Rationing Problems in Bipartite Networks

HERVE J. MOULIN
Rice University
and
JAY SETHURAMAN
Columbia University

The standard theory of rationing problems is extended to the bipartite context. The focus is on *consistency*, a compelling rationality property of fair division methods in the standard setting.

Categories and Subject Descriptors: G.2.3 [**Mathematics of Computing**]: Discrete Mathematics—*Applications*

General Terms: Algorithms; Economics; Theory

Additional Key Words and Phrases: Games, Networks/Graphs, Fair Division, Consistency

We start with a brief description of standard rationing problems, followed by an overview of the main results of our paper [Moulin and Sethuraman 2011]. The final sections report on more general results (work in progress) and some directions for further research.

1. STANDARD RATIONING PROBLEMS

Standard rationing problems are the simplest type of fair division problems, where an amount r of some divisible resource (e.g., money) must be divided between the agents i of a given set N, who each have a *claim* (or *demand*) x_i on the resources, and the total claim exceeds the available resources: $\sum_N x_i > r$. Bankruptcy and the distribution of emergency supplies are typical examples.

A rationing method h selects for every N, every vector x of claims, and every non-negative r such that $r \leq x_N$ (with the standard notation $x_N = \sum_N x_i$), a vector of shares $y = h(N, x, r)$ such that $0 \leq y \leq x$ and $y_N = r$. Three benchmark methods emerge clearly from the abundant axiomatic (see the surveys [Moulin 2002] and [Thomson 2003]), as well as the experimental and social psychology [Cook and Hegtvedt 1983; Deutsch 1975] literatures. The *proportional* method allocates shares in proportion to claims, so agent i's share is $y_i = (x_i/x_N)r$. The *uniform gains* (a.k.a. *equal awards*) method equalizes shares as much as permitted by the constraint $y_i \leq x_i$: therefore $y_i = \min\{\lambda, x_i\}$, where λ is chosen so that $\sum_N y_i = r$. Finally the *uniform losses* (a.k.a. *equal losses*) method equalizes losses (i.e., the differences $x_i - y_i$) as much as permitted by the constraint $y_i \geq 0$: thus $y_i = \max\{x_i - \mu, 0\}$, where μ is chosen so that $\sum_N y_i = r$.

In addition to these three methods, we can of course define a great variety of rationing methods, for instance fixed convex combinations of the three benchmarks. However such combinations violate *consistency*—an axiomatic requirement widely

Authors' addresses: `moulin@rice.edu`, `jay@ieor.columbia.edu`

regarded as a compelling rationality property for fair division methods in a variety of contexts (including TU games, matching, assignment, etc.; see [Thomson 2005]). A rationing method is consistent if, when we take away one agent from the set of participants, and subtract his share from the available resources, the division among the remaining set of claimants does not change. Formally,

$$\text{for all} \quad N, i, x, r, \quad h_j\big(N \diagdown \{i\}, x_{-i}, r - h_i(N,x,r)\big) = h_j(N,x,r) \text{ for all } j \neq i,$$

In the words of Balinski and Young [1982], "any part of a fair division should be fair."

An important result [Young 1987] characterizes all rationing methods that are symmetric (w.r.t. individual claims), continuous (in x and r), and consistent. All such methods, and only those, have a *parametric representation* as follows. Choose a continuous function $\theta(x, \lambda)$ from $\mathbb{R}^2_+$ into $\mathbb{R}_+$, non decreasing in λ and such that $\theta(x, 0) = 0$ and $\theta(x, \infty) = x$. Then divide r units of resources as follows:

$$y_i = \theta(x_i, \lambda) \text{ for all } i, \text{ where } \lambda \text{ solves } \sum_N \theta(x_i, \lambda) = r. \tag{1}$$

It is easy to see that the three benchmark methods are consistent and have such a representation.

Our goal in the paper [Moulin and Sethuraman 2011], and in subsequent forthcoming work, is to extend the rich theory of fair rationing to the more complex environment involving restrictions in the access to multiple types of resources.

2. BIPARTITE RATIONING AND CONSISTENCY

We consider the fair selection of a maximum flow in a bipartite network. The nodes on one side are *agents*, who wish to consume *resources*, which are the nodes on the other side. Not every agent has access to every resource, and such restrictions are captured by the bipartite graph. Each agent has a global *demand*, and the various resources are available in arbitrary quantities. We assume that the resources are *substitutable*, so that an agent cares only about the total amount of the resources he receives. A typical example is load balancing, where agents are workers, resources are types of jobs, and each agent is only qualified for certain jobs. See [Moulin and Sethuraman 2011] for more examples.

To motivate the consistency property in this context, we start with the example shown in Figure 1. Two agents Ann and Bob (nodes A and B on the left) have a

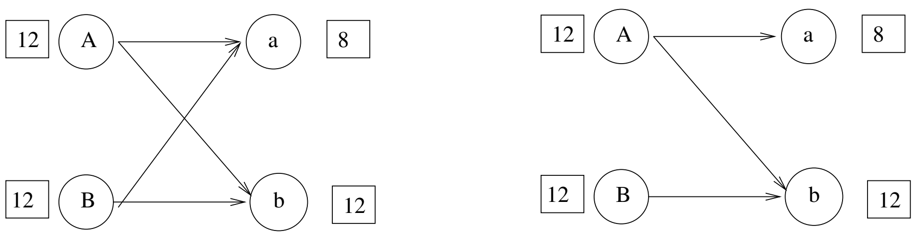

Fig. 1. An example with 2 sources and 2 sinks

demand of 12 units each ; node a has 8 units of the resource whereas node b has 12. In Figure 1a the bipartite graph is complete: every agent can access every resource node. Note that the resources are overdemanded: in any efficient allocation, all of the resources will be fully distributed. By symmetry Ann and Bob will receive 10 units each, which can be achieved by many different flows in the network.

Now assume Ann can access both a and b, whereas Bob can only access b (Figure 1b). It is still possible for Ann and Bob to receive 10 units each (by letting Ann receive only 2 units from b), and indeed this will be the *egalitarian* recommendation [Bochet et al. 2011]. Whether or not this is *fair* depends very much on our view of why Bob cannot access a. In this article, we postulate that agents ought to be held *responsible* for the compatibility constraints they impose. Therefore, as Ann and Bob have identical demands, but Ann is compatible with a larger set of resource nodes than Bob, she should receive more. More precisely Ann is entitled to all of the resource (8 units) at a, which she is the only one to claim. She still competes with Bob for the resources at b, but her claim on b is reduced from the original 12 to the *residual* claim of $12 - 8 = 4$ units. If we apply the proportional rule in this *standard* rationing problem for resource b, the 12 available units are split in proportion to the claims 4 and 12, thus Ann and Bob get 3 and 9 units respectively.

A version of consistency appropriate for bilateral rationing problems captures the idea that agents are responsible for their own compatibilities with the different types of resources. We can now take away either an agent or a type of resource: if the former, we subtract from each resource-type the share previously assigned to the departing agent; if the latter, we subtract from the claim of each agent the share of the departing resources he was previously receiving; in each case we insist that the division in the reduced problem remain as before. A stronger form of consistency can be applied to each edge of the graph: when we remove a certain edge, we subtract its flow from the capacity of both end nodes, and require as above that the solution choose the same flow in the reduced problem.

Our main goal is to understand which (and how) standard methods can be extended to bipartite methods while respecting consistency.

2.1 Model

A rationing problem involves a set N of n agents, a set Q of q *types*, and a bipartite graph $G \subseteq N \times Q$: an edge $(i, a) \in G$ indicates that agent i can consume the type a. We define $f(i)$ to be the (non empty) set of types that i is connected to, and $g(a)$ to be the (non empty) set of agents connected to type a. Each agent i has a claim x_i and each type a has a capacity (amount it can supply) r_a. We assume that the resources are (weakly) *overdemanded*, so that we have a rationing situation. Formally, this means

$$\textit{for all } B \subseteq Q\text{: } r_B \leq x_{g(B)}. \tag{2}$$

A bipartite *rationing problem* is a list $P = (N, Q, G, x, r)$ that satisfies (2). Let $\mathcal{P}$ denote the set of bipartite rationing problems.

A max-flow of P is a flow $\varphi = (\varphi_{ia})_{(i,a)\in G} \in \mathbb{R}_+^G$ sending all the resources to the agents, and such that no agent receives more than his demand:[1]

$$\varphi_{g(a)a} = r_a \text{ for all } a \in Q; \text{ and } \varphi_{if(i)} \leq x_i \text{ for all } i \in N$$

We write $\mathcal{F}(P)$ for the set of max-flows in P. A bipartite *rationing method* H associates to each problem $P \in \mathcal{P}$ a max-flow $\varphi = H(P) \in \mathcal{F}(P)$.[2]

We restrict attention to symmetric and continuous rationing methods: a method H is symmetric if the labels of the agents and resources do not matter; it is continuous if the mapping $(x, r) \to H(G, x, r)$ is continuous in the relevant subset of $\mathbb{R}_+^N \times \mathbb{R}_+^Q$. For brevity we use the informal definition of consistency given just before subsection 3.1: we speak of *node-consistency* when comparing a problem and its reduction after taking away an arbitrary agent or type, and of the stronger *edge-consistency* when we take away an arbitrary edge.

3. MAIN RESULTS

We find that the standard proportional method has a unique consistent extention to the bipartite case, whereas the two other benchmark methods have infinitely many such extensions.

We start with the proportional method. For any $z \geq 0$, define the function $En(z) = z \ln(z)$, with the convention that $En(0) = 0$ (so the sum $\sum_k En(z_k)$ is the familiar *entropy* of a vector z). Note that $En(z)$ is strictly convex.

Given a problem $P \in \mathcal{P}$, define $\widehat{\varphi}$ as

$$\widehat{\varphi} = \arg \min_{\varphi \in \mathcal{F}(P)} \sum_{ia \in G} En(\varphi_{ia}) + \sum_{i \in N} En(x_i - \phi_{if(i)}) \tag{3}$$

Problem (3) has a unique solution $\widehat{\varphi}$ because the objective function is strictly convex and finite. We define the *proportional method* H^{pro} as $H^{pro}(P) = \widehat{\varphi}$. This extends the standard proportional method: if G has exactly one resource node a, the problem is to minimize $\sum_{i \in N} En(y_i) + \sum_{i \in N} En(x_i - y_i)$ over all $y \geq 0$ such that $y_N = r_a$, and the Kuhn-Tucker conditions say that $\ln(\frac{y_i}{x_i - y_i})$ is independent of i.

THEOREM 3.1. *The proportional method H^{pro} is symmetric, continuous, and edge-consistent. Moreover, it is the only continuous and node-consistent method that is proportional for standard problems.*

While the above characterization is useful, an equivalent definition sheds further light on the structure of the proportional method. This definition is restricted to the subset of *irreducible* problems, defined by the property that all the inequalities in (2) for $B \subsetneqq Q$ are strict. These are the problems in which every edge in the graph carries a positive flow in some max-flow.

[1] We use the notation $\varphi_{g(a)a} = \sum_{i \in g(a)} \varphi_{ia}$, and $\varphi_{if(i)} = \sum_{a \in f(i)} \varphi_{ia}$.

[2] Any bipartite max-flow problem is the union of two separate rationing problems just described, one problem where all max-flows empty all resource nodes while agents are rationed, the other where max-flows fulfill all demands while resource nodes are not exhausted ([Moulin and Sethuraman 2011]). In this sense our theory applies to the selection of a fair max-flow in a general bipartite problem (N, Q, G, x, r) that may not satisfy (2).

THEOREM 3.2. *For any irreducible problem P the system of equations in z:*

$$\sum_i z_i = x_N - r_Q \text{ and } x_i = z_i + \sum_{a\in f(i)} \frac{z_i}{z_{g(a)}} r_a \text{ for all } i \in N \tag{4}$$

has a unique solution $\widehat{z} >> 0$, and the proportional flow is

$$\widehat{\varphi}_{ia} = \frac{\widehat{z}_i}{\widehat{z}_{g(a)}} r_a$$

Note that z_i is precisely agent i's loss: $z_i = x_i - \widehat{\varphi}_{if(i)}$. Hence a nice interpretation of Theorem 3.2: each resource is allocated proportionally to the agents who are connected to it, but the proportionality is with respect to the losses rather than the original claims. In the case of a standard method, losses $x_i - y_i$ are proportional to gains y_i, but this is not true any more in the bipartite context.

A straightforward generalization of Problem (3) delivers a large family of edge-consistent bipartite methods:

THEOREM 3.3. *Fix a strictly convex function W and a convex function B, both from $\mathbb{R}_+$ into itself. For any problem $P \in \mathcal{P}$ the flow*

$$\widetilde{\varphi} = \arg \min_{\varphi\in\mathcal{F}(P)} \sum_{ia\in G} W(\varphi_{ia}) + \sum_{i\in N} B(x_i - \varphi_{if(i)}) \tag{5}$$

defines an edge-consistent, symmetric, and continuous bipartite rationing method.

The bipartite proportional method corresponds to $W = B = En$, and we know from Theorem 3.1, that this is the only edge-consistent extension of the standard proportional method. In contrast, the bipartite rationing methods in Theorem 3.3 contain infinitely many extensions of the uniform gains, and, in a limit sense, of the uniform losses methods.

By taking any strictly convex W and $B \equiv 0$ in (5), we obtain a consistent extension of the uniform gains method. Formally:

PROPOSITION 3.4. *For any strictly convex function W from $\mathbb{R}_+$ into itself, and any problem $P \in \mathcal{P}$, the flow*

$$\mathring{\varphi} = \arg \min_{\varphi\in\mathcal{F}(P)} \sum_{ia\in G} W(\varphi_{ia}) \tag{6}$$

defines a bipartite rationing method H^W that is symmetric, continuous, edge-consistent and extends the uniform gains method for standard rationing problems. Different choices of W yield infinitely many different methods H^W.

Turning finally to the uniform losses method, we proceed in two steps; first we take $W \equiv 0$ and for B a strictly convex function in (5). Independently of the choice of B, this determines the net share $y_i = \varphi_{if(i)}$ of each agent i. Then we fix the edge-flows by one more minimization w.r.t. a strictly convex W in (5). In the following statement we write $\mathcal{Y}(P)$ for the set of feasible net shares $y = (\varphi_{if(i)})_{i\in N}$ when φ is a max-flow of P ($\varphi \in \mathcal{F}(P)$).

PROPOSITION 3.5. *Fix any two strictly convex functions W, B from $\mathbb{R}$ into itself. For any problem $P = (G, x, r) \in \mathcal{P}$, the net shares*

$$\dot{y} = \arg \min_{y \in \mathcal{Y}(P)} \sum_{i \in N} B(x_i - y_i)$$

and the flow

$$\dot{\varphi} = \arg \min_{\varphi \in \mathcal{F}(\dot{P})} \sum_{ia \in G} W(\varphi_{ia}) \text{ where } \dot{P} = (G, \dot{y}, r)$$

define a method $H^{B \succ W}$ that is symmetric, continuous, edge-consistent, and extends the uniform losses method for standard problems. The choice of B does not matter, but different choices of W yield infinitely many different methods.

4. TOWARD A GENERAL CHARACTERIZATION

It is puzzling that the proportional method admits a unique consistent extension, while uniform gains and uniform losses admit so many. Moreover in any consistent extension of uniform losses, the profile of net shares is fixed, whereas the various consistent extensions of uniform gains may yield different net shares.

A difficult, and more general, question is to understand which consistent standard methods can be extended at all to the bipartite framework as continuous, symmetric and node- (or edge-) consistent methods.

Some partial understanding of these questions follows from generalizing the system (4) in Theorem 3.2, where the unknown is the profile of losses z_i. To explain this approach, we first go back to the standard rationing model and note that a large subset of parametric methods (section 1) can be described in *loss format* as follows. System (1) explains how we compute shares given the profile of *demands*. Can we similarly compute shares if we know the profile of losses z? This amounts to asking if the following system in y, λ,

$$y_i = \theta(z_i + y_i, \lambda) \text{ for all } i, \text{ and } \sum_N y_i = r,$$

has a unique solution for a given vector z.

We impose two mild properties that all standard rationing methods discussed in the literature satisfy:

Ranking (RK): $x_i \leq x_j \Rightarrow y_i \leq y_j$
Ranking (RK*)*: $x_i \leq x_j \Rightarrow x_i - y_i \leq x_j - y_j$

In the parametric format this amounts to saying that $\theta(x, \lambda)$ is weakly increasing and 1-Lipschitz in x. If in addition the method satisfies the strict version of RK*:

Strict Ranking (SRK*)*: $x_i < x_j \Rightarrow x_i - y_i < x_j - y_j$,
then for any z_i, λ the equation $y_i = \theta(z_i + y_i, \lambda)$ has at most one solution $y_i = \beta(z_i, \lambda)$, and the function β is continuous and weakly increasing in both variables. It is the *loss-parametrization* of our method:

$$\text{given } z, \text{ the shares are } y_i = \beta(z_i, \lambda) \text{ where } \lambda \text{ solves } \sum_N \beta(z_i, \lambda) = r.$$

The proportional method admits the gain-parametrization $y_i = \theta(x_i, \lambda) = \frac{\lambda x_i}{1+\lambda}$, equivalent to its loss-parametrization $\beta(z_i, \lambda) = \lambda z_i$.

The uniform gains and uniform losses method do not meet SRK*, and the equation $y_i = \theta(z_i + y_i, \lambda)$ may indeed have a full interval of solutions y_i, so these two methods are excluded from the analysis below.

But the rich families of parametric methods known as equal sacrifice, and dual equal sacrifice [Young 1988], all admit loss-parametrizations. For instance the following equal sacrifice methods h^1 :

$$h^1 : \frac{1}{y_i} - \frac{1}{x_i} = \frac{1}{\lambda} \Leftrightarrow \theta(x_i, \lambda) = \frac{\lambda x_i}{x_i + \lambda}$$

has the loss-parametrization

$$\beta(z_i, \lambda) = \frac{\lambda z_i}{(\frac{z_i^2}{4} + \lambda z_i)^{\frac{1}{2}} + \frac{z_i}{2}}$$

and the dual equal sacrifice method h^2

$$h^2 : \frac{1}{x_i - y_i} - \frac{1}{x_i} = \lambda \Leftrightarrow \theta(x_i, \lambda) = \frac{\lambda x_i^2}{1 + \lambda x_i}$$

has the loss-parametrization

$$\beta(z_i, \lambda) = \frac{\lambda z_i^2}{1 - \lambda z_i}$$

The latter example shows that the function β may not be defined on the entire orthant $\mathbb{R}^2_+$, a feature that brings only minor technical complications.

Fix now a standard method h in the loss format $\beta(z_i, \lambda)$, and suppose that it can be extended to a consistent bipartite method H. Consider a *strictly overdemanded* problem $P = (N, Q, G, x, r)$, i.e., such that the inequalities in (2) are strict for all B, including $B = Q$. Then the flow $H(P) = \varphi$ is

$$\varphi_{ia} = \beta(z_i, \lambda^a), \quad \text{for all } ia \in G,$$

where the profiles of losses $z \in \mathbb{R}^N_+$ and of parameters $\lambda \in \mathbb{R}^Q_+$ uniquely solve the following system

$$x_i = z_i + \sum_{a \in f(i)} \beta(z_i, \lambda^a) \text{ for all } i \in N \tag{7}$$

$$\sum_{i \in g(a)} \beta(z_i, \lambda^a) = r_a \text{ for all } a \in Q \tag{8}$$

Indeed agent i's loss z_i is defined by $x_i = z_i + \sum_{a \in f(i)} \varphi_{ia}$. In repeated applications of consistency, eliminating all resources but a, the loss profile is preserved, therefore when we are left with a single resource, its flow is derived from the loss-parametrization. Note that system (7) and (8) boils down to (4) for the proportional method.

This result shows that under the assumption SRK*, a standard rationing method admits no more than one consistent bipartite extension, and this extension is well defined for all strictly overdemanded problems. The difficult open question is whether or not we can extend continuously the method H to all irreducible, then to all overdemanded problems. If the answer to the latter is positive, we will conclude that all standard methods meeting SRK* have a unique consistent extension to the bipartite setting.

5. DIRECTIONS FOR FUTURE RESEARCH

All reasonable symmetric standard rationing methods, $y_i = h(x_i, r)$, including the three benchmark methods, meet the following monotonicity properties (in addition to RK and RK*):

(1) *Resource Monotonicity (RM):* $r \to y_i$ is weakly increasing for all i

(2) *Claim Monotonicity (CM):* $x_i \to y_i$ is weakly increasing for all i

(3) *Cross Monotonicity (CRM):* $x_i \to y_j$ is weakly decreasing for all $j \neq i$

The corresponding properties for bipartite rationing methods are equally compelling. Properties CM and CRM have the same definition, and can be strengthened as follows:
Edge-CM: $x_i \to \varphi_{ia}$ is weakly increasing for all i and all $a \in f(i)$
Edge-CRM: $x_i \to \varphi_{ja}$ is weakly decreasing for all $j \neq i$ and $a \in f(i) \cap f(j)$

Then RM, RK, and RK* are adapted as follows:[3]
RM: $r_a \to y_i$ is weakly increasing for all $i \in g(a)$
RK: $x_i \leq x_j \Rightarrow y_i \leq y_j$ for all i, j such that $f(i) = f(j)$
RK*: $x_i \leq x_j \Rightarrow x_i - y_i \leq x_j - y_j$ for all i, j such that $f(i) = f(j)$

We have shown that system (7) and (8) implies the properties CM, CRM, RK, and RK*, but cannot say anything about any of these other properties, even for the bipartite proportional method. For the extensions of uniform gains and uniform losses, it is possible that the answer depends upon the choice of the convex function W.

Another interesting direction for future research is to study similar questions in a strategic framework. Consider an environment in which agents have single-peaked preferences over their allocations. The peaks (or more generally, preferences) are not known to the mechanism designer, who must design a strategyproof mechanism to allocate the resources to the agents. This is a well-studied problem in the standard rationing model [Sprumont 1991]; in particular, the mechanism assigning to each agent their allocation under the uniform-gains method is strategyproof, in the sense that no agent has an incentive to misreport their peaks (and the mechanism only works with the information about the peaks of the agents). It is a happy coincidence that this method is also consistent. In a recent series of papers, Bochet et al. [2011; 2012] study these questions in the bipartite framework and propose the egalitarian mechanism that generalizes the uniform gains method to the bipartite setting. Unfortunately the method underlying their mechanism fails consistency. A natural open question is to design a strategyproof mechanism in this setting such that the associated method is node or edge-consistent.

[3]They also admit an "edge" version, omitted for brevity.

REFERENCES

BALINSKI, M. AND YOUNG, H. P. 1982. *Fair Representation: Meeting the Ideal of One Man One Vote.* Yale University Press, New Haven, Conn.

BOCHET, O., İLKILIÇ, R., AND MOULIN, H. 2011. Egalitarianism under earmark constraints. *Submitted for Publication*.

BOCHET, O., İLKILIÇ, R., MOULIN, H., AND SETHURAMAN, J. 2012. Clearing supply and demand under bilateral constraints. *Theoretical Economics*.

COOK, K. AND HEGTVEDT, K. 1983. Distributive justice, equity and equality. *Annual Reviews of Sociology 9*, 217–241.

DEUTSCH, M. 1975. Equity, equality and need: what determines which value will be used as the basis of distributive justice? *Journal of Social Issues 31,* 3, 137–149.

MOULIN, H. 2002. Axiomatic cost and surplus sharing. In *Handbook of Social Choice and Welfare.* North-Holland.

MOULIN, H. AND SETHURAMAN, J. 2011. The bipartite rationing problem. Submitted for publication.

SPRUMONT, Y. 1991. The division problem with single-peaked preferences: a characterization of the uniform allocation rule. *Econometrica 59,* 2, 509–519.

THOMSON, W. 2003. Axiomatic and game-theoretic analysis of bankruptcy and taxation problems: a survey. *Math. Social Sci. 45,* 3, 249–297.

THOMSON, W. 2005. Consistent allocation rules. Mimeo, University of Rochester.

YOUNG, H. P. 1987. On dividing an amount according to individual claims or liabilities. *Math. Oper. Res. 12,* 3, 398–414.

YOUNG, H. P. 1988. Distributive justice in taxation. *J. Econom. Theory 44,* 2, 321–335.