. Thus, $x'_k = x_k$ for all $k \in M_- \cap g(D')$. Also, whether the report is s or is s' , the suppliers in $M_- \cap g(D')$ send all of their flow only to the demanders in D' ; and that these demanders receive all of their flow only from the suppliers in $M_- \cap g(D')$. Therefore, removing the suppliers in $M_- \cap g(D')$ and the demanders in D' does not affect the egalitarian solution for either problem. As we picked a smallest counterexample, $D' = \emptyset$.

We now turn to the other case. Let $\tilde{X} := M_+ \cap M'_-$. Note that $f(\tilde{X}) \cap Q_- \subseteq Q'_+$, for otherwise there will be a supplier in M'_- linked to a demander in Q'_- . Consider the demanders in $f(\tilde{X}) \cap Q_-$:

- When the report is s' , every such demander can receive flow only from the suppliers in \tilde{X} : such demanders are linked only to the suppliers in M_+ and can receive flow only from the suppliers in M'_- . Therefore $\sum_{k \in \tilde{X}} x'_k \geq \sum_{j \in f(\tilde{X}) \cap Q_-} d_j$.
- When the report is s , the suppliers in \tilde{X} send flow only to the demanders in Q_- , and they can send flow only to the demanders they are connected to, so the suppliers in \tilde{X} can send flow only to the demanders in $f(\tilde{X}) \cap Q_-$. Therefore $\sum_{k \in \tilde{X}} x_k \leq \sum_{j \in f(\tilde{X}) \cap Q_-} d_j$.

Note also that $s'_k \geq x'_k$ and $x_k \geq s_k$ for any $k \in \tilde{X}$, and that $s'_k = s_k$ for all $k \notin A$. Putting all this together we have:

$$\sum_{k \in \tilde{X} \setminus A} s_k + \sum_{k \in \tilde{X} \cap A} x'_k = \sum_{k \in \tilde{X} \setminus A} s'_k + \sum_{k \in \tilde{X} \cap A} x'_k \geq \sum_{k \in \tilde{X}} x'_k \geq \sum_{j \in f(\tilde{X}) \cap Q_-} d_j, \quad (9)$$

and

$$\sum_{j \in f(\tilde{X}) \cap Q_-} d_j \geq \sum_{k \in \tilde{X}} x_k \geq \sum_{k \in \tilde{X} \setminus A} s_k + \sum_{i \in \tilde{X} \cap A} x_i. \quad (10)$$

For every agent in A to be (weakly) better off when reporting s' , we must have $x_k \leq x'_k$ for each $k \in A$. Combining this with inequalities (9) and (10), we conclude that $x'_k = x_k$ for each $k \in \tilde{X} \cap A$. Moreover, these inequalities also imply that $x'_k = x_k = s_k$ for each $k \in \tilde{X} \setminus A$. Thus, $x'_k = x_k$ for all $k \in \tilde{X}$. Note that the suppliers in \tilde{X} send all of their flow flow to the demanders in $f(\tilde{X}) \cap Q_-$, whether the report is s or s' ; also the demanders in $f(\tilde{X}) \cap Q_-$ receive all of their flow from the suppliers in \tilde{X} , whether the report is s or s' . Therefore, removing the suppliers in \tilde{X} and the demanders in $f(\tilde{X}) \cap Q_-$ does not affect the egalitarian solution for either problem. As we picked a smallest counterexample, $\tilde{X} = \emptyset$.

We now establish that the decomposition does not change in a smallest counterexample. We already know that $D' = \emptyset$, which implies $Q'_- \subseteq Q_-$. Suppose this containment is strict so that there is a demander $j \in Q_- \setminus Q'_-$. Then, $g(j) \subseteq M_+$. As $\tilde{X} = \emptyset$, $g(j) \subseteq M'_+$, which implies demander j cannot receive any flow when the report is s' . Therefore $Q'_- = Q_-$, which implies $Q'_+ = Q_+$, $M'_+ = M_+$, and $M'_- = M_-$.

To complete the argument, let A be as defined earlier. Let $A_+ = A \cap M_+$ and $A_- = A \cap M_-$. For any $i \in A_-$, $s'_i < x_i$ implies $x'_i \leq s'_i < x_i$, causing i to do worse by reporting s'_i . Likewise, any $i \in A_+$, $s'_i > x_i$ implies $x'_i \geq s'_i > x_i$, causing i to do worse by reporting s'_i . So any improving coalition A must be such that $s'_i \geq x_i$ for all $i \in A_-$ and $s'_i \leq x_i$ for all $i \in A_+$. But in this case the egalitarian solution does not change for either problem. \blacksquare

4 The Two-sided Model

4.1 Model

We now turn to the two-sided model, introduced by Bochet et al. [2]. This is closely related to the one-sided model, and yet there are some important differences, so our treatment will be concise, and will focus mostly on the aspects that make this model different. As in the one-sided model, we are given a bipartite graph G with suppliers (M) indexed by i and demanders by j . In the two-sided model, however, both the suppliers and demanders have single-peaked preferences: supplier i has a peak of s_i , and demander j has a peak of d_j . Define the graph $\Gamma(G, s, d)$ by adding a source σ connected to all suppliers, and a sink τ connected to all demanders; by orienting the

edges from source to sink; by setting the capacity of an edge in G to infinity, that of an edge $\sigma i, i \in S$, to s_i , and that of $j\tau, \tau \in D$, to d_j . A σ - τ cut (or simply a cut) in this graph is a subset C of nodes that contains σ but not τ . The capacity of a cut C is the total capacity of the edges that are oriented from a node in C to a node outside of C (such edges are said to be “in the cut”). It is well-known that the max-flow from σ to τ in this network is the same as the capacity of a min σ - τ cut. Fix an arbitrary min σ - τ cut C , and let X be the set of suppliers in C and Y be the set of demanders in C . It is not hard to see that $Y = f(X)$ in a min-cut. Bochet et al. [2] characterize the Pareto optimal solutions to this model using a network construction very similar to that used for the one-sided model. In particular they show that in any Pareto optimal allocation (x, y) , $x_i \leq s_i$ for $i \in X$ and $x_i \geq s_i$ for each supplier $i \notin X$; similarly, $y_j \leq d_j$ for all $j \in Y$, and $y_j \geq d_j$ for all $j \notin Y$. In addition they show that in a Pareto-optimal allocation, there is no flow from a supplier *not* in X to a demander in Y (even though links may exist between such suppliers and demanders). In effect, the problem decomposes into two subproblems, one involving the suppliers in X and the demanders in Y ; the other involving the suppliers *not* in X and the demanders *not* in Y . The egalitarian allocation is found independently for each subproblem, using exactly the same algorithm described earlier in §3.2. They also show that the egalitarian allocation can be found by an alternative algorithm as follows: first, set all the demands to be d_j and apply the one-sided egalitarian algorithm described earlier in §3.2 with the additional restriction that no supplier sends more than his peak, and that no demander receives more than hers; this gives the egalitarian allocation for the suppliers. The egalitarian allocation for the demanders is obtained by applying the same algorithm under the same additional restriction, but interchanging the roles of the suppliers and demanders (thus, we fix the supplies to be the s_i etc.) It is this equivalent definition that we will use in the proofs. A key result is that the egalitarian allocation is the Lorenz dominant element of the subset of the Pareto optimal allocations in which no agent sends or receives more than their peak (this set is called Pareto* by Bochet et al.). By picking a cut C with the largest X , we find a canonical decomposition for each problem, as such a cut is unique: to relate this with the notation of Bochet et al., the X corresponding to the largest C is S_- , and the demanders Y in C is the set D_+ ; the suppliers outside of X belong to S_+ , and the demanders outside of Y are in D_- .

We turn now to strategic issues. Bochet et al. [2] show that the egalitarian mechanism is both link strategyproof and peak strategyproof. Here we show that the egalitarian mechanism is in fact peak group strategyproof.

Theorem 3 *In the two-sided model, the egalitarian mechanism is peak group strategyproof.*

Proof. Suppose not. Focus on a counterexample G with the *smallest* number of nodes. Suppose the true peaks of the suppliers and demanders are s and d respectively, and suppose their respective misreports are s' and d' . We can assume that $d_j > 0$ for every demander j , as otherwise deleting j would result in a smaller counterexample. Fix a coalition A of suppliers and a coalition B of

demanders : note that A contains all the suppliers k with $s'_k \neq s_k$, and B includes all demanders ℓ with $d'_\ell \neq d_\ell$.

Let (x, y) and (x', y') be the respective allocations to the suppliers and demanders when they report (s, d) and (s', d') respectively. Let S_+, S_-, D_+, D_- be the decomposition when the agents report (s, d) , and let S'_+, S'_-, D'_+, D'_- be the decomposition when the agents report (s', d') . We shall show that when the agents report (s', d') rather than (s, d) , the only allocation in which each agent in $A \cup B$ is (weakly) better off, then $x'_k = x_k$ for all $k \in A$ and $y'_\ell = y_\ell$ for all $\ell \in B$. This establishes the required contradiction.

Let $Y' := D_+ \cap D'_-$. Note that $g(Y') \subseteq S'_+$, for, otherwise, there will be a supplier in S'_- with a link to a demander in D'_- . We now make two simple observations about the suppliers in $S_- \cap g(Y')$:

- For any such supplier k , $s'_k = x'_k$ and $x_k \leq s_k$. Also, $d_\ell = y_\ell$ and $y'_\ell \leq d'_\ell$ for any $\ell \in Y'$.
- When the report is s' , every such supplier can send flow only to the demanders in Y' : this is because $f(S_-) \subseteq D_+$, and each supplier in $g(Y')$ can send flow only to the agents in D'_- . Therefore $\sum_{k \in S_- \cap g(Y')} x'_k \leq \sum_{\ell \in Y'} y'_\ell$.
- When the report is s , the demanders in Y' can receive flow only from such suppliers: the demanders in Y' can receive flow only from the suppliers in S_- and they are connected only to the suppliers in $g(Y')$. Therefore $\sum_{k \in S_- \cap g(Y')} x_k \geq \sum_{\ell \in Y'} y_\ell$.

Finally, note that $s'_k = s_k$ for all $k \notin A$, and $d'_\ell = d_\ell$ for all $\ell \notin B$. These observations first lead to

$$\sum_{\substack{k \in S_- \cap g(Y') \\ k \notin A}} s_k + \sum_{\substack{k \in S_- \cap g(Y') \\ k \in A}} x'_k = \sum_{\substack{k \in S_- \cap g(Y') \\ k \notin A}} s'_k + \sum_{\substack{k \in S_- \cap g(Y') \\ k \in A}} x'_k = \sum_{k \in S_- \cap g(Y')} x'_k \leq \sum_{\ell \in Y'} y'_\ell. \quad (11)$$

Note that every demander ℓ in $Y' \cap B$ receives *exactly* his peak allocation d_ℓ for a truthful report, so for the coalition B of demanders to do weakly better in the (G, s', d') problem, $y'_\ell = d_\ell$ for each such ℓ . Therefore,

$$\sum_{\ell \in Y'} y'_\ell = \sum_{\ell \in Y' \setminus B} y'_\ell + \sum_{\ell \in Y' \cap B} y'_\ell \leq \sum_{\ell \in Y' \setminus B} d'_\ell + \sum_{\ell \in Y' \cap B} d_\ell = \sum_{\ell \in Y'} d_\ell. \quad (12)$$

Finally,

$$\sum_{\ell \in Y'} d_\ell = \sum_{\ell \in Y'} y_\ell \leq \sum_{k \in S_- \cap g(Y')} x_k \leq \sum_{\substack{k \in S_- \cap g(Y') \\ k \notin A}} s_k + \sum_{\substack{k \in S_- \cap g(Y') \\ k \in A}} x_k. \quad (13)$$

For every supplier in A to be (weakly) better off when reporting s' , we must have $x'_k \geq x_k$ for each $k \in S_- \cap g(Y')$. Combining this with inequalities (11) and (13), we conclude that all the inequalities in (11)-(13) hold as equations. In particular, $x'_k = x_k$ for all $k \in S_- \cap g(D')$, and $y'_\ell = y_\ell$ for $\ell \in Y'$. Therefore, whether the report is s or is s' , the suppliers in $S_- \cap g(Y')$

send all of their flow only to the demanders in Y' ; and that these demanders receive all of their flow only from the suppliers in $S_- \cap g(Y')$. Therefore, removing the suppliers in $S_- \cap g(Y')$ and the demanders in Y' does not affect the egalitarian solution for either problem. As we picked a smallest counterexample, Y' must be empty.

We now turn to the other case. Let $\tilde{X} := S_+ \cap S'_-$. Note that $f(\tilde{X}) \cap D_- \subseteq D'_+$, for otherwise there will be a supplier in S'_- linked to a demander in D'_- . Consider the demanders in $f(\tilde{X}) \cap D_-$:

- For any such demander ℓ , $d'_\ell = y'_\ell$ and $y_\ell \leq d_\ell$. Also, $s_k = x_k$ and $x'_k \leq s'_k$ for any $k \in \tilde{X}$.
- When the report is s' , every such demander can receive flow only from the suppliers in \tilde{X} : such demanders are linked only to the suppliers in S_+ and can receive flow only from the suppliers in S'_- . Therefore $\sum_{k \in \tilde{X}} x'_k \geq \sum_{\ell \in f(\tilde{X}) \cap D_-} y'_\ell$.
- When the report is s , the suppliers in \tilde{X} send flow only to the demanders in D_- , and they can send flow only to the demanders they are connected to, so the suppliers in \tilde{X} can send flow only to the demanders in $f(\tilde{X}) \cap D_-$. Therefore $\sum_{k \in \tilde{X}} x_k \leq \sum_{\ell \in f(\tilde{X}) \cap D_-} y_\ell$.

Finally, note that $s'_k = s_k$ for all $k \notin A$, and $d'_\ell = d_\ell$ for all $\ell \notin B$. Putting all this together, we have:

$$\sum_{\substack{\ell \in f(\tilde{X}) \cap D_- \\ \ell \notin B}} d_\ell + \sum_{\substack{\ell \in f(\tilde{X}) \cap D_- \\ \ell \in B}} d'_\ell = \sum_{\ell \in f(\tilde{X}) \cap D_-} d'_\ell = \sum_{\ell \in f(\tilde{X}) \cap D_-} y'_\ell, \quad (14)$$

and

$$\sum_{\ell \in f(\tilde{X}) \cap D_-} y'_\ell \leq \sum_{k \in \tilde{X}} x'_k \leq \sum_{k \in \tilde{X} \setminus A} s'_k + \sum_{k \in \tilde{X} \cap A} x'_k = \sum_{k \in \tilde{X} \setminus A} s_k + \sum_{k \in \tilde{X} \cap A} x'_k. \quad (15)$$

Note that every supplier k in $\tilde{X} \cap A$ receives *exactly* his peak allocation s_k for a truthful report, so for the coalition A of suppliers to do weakly better in the (G, s', d') problem, $x'_k = s_k$ for each such k . Thus,

$$\sum_{k \in \tilde{X} \setminus A} s_k + \sum_{k \in \tilde{X} \cap A} x'_k = \sum_{k \in \tilde{X}} s_k = \sum_{k \in \tilde{X}} x_k \leq \sum_{\ell \in f(\tilde{X}) \cap D_-} y_\ell \leq \sum_{\substack{\ell \in f(\tilde{X}) \cap D_- \\ \ell \notin B}} d_\ell + \sum_{\substack{\ell \in f(\tilde{X}) \cap D_- \\ \ell \in B}} y_\ell \quad (16)$$

For every demander in B to be (weakly) better off, we must have $y'_\ell \geq y_\ell$ for each $\ell \in f(\tilde{X}) \cap D_-$. Combining this with inequalities (14)-(16), we conclude that all the inequalities in (14)-(16) hold as equations. In particular, $x'_k = x_k$ for all $k \in \tilde{X}$, and $y'_\ell = y_\ell$ for $\ell \in f(\tilde{X}) \cap D_-$. Therefore, whether the report is s or is s' , the suppliers in \tilde{X} send all of their flow only to the demanders in $f(\tilde{X}) \cap D_-$; and that these demanders receive all of their flow only from the suppliers in \tilde{X} . Therefore, removing the suppliers in \tilde{X} and the demanders in $f(\tilde{X}) \cap D_-$ does not affect the egalitarian solution for either problem. As we picked a smallest counterexample, \tilde{X} must be empty.

We now establish that the decomposition does not change in a smallest counterexample. We already know that $Y' = \emptyset$, which implies $D'_- \subseteq D_-$. Suppose this containment is strict so that there is a demander $j \in D_- \setminus D'_-$. Then, $g(j) \subseteq S_+$. As $\tilde{X} = \emptyset$, $g(j) \subseteq S'_+$, which implies demander j cannot receive any flow when the report is s' (i.e. $x'_j = 0$). This is a contradiction since, $d'(j) > 0$, then the egalitarian solution allocates the pareto value $x'_j = d'_j$ for all $j \in D'_+$. (w.l.o.g we can skip the case $d'_j = 0$ as we can delete such a j to obtain the new decomposition or just place it in D_-). Therefore $D'_- = D_-$, which implies $D'_+ = D_+$, $S'_+ = S_+$, and $S'_- = S_-$.

To complete the argument, let A be as defined earlier. Let $A_+ = A \cap S_+$ and $A_- = A \cap S_-$, $B_+ = A \cap D_+$ and $B_- = A \cap D_-$. Now, for any $j \in B_+$, $d'_j \neq d_j$ implies $y'_j = d'_j \neq d_j$ causing j to do worse by reporting d'_j . Hence, it follows, $\forall j \in B_+$, $d'_j = d_j$. By a similar argument, we could establish $s'_j = s_j \forall j \in A_+$.

For any $i \in A_-$, $s'_i < x_i$ implies $x'_i \leq s'_i < x_i$, causing i to do worse by reporting s'_i . Likewise, any $i \in B_-$, $d'_i < y_i$ implies $y'_i \leq d'_i < y_i$, causing i to do worse by reporting d'_i . So any improving coalition A must be such that $s'_i \geq x_i$ for all $i \in A_-$ and $d'_i \geq y_i$ for all $i \in B_-$. But in this case the egalitarian solution does not change for either problem. ■




Figure 2: A counterexample for link GSP

We turn now to link group strategyproofness. That the egalitarian mechanism is *not* link group strategyproof in the two-sided model is not difficult to see. Consider the network shown in Figure 2. The network (a) represents the true network, with the peaks shown next to the agent labels. The egalitarian allocation gives 1 unit to each supplier and to each demander on this example. Suppose however supplier 1 and demander 2 collude, and supplier 1 does not report his link to demander 1. In the resulting network, shown in (b), each supplier still receives his peak allocation; demander 2 now receives her peak, and demander 1 receives nothing. Note that both members of the coalition weakly improve, and demander 2 strictly improves, proving that the egalitarian mechanism is in general *not* link group strategyproof³. The following result, however,

³This example may suggest that if we require each member of the deviating coalition to strictly improve their allocation, then the egalitarian mechanism may be link group strategyproof. However, this is also false, as shown by

shows that the egalitarian mechanism satisfies a limited form of link group strategyproofness.

Theorem 4 *In the two sided model, the egalitarian mechanism is link group strategy proof when the coalition is restricted to the set of suppliers only (demanders only).*

Proof. We prove the result for an arbitrary coalition of suppliers; the result for the demanders follow by a similar argument. Let A_i be the set of demanders that supplier i is linked to, and let A'_i be supplier i 's report. We may assume without loss of generality that any given demander finds all the suppliers acceptable: if demander j finds supplier i unacceptable, then supplier i cannot have a link to demander j regardless of his report, so clearly i 's manipulation opportunities are more restricted. Let ϕ and ϕ' be (any) egalitarian flows when the suppliers report A and A' respectively, and let x and x' be the corresponding allocation to the suppliers. We show that no coalition of suppliers can weakly benefit by misreporting their links unless each supplier in the coalition gets exactly their egalitarian allocation.

The proof is by induction on the number of type 2 breakpoints in the algorithm to compute the egalitarian allocation. Suppose the given instance has n type 2 breakpoints, and suppose X_1, X_2, \dots, X_n are the corresponding bottleneck sets of suppliers. If $n = 0$, every supplier is at his peak value in the egalitarian allocation, and clearly this allocation cannot be improved. Suppose $n \geq 1$. Define

$$\tilde{X}_\ell = \{i \in X_\ell \mid \sum_{j \in A_i} \phi'_{ij} \geq \sum_{j \in A_i} \phi_{ij}\},$$

and

$$\hat{X}_\ell = \{i \in X_\ell \mid \sum_{j \in A_i} \phi'_{ij} \leq \sum_{j \in A_i} \phi_{ij}\}.$$

We shall show, by induction on ℓ , that for each $\ell = 1, 2, \dots, n$:

- (a) $\phi'_{ij} = 0$ for any $i \in \tilde{X}_{\ell'}$, $j \in \cup_{i' \in X_{\ell'}} A_{i'}$, $\ell' > \ell$; and
- (b) $X_\ell \subseteq \hat{X}_\ell$.

The theorem follows from part (b) above.

Any supplier $k \in X_\ell \setminus \tilde{X}_\ell$ must have $A_k = A'_k$ as otherwise supplier k is part of the deviating coalition and does worse. Consider now a supplier $i \in \tilde{X}_\ell$ with $x_i < s_i$ and a supplier $k \in X_\ell \setminus \tilde{X}_\ell$. We have the following chain of inequalities:

$$\sum_{j \in A'_k} \phi'_{kj} = \sum_{j \in A_k} \phi'_{kj} < \sum_{j \in A_k} \phi_{kj} = x_k \leq x_i = \sum_{j \in A_i} \phi_{ij} \leq \sum_{j \in A_i} \phi'_{ij} \leq \sum_{j \in A'_i} \phi'_{ij}.$$

To see why, note that as $k \in X_\ell \setminus \tilde{X}_\ell$, the second inequality is true by definition, and also $A_k = A'_k$ (justifying the first equality). Also $k, i \in X_\ell$ and $x_i < s_i$, implies $x_k < s_i$, as suppliers k and i

Bogomolnaia and Moulin [3]. They construct an example involving 4 agents on each side with all peaks identically 1 in which a coalition of agents from both sides deviate and all *strictly* improve.

both belong to the same bottleneck set and supplier i is below his peak; this justifies the third inequality. The fourth and fifth inequalities follow from the fact that $i \in \tilde{X}_\ell$ and the fact that ϕ'_{ij} must be zero for all $j \in A_i \setminus A'_i$. This chain of inequalities implies that $x'_k < x_k \leq s_k$ and $x'_k < x'_i$. Therefore, when the suppliers report A' , supplier k must be a member of an “earlier” bottleneck set than supplier i . An immediate consequence is that demanders in $A'_k = A_k$ do not receive any flow from supplier i when the report is A' .

By the induction hypothesis, supplier $i \in X_\ell$ does not send any flow to the demanders in $\cup_{1 \leq i' \leq \ell-1} \cup_{k \in X_{i'}} A_{k'}$. Therefore

$$\{j \mid \phi'_{ij} > 0, j \in A_i\} \subseteq \{j \mid \phi_{ij} > 0, j \in A_i\}.$$

This observation, along with the fact that every $i \in \tilde{X}_\ell$ weakly improves, and the fact that X_ℓ is a type 2 breakpoint implies that $\sum_{j \in A_i} \phi'_{ij} = \sum_{j \in A_i} \phi_{ij}$, establishing (b). Furthermore, in such a solution, every demander $j \in A_i$ for $i \in \tilde{X}_\ell$ must receive all his flow from the suppliers in \tilde{X}_ℓ . In particular, the demanders in X_ℓ cannot receive any flow from suppliers in $X_{\ell'}$ for $\ell' > \ell$, establishing (a). To complete the proof we need to establish the basis for the induction proof, i.e., the case of $\ell = 1$. This, however, follows easily: it is easy to verify that the set $X_1 \setminus \tilde{X}_1$ must be empty, so $X_1 = \tilde{X}_1$. As X_1 is a type 2 bottleneck set, it is not possible for *every* member of X_1 to do weakly better unless the allocation remains unchanged. Thus, both (a) and (b) follow. ■

An easy implication is the following result, whose proof is an immediate consequence of the results we have already established.

Theorem 5 *In the two sided model, the egalitarian mechanism is group strategyproof w.r.t. to both links and peaks when the coalition is restricted to the set of suppliers only (demanders only).*

5 Conclusions




Figure 3: Bossiness of the egalitarian mechanism

Our main contribution is a proof of peak group strategyproofness of the egalitarian mechanism in the two settings considered by Bochet et al. A number of interesting open problems remain,

and we mention a few. First, Bochet et al. characterize the egalitarian mechanism using Pareto efficiency, strategyproofness, and an equity property that can be thought of as the equal treatment of equals adapted to this constrained setting. A natural question is if there is an alternative characterization that uses group strategyproofness, but weakens the efficiency or the equity requirement. Second, it will be interesting to characterize all (both link and peak) strategyproof mechanisms for this problem. Finally, note that the egalitarian mechanism has the somewhat undesirable property that an agent can change the allocation of other agents without altering his own, see the example of Figure 3. (Such mechanisms are said to be “bossy.”) Are there natural non-bossy mechanisms for this problem that are compelling from a normative point of view?

References

- [1] O. Bochet, H. Moulin, and R. Ilkılıç. Egalitarianism under earmark constraints,. Technical report, mimeo, 2010.
- [2] O. Bochet, H. Moulin, R. Ilkılıç, and J. Sethuraman. Clearing supply and demand under bilateral constraints. *Theoretical Economics*, forthcoming.
- [3] A. Bogomolnaia and H. Moulin. Random matching under dichotomous preferences. *Econometrica*, 72(1):257–279, 2004.
- [4] J. Randall Brown. The sharing problem. *Oper. Res.*, 27(2):324–340, 1979.
- [5] László Lovász and Michael D. Plummer. *Matching theory*. AMS Chelsea Publishing, Providence, RI, 2009. Corrected reprint of the 1986 original [MR0859549].
- [6] N. Megiddo. Optimal flows in networks with multiple sources and sinks. *Mathematical Programming*, 7(1):97–107, 1974.
- [7] N. Megiddo. A good algorithm for lexicographically optimal flows in multi-terminal networks. *American Mathematical Society*, 83(3), 1977.
- [8] Y. Sprumont. The division problem with single-peaked preferences: a characterization of the uniform allocation rule. *Econometrica*, 59(2):509–519, 1991.