

Fair Allocation of Production Externalities: Recent Results

Hervé Moulin* and Yves Sprumont†

December 7, 2005

Abstract

We survey recent axiomatic results in the theory of cost-sharing. In this literature, a method computes the individual cost shares assigned to the users of a facility for any profile of demands and any monotonic cost function.

We discuss two theories taking radically different views of the asymmetries of the cost function. In the *full responsibility* theory, each agent is accountable for the part of the costs that can be unambiguously separated and attributed to her own demand. In the *partial responsibility* theory, the asymmetries of the cost function have no bearing on individual cost shares, only the differences in demand levels matter.

We describe several invariance and monotonicity properties that reflect both normative and strategic concerns. We uncover a number of logical trade-offs between our axioms, and derive axiomatic characterizations of a handful of intuitive methods: in the full responsibility approach, the *Shapley-Shubik*, *Aumann-Shapley*, and *subsidy-free serial* methods, and in the partial responsibility approach, the *cross-subsidizing serial* method and the family of *quasi-proportional* methods.

*Department of Economics, Rice University MS 22, P.O. Box 1892, Houston, TX 77251-1892, USA.

†Département de Sciences Economiques and CIREQ, Université de Montréal, 3150 rue Jean-Brillant, Montréal H3C 3J7, Canada.

1 Introduction

Managing production facilities shared by multiple users is a vast empirical problem because externalities in production are the key to a host of environmental issues, from pollution to global warming, to the congestion of communication networks, roads, telephone, internet, and more generally to the limits on the rivalry and excludability in consumption of countless commodities. We review here a literature focusing on a simple yet foundational model of production externalities involving a single facility, in which each user's property rights are defined by the principle of user sovereignty. The goal is to define a fair allocation of the production externalities.

User i provides a quantity x_i of a certain idiosyncratic *input* and receives an amount y_i of a common *output* commodity. Alternatively, agents may demand various commodities produced from a single input: x_i is then interpreted as agent i 's idiosyncratic demand and y_i is her share of the total cost. We speak of an *output-sharing* problem in the first case, of a *cost-sharing* problem in the second. If N denotes the set of users, a *production* function F (or a *cost function* C) imposes the following technological constraint:

$$F(x) = \sum_{i \in N} y_i \text{ or } C(x) = \sum_{i \in N} y_i.$$

User Sovereignty means that each agent chooses freely the amount of input she wishes to contribute (or the amount of output she demands) and these contributions (or demands) are honored: no other input goes into the production process in the case of an output-sharing problem, all individual demands are exactly met in a cost-sharing problem. The design problem we address is how to share total output $F(x)$ (or total cost $C(x)$) among the agents of N .

Examples of the output-sharing problem include the following:

1) Exploitation of common property resources. A set of fishermen exploit a lake: they supply labor and share the catch. Because labor is homogeneous, the production function takes the form $F(z) = g(z(N))$, where we use the notation $z = (z_1, \dots, z_n)$ and $z(N) = z_1 + \dots + z_n$. When average product $g(z(N))/z(N)$ is decreasing, this is the standard model of the "tragedy of the commons" (Hardin, 1968, Case, 1979). When it is increasing, this is the "natural monopoly" problem (Sharkey, 1982).

2) Producers cooperatives. Agents supply labor and share profits. Skills differ: the production function is $F(z) = g(\alpha_1 z_1 + \dots + \alpha_n z_n)$, where $\alpha_1, \dots, \alpha_n$ are productivity parameters (Israelsen, 1968). Alternative assumptions on F are possible: complementary inputs (coordination games) $F(z) = g(\min_N \alpha_i z_i)$, a Cobb-Douglas technology $F(z) = \prod_N z_i^{\alpha_i}$, or general partnerships where no restriction on F can be imposed.

3) The problem of sharing a communication link with fixed total capacity can be approached as an output-sharing game where inputs are the users' bids for a share of the link: Johari and Tsitsiklis (2004), Yang and Hajek (2004), Sanghavi and Hajek (2004).

Examples of application of the cost-sharing model are more numerous:

1) Buyers cooperatives. The agents are interested in a homogeneous good and form a monopsony. This is a cost-sharing problem where $C(z) = c(z(N))$; average cost $c(z(N))/z(N)$ is typically decreasing.

2) Capacity sharing. Agents share a facility (ship channel, irrigation ditch, runway) and each requests a certain capacity z_i (depth of the channel or of the ditch, length of the runway). The facility is a nonrival good: the cost of meeting the demands is $C(z) = c(\max_N z_i)$. (Littlechild and Owen, 1973, Moulin, 1994, Aadland and Kolpin, 2004).

A number of cost-sharing problems emerged recently in the literature on networks:

3) Access to a network: each potential user chooses to be connected or not, $x_i = 0$ or 1 ; the mechanism divides the cost $C(S)$ of serving the subset S of agents. The function C is typically submodular (Moulin and Shenker, 2001, Feigenbaum et al., 2001, Archer et al., 2004, Immorlica et al., 2005).

4) Congestion and queues. Users share a link or a server, demands x_i are transmission rates and cost is delay in transmission or a monetary charge. In many queuing models, the server can achieve any profile of delays y_i such that $\sum_S y_i \geq c(\sum_S z_i)$ for all subsets S of users, where c is convex. For instance in the classic M/M/1 queue, $c(z(N)) = z(N)/(a - z(N))$: see Shenker (1995), Demers et al. (1990), Clarke et al. (1992). With quadratic or more general costs, see Johari and Tsitsiklis (2005), Johari et al. (2005), Chen and Zhang (2005).

We focus on the cost-sharing terminology throughout most of the paper. Although most of our axioms and results can be rephrased identically in the

output-sharing context, there are some discrepancies in the interpretation of some properties, particularly the upper bound property: see Section 7.1.

In our approach, the mechanism consists of a simple formula computing the cost shares y_i for any profile of demands x and any monotonic cost function C . We do not model individual preferences explicitly and think of each individual demand x_i as *inelastic*. This is in contrast with the axiomatic literature on cooperative production where a mechanism elicits individual preferences (and perhaps other characteristics) and selects an allocation (x_i, y_i) for each participant. See e.g., Moulin (1990), Roemer and Silvestre (1993), Fleurbaey and Maniquet (1996).

Two important assumptions are that cost shares are nonnegative and budget-balanced. The former is often a feasibility constraint, as when the cost y_i measures the waiting time until service is complete. When y_i represents cash, subsidies are often ruled out on normative grounds: because we limit attention to monotonic cost functions¹, non-negativity of the cost shares is a weak form of demand responsiveness.

It makes sense to relax the budget balance assumption when preferences are explicitly modeled and we consider the “demand game” where each participant chooses her demand strategically. If preferences are quasi-linear, the inherently budget-imbalanced Vickrey-Clarke-Groves mechanisms are of particular interest. Another argument is that the equilibrium of the demand game induced by a cost-sharing method violating budget balance may Pareto-dominate that of the game induced by a budget-balanced method (see e.g., Moulin (2005)).

In our axiomatic discussion below, we focus on the role of variable individual demands. We discuss two theories taking a radically different view of the asymmetries of the cost function C with respect to the various inputs. In the *full responsibility* theory, each agent is accountable for “her costs”, namely for the part of the costs that can be unambiguously separated and attributed to her own demand. In the *partial responsibility* theory, the asymmetries of the cost function have no bearing on individual cost shares, only the differences in demand levels matter. We also define several monotonicity properties of the mapping $x \mapsto y$ for a given function C , that have a dual normative and incentive compatibility interpretation. We uncover a number of logical trade-offs between our axioms, and end up recommending

¹There are natural examples of non monotonic cost functions. See, e.g., Flam and Jourani (2003).

a handful of intuitive methods, each one defining a coherent interpretation of fairness for managing production externalities. These methods are, in the full responsibility approach, the *Shapley-Shubik*, *Aumann-Shapley*, and *subsidy-free serial* methods, and in the partial responsibility approach, the *cross-subsidizing serial* method and the family of *quasi-proportional* methods. We state the main axiomatic characterizations of these methods. We conclude with a brief overview of the sources of the results we survey.

2 The model

Each agent i in a set $N = \{1, \dots, n\}$ demands an integer quantity $x_i \in \mathbb{N} = \{0, 1, \dots\}$ of an indiosyncratic good. The cost of meeting the demand profile $x \in \mathbb{N}^N$ must be split among the members of N . A cost function is a mapping $C : \mathbb{N}^N \rightarrow \mathbb{R}_+$ that is nondecreasing and satisfies $C(0) = 0$. A (cost-sharing) method φ assigns to each problem (C, x) a vector of nonnegative cost shares $\varphi(C, x) \in \mathbb{R}_+^N$ such that $\sum \varphi_i(C, x) = C(x)$.

The key axiom is the following invariance property.

Additivity. For all C, C' and x , $\varphi(C + C', x) = \varphi(C, x) + \varphi(C', x)$.

Additive cost-sharing methods are very convenient in practice. When production can be decomposed into the sum of several independent processes (like research, production and marketing; or construction and maintenance), applying the method to each subprocess and adding the resulting cost shares is equivalent to applying the method to the consolidated cost function. The proper level of application of the method is not a matter of dispute.

While devoid of ethical content, Additivity drastically reduces the set of cost-sharing methods. One consequence is that $\varphi(C, x)$ only depends upon the restriction of C to the “rectangle” $[0, x]$. Moulin and Vohra (2003) offer a representation of all additive methods in terms of flow on $[0, x]$, that plays an important role in our results below.

For simplicity we only discuss symmetric methods, even though most of the theory extends to possibly asymmetric methods. If π is a permutation on N and z is a demand profile, define πz by $(\pi z)_{\pi(i)} = z_i$ for all $i \in N$. For every cost function C , define πC by $\pi C(\pi z) = C(z)$ for all z .

Symmetry. For every permutation π on N , and all C and x , $\varphi(\pi C, \pi x) = \pi \varphi(C, x)$.

This axiom expresses the familiar idea that the names of the agents should be ignored when computing the cost shares: it is generally accepted as a basic notion of fairness. It implies weaker symmetry properties frequently used in the literature: for instance, agents with equal demands pay the same cost share when the cost function is symmetric.

The early work on axiomatic cost sharing was developed in a *continuous* variant of the above model: demands are real numbers and a cost function is a mapping $C : \mathbb{R}^N \rightarrow \mathbb{R}_+$ that is nondecreasing, continuously differentiable, and satisfies $C(0) = 0$. This model is technically more demanding than the discrete model, but the results are often parallel; many are identical in both models. However the literature on the continuous model has discussed exclusively the “full responsibility” theory: so far the alternative “partial responsibility” theory has only been developed in the discrete model. In order to discuss the latter approach in the continuous model, a required step is to adapt the Moulin-Vohra characterization of Additivity.

3 Full responsibility in costs and demands

Externalities in production render the naive principle “I am responsible for my own costs” ambiguous because Donald’s level of demand affects the (marginal) cost of Daisy’s demand. However, the interpretation of the principle is straightforward in one simple case, namely, when the cost function is additively separable. This leads to the following requirement.

Separability. For all C and x , $\{C(z) = \sum_{i \in N} c_i(z_i) \text{ for all } z\} \Rightarrow \{\varphi_i(C, x) = c_i(x_i) \text{ for all } i\}$.

Under Additivity, Separability is equivalent to the more familiar and intuitively less demanding Dummy axiom. For all i and C , define i ’s marginal cost function $\partial_i C : \mathbb{N}^N \rightarrow \mathbb{R}_+$ by $\partial_i C(z) = C(z + e^i) - C(z)$, where $e^i_i = 1$ and $e^i_j = 0$ for $j \neq i$.

Dummy. For all C , x , and i , $\{\partial_i C(z) = 0 \text{ for all } z\} \Rightarrow \{\varphi_i(C, x) = 0\}$.

The simplest cost-sharing problems (C, x) are those where individual demands are binary, that is, x belongs to $\{0, 1\}^N$. All the relevant information is then captured in the corresponding “stand-alone cost cooperative game” (N, C_x) where the cost associated with a coalition $S \subseteq N$ is just the cost

of meeting the consumptions of its members: in straightforward notation, $C_x(S) = C(x_S, 0_{N \setminus S})$.

The *Shapley cost-sharing rule* φ^{Sh} , which is restricted to such binary-demand problems, solves the original problem by applying the familiar Shapley value Sh to its stand-alone game. We use the notation e^S for the vector $e_i^S = 1$ if $i \in S$, $e_i^S = 0$ otherwise. For every permutation π on N and every $i \in N$, let $S(i, \pi) = \{j \in N | \pi(j) < \pi(i)\}$ and compute agent i 's cost share as

$$\varphi_i^{Sh}(C, x) = Sh_i(N, C_x) = \frac{1}{n!} \sum_{\pi} \partial_i C(e^{S(i, \pi)}), \quad (1)$$

where the summation is taken over all $n!$ permutations on N .

Theorem 1. *On the subset of binary-demand problems, the Shapley cost sharing method is the only method satisfying Additivity, Dummy and Symmetry.*

Removing Symmetry leads to the so-called *random-order values*. The incremental method based on the permutation π charges to agent i the share $\varphi_i^{\pi}(C, x) = \partial_i C(e^{S(i, \pi)})$; Additivity and Dummy characterize the convex combinations of the $n!$ incremental methods.

Theorem 1 is the founding result of the literature we are surveying. When demands can take more than two values, Additivity, Separability and Symmetry allow many more methods: the goal of this survey is to discuss the most important ones.

A stronger requirement than Separability applies to the fairly large classes of submodular and supermodular cost functions. For any C , i, j , and z , define $\partial_{ij}C(z) = \partial_i C(z + e^j) - \partial_i C(z)$. A cost function C is *submodular* if $\partial_{ij}C(z) \leq 0$ for all z and distinct i, j , and *supermodular* if $\partial_{ij}C(z) \geq 0$ for all z and distinct i, j .

Under such cost functions, externalities in production are unambiguous: positive in the case of submodular costs, negative in the case of supermodular costs.

Stand-Alone Principle. For all C, x and i , $\{C \text{ is submodular}\} \Rightarrow \{\varphi_i(C, x) \leq C(x_i, 0_{N \setminus i})\}$, and $\{C \text{ is supermodular}\} \Rightarrow \{\varphi_i(C, x) \geq C(x_i, 0_{N \setminus i})\}$.

This is an appealing requirement because the stand-alone cost is a natural benchmark. Note that the Stand-Alone Principle implies Separability since an additively separable function is both submodular and supermodular.

Conversely, one can show that Additivity and Separability together imply the Stand-Alone Principle.

Even stronger is the requirement that, when externalities are unambiguous, any increase in an agent's demand affects all other users in the "correct" direction.

Unambiguous Externalities. Fix C and let x, x' and i be such that $x_i < x'_i$ and $x_j = x'_j$ for all $j \in N \setminus i$. Then $\{C \text{ is submodular}\} \Rightarrow \{\varphi_j(C, x) \geq \varphi_j(C, x') \text{ for all } j \in N \setminus i\}$, and $\{C \text{ is supermodular}\} \Rightarrow \{\varphi_j(C, x) \leq \varphi_j(C, x') \text{ for all } j \in N \setminus i\}$

We call the first part of the axiom **Negative Externalities** and the second part **Positive Externalities**. Unambiguous Externalities guarantees that all agents face common and correct incentives: if the cost function is submodular (respectively, supermodular), everyone supports an increase (respectively, a decrease) in everyone else's demand. This is also appealing from a normative point of view: because i 's cost share depends only upon the marginal cost function $\partial_i C$ (as explained below), an increase in j 's demand must not increase i 's share if C is submodular (or decrease it if C is supermodular).

Note that, even under Additivity, Unambiguous Externalities is a strictly stronger requirement than the Stand-Alone Principle.

Finally, we consider one more simple-minded principle: "I am responsible for my own demand". This is even more vague than "I am responsible for my own costs", yet it suggests another normatively compelling property. Because raising my demand never decreases the total cost shared by all users -for *any* cost function-, this move should not allow me to lower my share of the cost.

Monotonicity. For all C, x, x' , and i , $\{x_i < x'_i \text{ and } x_j = x'_j \text{ for all } j \in N \setminus i\} \Rightarrow \{\varphi_i(C, x) \leq \varphi_i(C, x')\}$.

Monotonicity is a compelling ethical requirement in any cost-sharing theory holding agents responsible for their demand; it is meaningful in both the full and partial responsibility approaches. Alternatively, it may be defended on strategic grounds: a monotonic method is not vulnerable to artificial inflation of individual demands.

The combination of all axioms above leaves us with a rich class of cost-sharing methods.

Consider first the combination of Additivity and Dummy: it implies that an agent's cost share must be an average of her marginal costs at the demand profiles below the actual demand x . In order to express the budget balance condition, define a *flow* to x to be a mapping $f(., x) : [0, x[\rightarrow \mathbb{R}_+^N$ satisfying the convention that $f_i(z, x) = 0$ whenever $z_i = x_i$, the normalization $\sum_{i \in N} f_i(0, x) = 1$, and the conservation constraints $\sum_{i \in N} f_i(z, x) = \sum_{i \in N(z)} f_i(z - e^i, x)$ for all $z \in]0, x[$, where $N(z) = \{i \in N \mid z_i > 0\}$. An important lemma states that a method φ satisfies Additivity and Separability if and only if, for every x , there is a (necessarily unique) flow $f(., x)$ to x such that

$$\varphi_i(C, x) = \sum_{z \in]0, x[} f_i(z, x) \partial_i C(z) \quad (2)$$

for all C and all i ; we call f the *flow system* associated with φ .

Next, Unambiguous Externalities and Monotonicity connect the flows to the various demand profiles in a simple way: the flow to a demand profile is simply the projection of the flow to any higher demand profile. Formally, the *projection* of the flow $f(., x)$ to x on $[0, x'[\subseteq [0, x[$, denoted $p_{x'} f(., x)$, is defined as follows: for any i and $z \in [0, x'[\$ write $K = \{j \in N \mid z_j = x'_j\}$ and let

$$p_{x'} f_i(z, x) = f_i(z, x) \text{ if } K = \emptyset$$

and

$$\begin{aligned} p_{x'} f_i(z, x) &= 0 \text{ if } K \neq \emptyset \text{ and } i \in K, \\ &= \sum_{w_K \in [x'_K, x_K]} f_i((w_K, z_{N \setminus K}), x) \text{ if } K \neq \emptyset \text{ and } i \notin K. \end{aligned}$$

Note that $p_{x'} f(., x)$ is a flow to x' . A method φ is a *fixed-flow method* if its associated flow system f is such that $f(., x') = p_{[0, x']} f(., x)$ whenever $x' \leq x$. Note that φ is entirely determined by the subsystem $\{f(., ke^N) \mid k \in \mathbb{N}\}$.

Theorem 2. *The fixed-flow methods are the only methods satisfying Additivity, Unambiguous Externalities, and Monotonicity.*

In this result Unambiguous Externalities may be replaced by the combination of Dummy and Positive Externalities, or by the combination of Dummy and Negative Externalities.

4 Ordinality: the Shapley-Shubik method

The stand-alone cost game associated with a cost-sharing problem remains well defined when some agents consume several units. Shubik (1962) recommended the Shapley value of that game as a solution to the underlying cost-sharing problem. Formally, the so-called *Shapley-Shubik method* φ^{ShSh} simply applies the formula

$$\varphi^{ShSh}(C, x) = Sh(N, C_x)$$

to all problems (C, x) . Interestingly, the formula is identical in the continuous variant of the cost-sharing model.

This very simple solution is a fixed-flow method. For any demand profile x , the corresponding flow $f^{ShSh}(\cdot, x)$ is symmetrically spread along the edges of the cube $[0, x[$. Letting $E_i(x) = \{z \in [0, x[\mid \forall j \in N \setminus i, z_j \in \{0, x_j\}\}$, $f_i^{ShSh}(z, x) = \frac{(n(x)-n(z,x))!(n(z,x)-1)!}{n(x)!}$ if $z, z + e^i \in E_i(x)$ and $f_i^{ShSh}(z, x) = 0$ otherwise, where $n(z, x) = |\{j \in N \mid z_j < x_j\}|$.

The qualitative features of a commodity (such as temperature, height, weight, viscosity, brightness, resistance to shocks) may often be measured on different scales. The Ordinality property emerges when the choice of a particular numerical scale to measure demands (or inputs) should not matter to the eventual division of shares.

The property is easy to formulate in the continuous model. An *ordinal transformation* is an increasing and continuously differentiable bijection from $\mathbb{R}_+$ onto itself. Given a demand profile x , an agent i , and an ordinal transformation τ_i , we write $x^{\tau_i} = (\tau_i(x_i), x_{N \setminus i})$. If C is a valid cost function, then so is the function $z \mapsto C^{\tau_i}(z) = C(\tau_i(z_i), z_{N \setminus i})$. Because the problems (C, x^{τ_i}) and (C^{τ_i}, x) are ordinally equivalent, they should receive the same solution.

Ordinality. For all C , $x \in \mathbb{R}_+^N$, $i \in N$, and all ordinal transformations τ_i , $\varphi(C, x^{\tau_i}) = \varphi(C^{\tau_i}, x)$.

A corresponding property in the discrete model is the following: if it happens that C is flat between w_i and $w_i + 1$, then erasing unit $w_i + 1$ from the books should have no impact on cost shares. Formally, given a cost function C define $C^{w_i}(z) = C(z)$ if $z_i < w_i$ and $C^{w_i}(z) = C(z + e^i)$ otherwise.

Ordinality. For all C , $x \in \mathbb{N}^N$, $i \in N$, and all $w_i \leq x_i$, $\{\partial_i C(z) = 0 \text{ whenever } z_i = w_i\} \Rightarrow \{\varphi(x, C) = \varphi(x - e^i, C^{w_i})\}$

The interpretation of Ordinality in the discrete model is slightly different. When goods come in indivisible units, the property allows us to add artificial costless “half units” without affecting the final allocation of costs. We cannot use an “ordinal transformation” as above to change the measurement scale of an individual demand, because the only increasing bijection of $\mathbb{N}$ into itself is the identity.

Note that for an ordinal method, the choice of the zero in the scale of each demand does matter. When each scale is qualitative, the zero is a benchmark lower bound of the demand: it may be chosen to represent the status quo *ex ante*, as when the users of a facility share the cost of its renovation, each user requesting an upgrade of a different attribute (sound-proofing, air-conditioning system, light,...).

Theorem 3. *In both the discrete and continuous models, the Shapley-Shubik method is the only method satisfying Additivity, Dummy, Symmetry, and Ordinality.*

In particular, the Shapley-Shubik method is the only symmetric fixed-flow method meeting Ordinality. As in Theorem 1, removing Symmetry quickly leads to a characterization of the random-order values: they are characterized by Additivity, Dummy, Ordinality, and the mild condition of Independence of Zero Demands discussed in Section 5.1 below.

In Section 5.2, we give a related characterization of the Shapley-Shubik method in the continuous model, where Ordinality is replaced by the combination of Scale Invariance (a weakening of Ordinality) and Monotonicity.

5 Merging, splitting and reshuffling: the Aumann-Shapley method

When some agents consume several units, the stand-alone cost game (N, C_x) associated with a cost-sharing problem (C, x) ignores a lot of the potentially relevant information contained in the cost function C . The Aumann-Shapley method exploits this information by constructing a game (N_x, Γ_x) where each *unit* consumed by each agent is regarded as a separate player. The Shapley value of that game determines a price for every unit of consumption. The Aumann-Shapley method charges to agent i the sum of the prices attached to the units she consumes.

5.1 Discrete formulation

It will be useful to consider a variable-population version of the discrete cost-sharing model described earlier: a cost-sharing problem is now a triple (N, C, x) , where the agent set N is any nonempty finite subset of $\mathbb{N}$.

Given such a problem, we construct a cooperative game with transferable utility among $x(N)$ players, each player representing one unit of an individual demand. Write $N = \{1, \dots, n\}$, choose pair-wise disjoint sets $N_1, \dots, N_n$ such that $|N_i| = x_i$ for each $i \in N$, and let $N_x = \cup_{i \in N} N_i$. For each $S \subseteq N_x$ define $\Gamma_x(S) = C(|S \cap N_1|, \dots, |S \cap N_n|)$.

The *Aumann-Shapley method* φ^{Ash} computes the cost shares in the problem (N, C, x) according to the formula

$$\varphi_i^{Ash}(N, C, x) = \sum_{j \in N_i} Sh_j(N_x, \Gamma_x) \text{ for all } i \in N,$$

where we recall that Sh denotes the familiar Shapley value of for cooperative games.

The anonymity of the Shapley value guarantees that this definition is meaningful: the game (N_x, Γ_x) is not uniquely defined but all possible choices are equivalent. We omit N as an argument of the method φ^{Ash} whenever this causes no confusion.

The Aumann-Shapley method admits a simple flow representation as in (2). For any demand profile x , the corresponding flow $f^{Ash}(\cdot, x)$ is equally spread on all the nondecreasing paths in the rectangle $[0, x[$. Thus $f_i^{Ash}(z, x)$ equals the proportion of paths from 0 to x which go through z and $z + e^i$: straightforward computations yield $f_i^{Ash}(z, x) = \alpha(z)\alpha(x - z - e^i)/\alpha(x)$ for all $i \in N$ and $z \in [0, x - e^i]$, where $\alpha(z) = z(N)! / \prod_{j \in N} z_j!$.

This method, however, is not a fixed-flow method. It satisfies all axioms in Theorem 2 but Monotonicity, as shown in the following example.

Example 1. Returning to the output-sharing interpretation of our model, consider a production function with two perfectly complementary inputs: $F(z) = g(\min\{z_1, z_2\})$. If $x_1 = x_2 = a$, the Aumann-Shapley output shares are $y_1 = y_2 = \frac{g(a)}{2}$ since the method is symmetric. But when $x_1 < a = x_2$, we get $y_1 > g(a)/2 > y_2$: for instance, the output shares are $(y_1, y_2) = (\frac{2}{3}g(1), \frac{1}{3}g(1))$ when $(x_1, x_2) = (1, 2)$. This violates Monotonicity. Also, the high-input agent gets a smaller share of output, despite the symmetry of

the production function. In contrast, the Shapley-Shubik method gives $y_1 = y_2 = \frac{1}{2}g(\min\{x_1, x_2\})$, a much more palatable division of output.

This observation notwithstanding, the Aumann-Shapley method has several extremely appealing features. First, cost shares are *proportional* to demands whenever all goods are perfect substitutes. This property is known as Weak Consistency in the literature on the continuous model.

Weak Consistency. Let (N, C, x) be a cost-sharing problem such that $C(z) = c(\sum_{i \in N} z_i)$ for all z . Then $\varphi_i(N, C, x) = (x_i / \sum_{j \in N} x_j)C(x)$ for all $i \in N$.

Weak Consistency is intuitive from an ethical viewpoint. Perhaps more importantly, it is necessary and sufficient to rule out certain natural strategic maneuvers. In a context where demands cannot easily be traced to their actual consumers, a coalition of agents may contemplate *merging* into a single large consumer whose demand is the sum of the individual demands; dually, a single agent may *split* his demand between a number of “virtual” consumers. It is well known that when all goods are perfect substitutes, the only cost-sharing method for which such merging or splitting tactics are never profitable is the proportional method: see for instance Moulin (2002) or, for a more comprehensive treatment, Ju, Miyagawa, and Sakai (2005).

The idea of preventing merging or splitting maneuvers can be applied even when not all individual demands are perfect substitutes. Suppose this is true for a strict *subset* S of the goods, namely the cost function takes the form $C(z) = c(z_{N \setminus S}, \sum_{i \in S} z_i)$. In this case one would like to ensure that the agents consuming the goods in S have no incentive to merge or split. The No Merging or Splitting axiom strengthens Weak Consistency by imposing this requirement for any subset of goods.

No Merging or Splitting. Fix N, C, x and $i \in N$, and let I be a finite subset of $\mathbb{N}$ such that $N \cap I = \{i\}$. Write $N' = (N \setminus i) \cup I$, define the cost function C' on $\mathbb{N}^{N'}$ by $C'(z) = C(z_{N \setminus i}, \sum_{i' \in I} z_{i'})$ for all $z \in \mathbb{N}^{N'}$, and let $x' \in \mathbb{N}^{N'}$. Then $\{\sum_{i' \in I} x'_{i'} = x_i \text{ and } x'_j = x_j \text{ for all } j \in N \setminus i\} \Rightarrow \{\sum_{i' \in I} \varphi_{i'}(N', C', x') = \varphi_i(N, C, x)\}$.

This property connects cost shares in two problems with different sets of agents: N' obtains from N by “splitting” agent i into a set I of agents i' whose aggregate demand equals i ’s original demand. The other agents in N and N' are identical, and their demands do not change. The cost function C' for N' expresses the same technology as C : the sum of the consumptions

by the agents i' in I merely plays the role of agent i 's original consumption. The condition prevents manipulations of identity: agent i in N does not gain by splitting her consumption into smaller pieces; the agents i' in N' do not gain by merging theirs into one larger block.

The Aumann-Shapley method satisfies the axiom: when a subset S of the goods are perfect substitutes, the total cost share paid by the agents demanding those goods is $\sum_{j \in \cup_{i \in S} N_i} Sh_j(N_x, \Gamma_x)$, which is completely insensitive to merging or splitting maneuvers within S . In fact, such maneuvers also leave the cost shares of the agents in $N \setminus S$ unchanged (see, e.g., Monderer and Neyman, 1988).

Because No Merging or Splitting compares problems with different sets of agents, it is a somewhat complex condition. A related fixed-population condition which is a little simpler to grasp is No Reshuffling.

No Reshuffling. Let (N, C, x) be a cost-sharing problem such that $C(z) = c(z_{N \setminus S}, \sum_{i \in S} z_i)$ for all z and some $S \subseteq N$. For all x' , $\{\sum_{i \in S} x_i = \sum_{i \in S} x'_i$ and $x_{N \setminus S} = x'_{N \setminus S}\} \Rightarrow \{\sum_{i \in S} \varphi_i(N, C, x) = \sum_{i \in S} \varphi_i(N, C, x')\}$.

This condition says that the aggregate cost share of a group of agents consuming essentially the same good depends only on their aggregate consumption. If this condition were violated, all members of S could benefit by reshuffling individual consumptions within S and performing suitable monetary transfers. It is easy to check that No Reshuffling is implied by No Merging or Splitting. The converse is not true, however, even in the presence of Additivity and Dummy, unless we add the following mild variable-population requirement.

Independence of Zero Demands. For all N, C, x and i , $\{x_i = 0\} \Rightarrow \{\varphi_i(N, C, x) = 0 \text{ and } \varphi_{N \setminus i}(N, C, x) = \varphi(N \setminus i, C_{N \setminus i}, x_{N \setminus i})\}$,

where $C_{N \setminus i}(z_{N \setminus i}) = C(0_i, z_{N \setminus i})$.

Theorem 4. *In the discrete model, the Aumann-Shapley method is the only cost-sharing method satisfying Additivity, Dummy, and No Merging or Splitting.*

Alternatively, it is the only method satisfying Additivity, Dummy, No Reshuffling, and Independence of Zero Demands.

In contrast to Theorems 1 and 3, Theorem 4 does not require Symmetry. A consequence is that all asymmetric methods satisfying Additivity and Dummy are vulnerable to merging or splitting maneuvers.

5.2 Continuous formulation

We return to the fixed-population framework but assume that goods are perfectly divisible. While the literature often casts the problem in pricing terms –rather than sharing the total cost among agents, one seeks to price the goods that they consume–, the two formulations are equivalent.

Viewing this continuous model as the limit of the discrete one when the discretization grid grows finer and finer, Aumann and Shapley (1974) obtained the following sleek formula: for every cost-sharing problem (C, x) , agent i 's cost share is the integral of her marginal costs along the ray to x :

$$\varphi_i^{Ash}(C, x) = x_i \int_0^1 \partial_i C(tx) dt.$$

This simple “diagonal” formula naturally lends itself to numerous applications: from the early work of Billera, Heath, and Raanan (1978) and Samet, Tauman, and Zang (1984) to the recent contributions of Castano-Pardo and Garcia-Diaz (1995), Haviv (2001) or Lee (2002), they range from the pricing of utilities such as water, phone or electricity to the allocation of highway construction costs, and the sharing of waiting time at a congested server.

The diagonal formula also exacerbates the drawbacks of the discrete Aumann-Shapley method. We start with two examples where it behaves poorly, contrary to the Shapley-Shubik method.

Example 1 (continued). Consider a continuously differentiable production function arbitrarily close to $F(z) = g(\min_N z_i)$. Any user j such that $x_j > \min_N x_i$ receives an output share arbitrarily close to zero: virtually all the output goes to the smallest contributors. The Shapley-Shubik method recommends equal shares.

Example 2. Consider a continuously differentiable cost function arbitrarily close to $C(z) = c(\max_N z_i)$. Any user j such that $x_j < \max_N x_i$ pays a cost share arbitrarily close to zero: all the cost is borne by the largest users. In contrast, the Shapley-Shubik method yields the incremental shares: ranking agents by increasing order of demands, that is, $x_1 \leq x_2 \leq \dots \leq x_n$, we get

$$\varphi_i(C, x) = \sum_{k=1}^i \frac{C(x_k) - C(x_{k-1})}{n - k + 1} \quad (3)$$

for each i . This celebrated “runway” formula for the capacity game (derived by Littlechild and Owen, 1973) is the natural solution. In particular, it is

recommended by all the symmetric fixed-flow methods.

In the next examples, the Aumann-Shapley delivers a plausible proportional division.

Example 3. The production function takes the form $F(z) = g(\sum_N \alpha_i z_i)$. Here the Aumann-Shapley output shares are proportional to the input contributions measured in efficiency units: agent j receives $y_j = \frac{\alpha_j x_j}{\sum_N \alpha_i x_i} F(x)$. The Shapley-Shubik shares are much less intuitive.

Example 4. The production function takes the form $F(z) = g(\prod_N z_i^{\alpha_i})$. The Aumann-Shapley method allocates $y_j = \frac{\alpha_j}{\sum_N \alpha_i} F(x)$ to agent j , whereas the Shapley-Shubik method simply shares output equally. In both methods the relative output shares of any two agents are entirely insensitive to individual input contributions.

Like its discrete counterpart, the continuous Aumann-Shapley method is invulnerable to merging, splitting and reshuffling. The exact counterpart of No Merging or Splitting for pricing rules appears in Tauman (1988), who does not use the axiom, however². The standard axiomatization of this method relies on the weaker axiom of Weak Consistency (asking only that cost shares be proportional to demands when *all* goods are perfect substitutes) and a weakening of Ordinality known as Scale Invariance.

A *linear transformation* is an increasing linear mapping from $\mathbb{R}_+$ onto itself. Using the same notation as when defining Ordinality, the axiom reads as follows.

Scale Invariance. For all $C, x \in \mathbb{R}_+^N$, $i \in N$, and all linear transformations τ_i , $\varphi(C, x^{\tau_i}) = \varphi(C^{\tau_i}, x)$.

Theorem 5. *In the continuous model, the Aumann-Shapley method is the only method satisfying Additivity, Weak Consistency, and Scale Invariance.*

Observe that Dummy does not appear in this statement.

Recall that the Aumann-Shapley method fails Monotonicity. If we replace Weak Consistency by Monotonicity in the above list of axioms, a new characterization of the Shapley-Shubik method emerges.

Theorem 6. *In the continuous model, the Shapley-Shubik method is the only method satisfying Additivity, Dummy, Monotonicity, Scale Invariance, and Symmetry.*

²It is not known whether the axioms in Theorem 4 characterize the Aumann-Shapley method in the continuous model.

Symmetry is not needed in Theorem 5 (just like in Theorem 4); if we remove it in Theorem 6, a characterization of the random order values obtains.

Ordinality and, to a lesser degree, Scale Invariance, allow us to deal with division problems where the individual demands are not interpersonally comparable. This is a fine help in many examples, and a great asset of the full responsibility approach. In the next section, we introduce a radically different approach, relying crucially on interpersonal comparisons of demands.

6 Partial responsibility: in demands only

Most public utilities routinely cross-subsidize among users with a different impact on costs. The same price is charged to deliver mail, or water, to a rural or an urban domestic address; the universal service constraint for telephone implies, among other things, that the connecting charge to a residential customer is the same whether the house is pre-wired or not; special transportation services are offered to handicapped persons at the same price as public transportation for non-handicapped persons. The underlying ethical principle is that individuals are responsible for their own demand, but not for cost asymmetries, because the latter are beyond their control. The farmer should not pay more for his mail, because he cannot farm in town, the resident is not responsible for the location of the water treatment facility, the handicapped person is unable to use the regular bus but should not be penalized for it, and so on. See Fleurbaey and Trannoy (1998) for further discussion of cross-subsidization in the provision of services to geographically dispersed agents.

Formally, this viewpoint requires that, irrespective of the asymmetries of the cost function C , individual demands and cost shares be *co-monotonic*: whenever i demands as much as j , i 's cost share is not smaller than j 's. Thus individual demands must be measured on a common scale, allowing interpersonal comparisons. In the list of examples in Section 1, this will be true for the exploitation of common property resources (substitutable inputs), the buyers cooperative (substitutable demands), and in the examples of congestion and queues, capacity sharing, and cost sharing on graphs. On the other hand, such comparisons will not be meaningful when inputs are labor and capital (as in example 4 in Section 5.2), when inputs are idiosyncratic (as in the producers cooperatives example), or when demands are not comparable (as in Lee's (2002) highway maintenance example where each user is a

different category of vehicles).

Moulin and Sprumont (2004, 2005) develop a *partial responsibility theory* of cost sharing in which the central requirement is the following.

Strong Ranking. For all C , x , and i, j , $\{x_i \leq x_j\} \Rightarrow \{\varphi_i(C, x) \leq \varphi_j(C, x)\}$.

Both the Shapley-Shubik and the Aumann-Shapley methods satisfy only the weaker property that demands and cost shares are co-monotonic when the cost function is symmetric. Strong Ranking says that agents who ask more should pay more, regardless of the cost function. Thus, users are not held responsible for the unequal contributions of their own demands to the total cost, even if the cost function is additively separable. In particular, Strong Ranking and Separability are at once incompatible, even for non additive methods. For instance, if $C(z) = c(z_1) + 2c(z_2)$ and $x_1 = x_2$, the former axiom implies $y_1 = y_2$ whereas the latter requires $y_2 = 2y_1$. A fortiori, Strong Ranking is incompatible with the Stand-Alone principle or Unambiguous Externalities.

The simplest additive methods meeting Strong Ranking are the *quasi-proportional* methods. Choose a non-decreasing real-valued “weight” function β such that $0 < \beta(1)$ and define

$$\varphi_i^\beta(C, x) = \frac{\beta(x_i)}{\sum_{j \in N} \beta(x_j)} C(x)$$

for all C , x , and i .

These crude methods only use the cost data at the actual demand profile x , ignoring any “counterfactual” information about what the cost would have been at other levels of demands. This is why they are so important in practice. The two most common weight functions are $\beta(x_i) = x_i$ for all x_i , giving the *proportional method* $\varphi_i(C, x) = \frac{x_i}{\sum_{j \in N} x_j} C(x)$, and $\beta(x_i) = 1$ for all x_i , giving the *egalitarian method* $\varphi_i(C, x) = \frac{1}{n} C(x)$.

The egalitarian method is ordinal. In both the discrete and the continuous models, there is no other method meeting Strong Ranking and Ordinality, a very simple counterpart to Theorem 3 in the “partial responsibility” context³.

The proportional method meets No Merging and Splitting. In both models, there is no other method meeting Strong Ranking and No Merging and

³Note that the statement does not require Additivity. The proof is elementary: given C, x with $x_1 \leq x_2$, we can pick τ_1 such that $x_1^{\tau_1} \geq x_2$, ensuring $y_1 = y_2$.

Splitting, again a simple counterpart to Theorem 4 under “partial responsibility”⁴.

All quasi-proportional methods are monotonic by virtue of the monotonicity of the functions s and C . In fact they meet a much more demanding form of responsiveness to shifts in individual demands. Whenever user i raises her demand x_i , the total cost share of any coalition S containing her cannot decrease.

Strong Group Monotonicity. For all C, x, x' , any nonempty $S \subseteq N$ and $i \in S$, $\{x_i < x'_i \text{ and } x_j = x'_j \text{ for all } j \in N \setminus i\} \Rightarrow \{\sum_{j \in S} \varphi_j(C, x) \leq \sum_{j \in S} \varphi_j(C, x')\}$.

This axiom rules out manipulations whereby a coalition S forms, one of its members artificially raises her demand, and side-payments are performed within S so that all members of the coalition end up paying less. Clearly, Strong Group Monotonicity also rules out the profitability of a coordinated increase by *several* members of S followed by side-payments.

Like Strong Ranking, Strong Group Monotonicity is incompatible with the full responsibility approach: the latter cannot avoid the kind of coalitional maneuvers just described.

For a natural characterization of the quasi-proportional methods, consider how a shift in a user’s demand affects the cost shares of other users. All quasi-proportional methods satisfy a powerful cross effect property: a raise in an agent’s demand always affects other users’ shares in the same direction.

Solidarity. For all C, x, x' , and i , $\{x_j = x'_j \text{ for all } j \in N \setminus i\} \Rightarrow \{\varphi_j(C, x) \leq \varphi_j(C, x') \text{ for all } j \in N \setminus i, \text{ or } \varphi_j(C, x) \geq \varphi_j(C, x') \text{ for all } j \in N \setminus i\}$.

Solidarity bears on all cost functions but the common direction of the cross-effects is left unspecified⁵. The axiom is a particular instance of the widely applicable principle of solidarity (Thomson, 1999): because agents should bear responsibility only for their own actions, they should be treated “similarly” when a change occurs for which none of them is responsible. A

⁴The statement is strictly true in the discrete model. It follows at once for any C, x by splitting each agent i into x_i agents demanding one unit each. In the continuous model, the same argument establishes that the method is proportional for every x with rational components ($x_i \in \mathbb{Q}$). This extends to all real demands if we assume that $\varphi(C, x)$ is continuous in x .

⁵Unlike in Unambiguous Externalities, an axiom that bears no logical or ethical relation to Solidarity.

minimal condition of “similar” treatment is that nobody benefit from such a change if someone else suffers from it.

Solidarity is related to Strong Group Monotonicity: an additive method meeting Solidarity and Monotonicity must satisfy Strong Group Monotonicity as well.

Additivity, Strong Ranking, Solidarity, and Monotonicity essentially characterize the quasi-proportional methods. The precise result makes use of two additional conditions.

Zero Charge for Zero Demand. For all C , x , and i , $\{x_i = 0\} \Rightarrow \{\varphi_i(C, x) = 0\}$.

Positive Cost for Positive Demand. For all C , x , and i , $\{x_i > 0 \text{ and } C(x) > 0\} \Rightarrow \{\varphi_i(C, x) > 0\}$.

Theorem 7. *In the discrete model and provided $n \geq 3$, the quasi-proportional methods such that $\beta(0) = 0$ are the only method satisfying Additivity, Strong Ranking, Solidarity, Monotonicity, Zero Cost for Zero Demand, and Positive Cost for Positive Demand.*

A variant of this result drops the last two conditions but uses a strict version of Solidarity. The resulting combination of axioms characterize the quasi-proportional methods such that $\beta(0) > 0$ (for instance the egalitarian method).

In the next section, we introduce a cost-sharing method translating the partial responsibility viewpoint into a much less crude formula than the quasi-proportional ones. This method uses a fair amount of counterfactual information on costs, and shares several important features of the full responsibility approach.

7 Limited liability: the serial methods

We introduce further normative requirements, valid in both the full and the partial responsibility approaches to cost sharing. These properties lead to a characterization of two important methods linking the full and partial responsibility theories.

7.1 Upper bounds on individual shares

Several of the axioms introduced in the previous sections, such as Unambiguous Externalities, Solidarity and Strong Group Monotonicity, evaluate the direction of the cross effects induced by a demand shift: when Donald raises his demand, does Daisy’s cost share go up or down? Do Daisy’s and Deborah’s shares move in the same direction? What about the sum of Donald and Daisy’s shares?

Another important consideration is the range of such cross effects. When I share the cost of a facility with a number of other users about whom I know nothing, the highest cost share I may end up paying for my own demand is an important criterion. When the size of the various individual demands may vary widely, guaranteeing a reasonable upper bound on an agent’s cost share is a way of protecting her⁶.

In that spirit, a weak requirement is that the *liability* of a user, given her demand, the cost function, and the number of other users, be finite: for all C , i , and x_i , $\sup_{x_{N \setminus i}} \varphi_i(C, (x_i, x_{N \setminus i})) < +\infty$. Note that the Shapley-Shubik and Aumann-Shapley methods fail this test. The same is true for any of the quasi-proportional methods.

Next we place an explicit upper bound on the liability of each user. Recall that e^N denotes the demand profile where each agent in N demands one unit. From the monotonicity of C , we know that $C(x) \leq \sum_{i \in N} C(x_i e^N)$ for all x . The following bounds are therefore feasible.

Universal Upper Bounds. For all C , x , and i , $\varphi_i(C, x) \leq C(x_i e^N)$.

These upper bounds are plausible only when demands are comparable, which is not always the case in the full responsibility approach. An important variant applies only to the much smaller domain of symmetric supermodular cost functions, but imposes tighter bounds.

Unanimity Upper Bounds. For all C , x , and i , $\{C \text{ symmetric and supermodular}\} \Rightarrow \{\varphi_i(C, x) \leq \frac{1}{n} C(x_i e^N)\}$.

If C is symmetric and all agents other than i demand x_i as well, Symmetry imposes that everyone including i pays $\frac{1}{n} C(x_i e^N)$. This charge is called the “unanimity” cost share of agent i , hence our terminology.

⁶In the output-sharing context, the same concern leads to place a *lower* bound on the share of every agent. The corresponding axioms do not have the same mathematical “bite” as the upper bounds described below. Thus the results in Section 7.3 do not have an obvious counterpart in that context.

7.2 Two serial formulas

We begin by defining the serial cost-sharing method for symmetric cost functions only. Fix a demand profile x such that $x_1 \leq x_2 \leq \dots \leq x_n$, and construct the related profiles x^k as follows: $x^0 = 0$, $x^k = (x_1, x_2, \dots, x_{k-1}, x_k, \dots, x_k)$ for $k = 1, \dots, n$. Note that $x^k \leq x^{k+1}$ for all k , and $x^n = x$. Let

$$\varphi_i^{ser}(C, x) = \sum_{k=1}^i \frac{C(x^k) - C(x^{k-1})}{n - k + 1} = \frac{C(x^i)}{n - i + 1} - \sum_{k=1}^{i-1} \frac{C(x^k)}{(n - k)(n - k + 1)} \quad (4)$$

for all i and all symmetric C . The formula is then extended to all profiles x by Symmetry.

The summation in (4) is quite similar to that in (3). Indeed for the capacity cost function $C(z) = c(\max_N z_i)$, we have $C(x^k) = c(x_k)$ for all k , therefore the serial methods coincide with any symmetric fixed flow method. The same is true for the production function $F(z) = g(\min_N z_i)$, as $F(x^k) = g(x_1)$ for all k .

A fairly simple though not entirely trivial observation is that a method given by (6) for symmetric cost functions, must meet the Unanimity Upper Bound.

There are two canonical extensions of the serial formula (4) to the full domain of cost functions: they follow respectively the partial and full responsibility approaches.

7.2.1 Partial responsibility

Consider the method defined by formula (4) for *every* cost function C , symmetric or not. We call it the *cross-subsidizing serial* method. It is clearly additive and satisfies Strong Ranking. Note that it also meets the Universal Upper Bounds.

Proposition 1. *In the discrete model, the cross-subsidizing serial method is characterized by the combination of Additivity, Strong Ranking, and formula (4) for symmetric cost functions.*

Of course the cross-subsidizing serial method is defined in exactly the same way in the continuous model, and meets the same properties. Its characterization by the three properties above is a plausible conjecture.

7.2.2 Full responsibility

Our second extension, the *subsidy-free serial* method, has two slightly different definitions in the discrete and the continuous models.

In the discrete model, the natural extension of (4) is the symmetric fixed flow method of which the flow is concentrated on the “diagonal” of $\mathbb{N}^N$. For two agents, this fixed flow is depicted on Figure 1; Figure 2 shows how the cost shares are obtained when $(x_1, x_2) = (2, 3)$ by projecting the fixed flow on $[0, (2, 3)]$: agent 1 pays $\frac{1}{2}(\partial_1 C(0, 0) + \partial_1 C(1, 1) + \partial_1 C(0, 1) + \partial_1 C(1, 2))$. More generally, assuming without loss of generality $x_1 \leq x_2 \leq \dots \leq x_n$, user i pays essentially the sum of her marginal costs along a path linking $x^1, x^2, \dots, x^i$ by means of “diagonal” segments.

For a precise definition recall, as noted before Theorem 2, that a fixed-flow method is entirely determined once the flow $f(., ke^N)$ is defined for all $k \in \mathbb{N}$. Call such a flow *symmetric* if $f(\pi z, ke^N) = \pi f(z, ke^N)$ for all $z \in [0, ke^N[$ and every permutation π on N . Call it *diagonal* if $f(z, ke^N) \neq 0$ only if $|z_i - \frac{z(N)}{n}| < 1$ for all $i \in N$: the “support” of a diagonal flow to ke^N is included in the union of the unit cubes $[0, e^N], [e^N, 2e^N], \dots, [(k-1)e^N, ke^N]$. Because there is clearly only one symmetric flow to the one-unit-demand profile e^N , it follows that there is also a unique symmetric diagonal flow to ke^N : we denote it $f^{ser}(., ke^N)$. The subsystem $\{f^{ser}(., ke^N) | k \in \mathbb{N}\}$ completely determines a unique fixed-flow system, which we denote f^{ser} . The subsidy-free serial method $\varphi^{f^{ser}}$ is the cost-sharing method represented by this symmetric diagonal fixed-flow system.

In the continuous model, the subsidy-free method computes an agent’s cost share by integrating her marginal cost along the path to x obtained by projecting the diagonal of $\mathbb{R}_+^N$ onto $[0, x]$. Assuming without loss $x_1 \leq x_2 \leq \dots \leq x_n$ and setting $x_0 = 0$, this means

$$\varphi_i^{f^{ser}}(C, x) = \sum_{j=1}^i \int_{x_{j-1}}^{x_j} \partial_i C(x_1, x_2, \dots, x_{j-1}, t, \dots, t) dt$$

for all i and C .

It is easy to check that the subsidy-free serial method meets Monotonicity and the Universal Upper Bounds.

Proposition 2. *In the discrete and continuous models, the subsidy-free serial method is characterized by the combination of Additivity, Dummy, Monotonicity, and formula (4) for symmetric cost functions.*

7.2.3 Comparing the two methods

In view of their common origin, the cross-subsidizing and subsidy-free serial methods form a bridge between the partial and full responsibility theories of cost sharing. In particular, the cross-subsidizing serial method meets the following weak version of Separability.

Weak Separability. For all C and x , $\{C(z) = \sum_{i \in N} c(z_i) \text{ for all } z\} \Rightarrow \{\varphi_i(C, x) = c(x_i) \text{ for all } i\}$.

This means that if the cost function is not only additively separable but also symmetric, each agent pays her own separable cost. The axiom is very appealing from the partial responsibility viewpoint. Since the purpose of cross-subsidization is only to correct for cost asymmetries, subsidization is not justified when the cost function is symmetric: in such cases, the separability principle should still apply. All quasi-proportional methods violate Weak Separability: they arguably perform too much cross-subsidization.

Another common feature is the responsiveness of the two methods to shifts in individual demands. Both methods fail Strong Group Monotonicity⁷ but both meet the following compromise between that axiom and Monotonicity.

Group Monotonicity. For all C , all x, x' , and any nonempty $S \subseteq N$, $\{x_i < x'_i \text{ for all } i \in S \text{ and } x_i = x'_i \text{ for all } i \in N \setminus S\} \Rightarrow \{\varphi_i(C, x) \leq \varphi_i(C, x') \text{ for at least one } i \in S\}$.

A violation of this axiom leads to the possibility of coordinated strategic inflation of demands: all agents in some group could pay a strictly smaller individual cost share by agreeing to simultaneously inflate their demands. Such manipulations do not require side-payments: Group Monotonicity is therefore a compelling incentive-compatibility requirement whenever agents can communicate easily.

The axiom has much bite in the full responsibility approach, where it eliminates most fixed-flow methods. In particular, the Shapley-Shubik method is monotonic, but not group-monotonic. Let $N = \{1, 2, 3\}$ and C be the cost

⁷Recall that among additive methods, Strong Group Monotonicity is incompatible with Dummy. It is also incompatible with the combination of Weak Separability and Strong Ranking (Proposition 7 in Moulin and Sprumont, (2005)).

function:

$$C(z) = \begin{cases} 1 & \text{if } z \geq (2, 0, 1) \text{ or } z \geq (1, 1, 1) \text{ or } z \geq (0, 2, 1), \\ 0 & \text{otherwise.} \end{cases}$$

Check that $\varphi_1^{ShSh}(C, (1, 1, 1)) = \varphi_2^{ShSh}(C, (1, 1, 1)) = \frac{1}{3} > \frac{1}{6} = \varphi_1^{ShSh}(C, (2, 2, 1)) = \varphi_2^{ShSh}(C, (2, 2, 1))$, which contravenes Group Monotonicity.

7.3 Characterizing the serial methods

The upper bounds we defined in Section 7.1 are a key to characterizing the serial methods. Starting with the partial responsibility approach, we note that while no quasi-proportional method guarantees finite liability (that is, $\sup_{x_{N \setminus i}} \varphi_i(C, (x_i, x_{N \setminus i})) < +\infty$ for all C , i , and x_i), there are methods meeting Solidarity and Strong Group Monotonicity for which the Universal Upper Bounds hold. Here is an example: for every C and x , divide $C(x)$ equally among all agents whose demand is largest. In view of Theorem 6, it is not surprising that this example can be viewed as a limit of quasi-proportional methods, where the ratios $\frac{\beta(z+1)}{\beta(z)}$ are arbitrarily large.

But the combination of the Upper Bounds and Weak Separability is very powerful. For our last two results, we introduce one axiom that is a considerable weakening of both the Universal and the Unanimity Upper Bounds.

Weak Upper Bounds. For all C , x , and i , $\{C(z) = c(z(N)) \text{ for all } z \text{ and } c \text{ convex}\} \Rightarrow \{\varphi_i(C, x) \leq c(nx_i) = C(x_i e^N)\}$.

Theorem 8. *In the discrete model, the cross-subsidizing serial method is the only method satisfying Additivity, Strong Ranking, Weak Separability, and Weak Upper Bounds.*

Turning to the full responsibility approach, we first note that many fixed-flow methods guarantee finite liability. The Upper Bounds are much more demanding.

Theorem 9. *In the discrete model, the subsidy-free serial method is the only method satisfying Additivity, Dummy, Monotonicity, Symmetry, and Weak Upper Bounds.*

In the continuous model, it is the only method satisfying Additivity, Dummy, Monotonicity, Symmetry, and Universal Upper Bounds.

It is not clear whether Weak or Unanimity Upper Bounds can replace Universal Upper Bounds in the second statement of this theorem.

We conclude with an alternative characterization based on Group Monotonicity. In the discrete model, this axiom leads to small family of methods based on a fixed flow concentrated “near” the diagonal. Given $k \in \mathbb{N}$, call a flow $f(\cdot, ke^N)$ *nearly diagonal* if $f(z, ke^N) \neq 0$ only if $|z_i - \frac{z^{(N)}}{n}| \leq 1$ for all $i \in N$. A subsidy-free *nearly serial* method is a cost-sharing method represented by a symmetric nearly diagonal fixed-flow system.

Figure 3 depicts a two-agent example. Each method is in fact fully characterized by a sequence $\{\alpha_r\}_{r \in \mathbb{N} \setminus \{0\}}$ in $[0, 1]$, where α_r is the fraction of the total (unit) fixed flow that goes through re^N .

Theorem 10. *In the discrete model, the subsidy-free nearly serial methods are the only fixed-flow methods satisfying Symmetry and Group Monotonicity.*

In view of Theorem 2, the subsidy-free nearly serial methods are characterized by the combination of *Additivity, Unambiguous Externalities, Symmetry and Group Monotonicity*.

We conjecture that the same axioms characterize the subsidy-free serial method in the continuous model.

8 Bibliographical notes

Section 2. The continuous version of the cost-sharing model with variable demands was developed by Billera and Heath (1982), Mirman and Tauman (1982), and Samet and Tauman (1982). The discrete model is more recent: the early contributions are Moulin (1995), who formulated the Monotonicity axiom, and van den Nouweland et al. (1995).

Section 3. Theorem 1 is Shapley’s (1953) classic characterization of the value for cooperative games with side-payments.

The characterization of Additivity and Dummy in terms of flows, given in formula (2), is due to Moulin and Vohra (2003). They built upon a result of Wang (1999) describing these methods as convex combinations of path-generated methods. A characterization parallel to Wang’s is proved in Haimanko (2000) and Friedman (2004) for the continuous model.

The fixed-flow methods were first defined in Moulin and Sprumont (2005), who proved a variant of Theorem 2. A slightly different characterization is offered in Sprumont (2004).

Section 4. Moulin (1995) proposed the discrete version of Ordinality and proved the corresponding version of Theorem 3. Proposition 5 in Moulin

(1995) involves Monotonicity but the axiom is not needed in the proof. The continuous version of the theorem is due to Sprumont (1998).

Section 5. The formal expression of the Aumann-Shapley method in the discrete model can be found in Moulin (1995). The axiomatization stated in Theorem 4 is due to Sprumont (2005).

The large literature on the continuous Aumann-Shapley method is rooted in Aumann and Shapley's (1974) theory of value for nonatomic games. The economic reinterpretation of the theory, the formulation of the Scale Invariance and Weak Consistency axioms, as well as the characterization of the diagonal formula offered in Theorem 5, are due to Billera and Heath (1982), Mirman and Tauman (1982), and Samet and Tauman (1982). A quite different axiomatization is proposed in Young (1985). For further references, see Tauman (1988) or Haimanko and Tauman (2002).

The characterization of the Shapley-Shubik method based on Monotonicity and Scale Invariance, Theorem 6, is a slight variant of Theorem 1 in Friedman and Moulin (1999).

Section 6. The partial responsibility approach to cost sharing is inspired by the more general discussion of responsibility in distributive justice initiated by Fleurbaey (2004). Variants of the Strong Ranking axiom have appeared in a number of papers concerned with methods restricted to the case of additively separable production functions: see Fleurbaey (1995), Bossert and Fleurbaey (1996). Moulin and Sprumont (2004) are the first to handle full-fledged externalities. Theorem 7 is in Moulin and Sprumont (2005). Solidarity, of course, is an old idea: see Thomson (1999) for references.

Section 7. The earliest instance of an upper or lower bound in cost sharing is the Stand Alone bound, discussed first by Shubik (1962). The Unanimity bound is a more recent idea (Moulin (1990)). Shenker (1995, circulated 1989) invented the serial formula for the case of perfect substitutes, $C(z) = c(z(N))$; in this domain Moulin and Shenker (1992, 1994) offer characterizations based respectively on incentive-compatibility and Weak Upper Bounds.

The general subsidy-free formula appears in Moulin (1995) for the discrete model, and in Friedman and Moulin (1999) for the continuous one, where the corresponding statements in Theorem 9 are proven (with a redundant cross-monotonicity requirement in the former). Theorem 10 is the central result of Sprumont (2004).

The cross-subsidizing serial method is due to Sprumont (1998). Its characterization in Theorem 8 is easily derived from the results in Moulin and

Sprumont (2005) (in particular Lemma A.2). The same applies to Propositions 1 and 2.

References

- [1] D. Aadland and V. Kolpin (2004). Environmental determinants of cost sharing, *J. Econ. Beh. Org.* 53, 495-511, 2004.
- [2] A. Archer, J. Feigenbaum, A. Krishnamurthy, R. Sami, and S. Shenker (2004). Approximation and collusion in multicast costsharing, *Games Econ. Beh.* 47, 36-71.
- [3] R.J. Aumann and L. Shapley (1974). *Values of Nonatomic Games*. Princeton: Princeton Univ. Press.
- [4] L. Billera and D. Heath (1982). Allocation of shared costs: A set of axioms yielding a unique procedure, *Math. Oper. Res.* 7, 32-39.
- [5] L. Billera, D. Heath, and J. Raanan (1978). Internal telephone billing rates: A novel application of nonatomic game theory, *Oper. Res.* 26, 956-965.
- [6] W. Bossert and M. Fleurbaey (1996). Redistribution and compensation, *Soc. Choice Welfare* 13, 343-356.
- [7] J.H. Case (1979). *Economics and the Competitive Process*, New York: New York University Press.
- [8] A. Castano-Pardo and A. Garcia-Diaz (1995). Highway cost allocation: an application of the theory of non-atomic games, *Transport. Res.* 29A, 187-203.
- [9] Y-J. Chen and J. Zhang (2005). Design of Price mechanisms for network resource allocation via price of anarchy, mimeo, New york University.
- [10] D. Clarke, S. Shenker, and L. Zhang (1992). Supporting real-time applications in an integrated services packet network, *Comput. Comm. Rev.* 22, 14-26.

- [11] A. Demers, S. Keshaw and S. Shenker (1990). Analysis and simulation of a fair queuing algorithm, *Internetworking: Research and experience*, 1, 3-26.
- [12] J. Feigenbaum, C. Papadimitriou, and S. Shenker (2001). Sharing the cost of multicast transmissions, *J. Computer and Syst. Sc.* 63, 21-41.
- [13] S.D. Flam and A. Jourani (2003). Strategic behavior and partial cost sharing, *Games and Econ. Beh.* 43, 44-56.
- [14] M. Fleurbaey (1994). On fair compensation, *Theory Dec.* 36, 277-307.
- [15] M. Fleurbaey (1995). Three solutions for the compensation problem, *J. Econ. Theory* 65, 505-521.
- [16] M. Fleurbaey and F. Maniquet (1996). Cooperative production, a comparison of welfare bounds, *Games and Econ. Beh.* 17, 200-208.
- [17] M. Fleurbaey and A. Trannoy (1998). La péréquation territoriale en question, *Flux* 31, 91-98.
- [18] E. Friedman (2004). Paths and consistency in additive cost sharing, *Int. J. Game Theory* 32, 501-518.
- [19] E. Friedman and H. Moulin (1999). Three methods to share joint costs or surplus, *J. Econ. Theory* 87, 275-312.
- [20] O. Haimanko (2000). Value theory without symmetry, *Math. Oper. Res.* 25, 573-590.
- [21] O. Haimanko and Y. Tauman (2002). Recent developments in Aumann-Shapley pricing, Part 1: The nondifferentiable case, forthcoming in *Game Theory and Resource Allocation: The Axiomatic Approach*, ed. by W. Thomson, NATO ASI Series.
- [22] G. Hardin (1968). The tragedy of the commons, *Science* 162, 1243-1248.
- [23] M. Haviv (2001). The Aumann-Shapley price mechanism for allocating congestion costs, *Oper. Res. Letters* 29, 211-215.
- [24] N. Immorlica, M. Mahdian, and V.S. Mirrokni (2005). Limitations of cross-monotonic cost sharing schemes, in *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, 602-605.

- [25] L.D. Israelsen (1980). Collectives, communes and incentives, *J. Comp. Econ.* 4, 99-124.
- [26] R. Johari and J. Tsitsiklis, Efficiency loss in a network resource allocation game, *Maths of Operations Research*, 29, 3, 407-435, 2004.
- [27] R. Johari and J. Tsitsiklis (2005). A scalable network resource allocation mechanism with bounded efficiency loss, mimeo, Stanford and MIT.
- [28] R. Johari, S. Mannor and J. Tsitsiklis (2005). Efficiency loss in a network resource allocation game: the case of elastic supply, *IEEE Transactions on Automatic Control*, forthcoming.
- [29] B.G. Ju, E. Miyagawa, and T. Sakai (2005). Non-manipulable division rules in claims problems and generalizations, *J. Econ. Theory*, forthcoming.
- [30] M. Koster, S. Tijs, and P. Borm (1998). Serial cost sharing methods for multi-commodity situations, *Math. Soc. Sci.* 36, 229-242.
- [31] D. Lee (2002). Game-theoretic procedures for determining pavement thickness and traffic lane costs in highway cost allocation, Ph. D. thesis, Texas A&M University.
- [32] S. Littlechild and G. Owen (1973). A simple expression for the Shapley value in a simple case, *Manage. Sci.* 20, 370-372.
- [33] L. Mirman and Y. Tauman (1982). Demand compatible equitable cost sharing prices, *Math. Oper. Res.* 7, 40-56.
- [34] D. Monderer and A. Neyman (1988). Values of nonatomic games: The method of multilinear approximation, in *The Shapley Value: Essays in Honor of Lloyd Shapley*, ed. by A. Roth. Cambridge: Cambridge University Press, 217-234.
- [35] H. Moulin (1990). Joint ownership of a convex technology: comparison of three solutions, *Rev. Econ. Stud.* 57, 439-452.
- [36] H. Moulin (1994). Serial cost-sharing of an excludable public good, *Rev. Econ. Stud.* 61, 305-325.

- [37] H. Moulin (1995). On additive methods to share joint costs, *Japan. Econ. Rev.* 46, 303-332.
- [38] H. Moulin (2002). Axiomatic cost and surplus-sharing, in *Handbook of Social Choice and Welfare*, ed. by K. Arrow, A. Sen and K. Suzumura, Amsterdam: North-Holland, 289-357.
- [39] H. Moulin (2005). The price of anarchy of serial cost sharing and other methods, mimeo, Rice University.
- [40] H. Moulin and S. Shenker (1992). Serial cost sharing, *Econometrica* 60, 1009-1037.
- [41] H. Moulin and S. Shenker (1994). Average cost pricing versus serial cost sharing: An axiomatic comparison, *J. Econ. Theory* 64, 178-201.
- [42] H. Moulin and S. Shenker (2001). Strategy-proof sharing of submodular costs: budget-balance versus efficiency, *Economic Theory* 87, 275-312.
- [43] H. Moulin and Y. Sprumont (2004). Responsibility and cross-subsidization in cost sharing, *Games and Econ. Beh.*, forthcoming.
- [44] H. Moulin and Y. Sprumont (2005). On demand responsiveness in additive cost-sharing, *J. Econ. Theory* 125, 1-35.
- [45] H. Moulin and R. Vohra (2003). A representation of additive cost sharing methods, *Econ. Letters*, 80, 399-407.
- [46] A. van den Nouweland, J. Potters, S. Tijs, and J. Zarzuelo (1995). Cores and related solution concepts for multi-choice games, *ZOR-Math. Meth. Oper. Res.* 41, 289-311.
- [47] J. Roemer and J. Silvestre (1993). The proportional solution for economies with both private and public ownership, *J. Econ. Theory* 59, 426-444.
- [48] Samet, D. and Y. Tauman (1982). The determination of marginal cost prices under a set of axioms, *Econometrica* 50, 895-909.
- [49] D. Samet, Y. Tauman, and I. Zang (1984). An application of the Aumann-Shapley prices to transportation problems, *Math. Oper. Res.* 10, 25-42.

- [50] S. Sanghavi and B. Hajek (2004). Optimal allocation of a divisible good to strategic buyers, *Proceedings of the 43d IEEE conference on Decision and Control*.
- [51] L. Shapley (1953). A value for n -person games, in *Contributions to the Theory of Games II* ed. by H.W. Kuhn and W. Tucker, Ann. Math. Stud. 28, Princeton: Princeton Univ. Press.
- [52] W. Sharkey (1982). *The Theory of natural Monopoly*, Cambridge: Cambridge Univ. Press.
- [53] S. Shenker (1995). Making greed work in networks: a game-theoretic analysis of gateway service discipline, *IEEE/ACM Trans. Networking* 3, 819-831.
- [54] M. Shubik (1962). Incentives, decentralized control, the assignment of joint costs, and internal pricing, *Manage. Sci.* 8, 325-343.
- [55] Y. Sprumont (1998). Ordinal cost sharing, *J. Econ. Theory* 81, 126-162.
- [56] Y. Sprumont (2004). Nearly serial sharing methods, mimeo, Université de Montréal.
- [57] Y. Sprumont (2005). On the discrete version of the Aumann-Shapley cost-sharing method, *Econometrica* 73, 1693-1712.
- [58] Y. Tauman (1988). The Aumann-Shapley prices: A survey, in *The Shapley Value: Essays in Honor of Lloyd Shapley*, ed. by A. Roth. Cambridge: Cambridge University Press, 279-304.
- [59] W. Thomson (1999). Welfare-domination under preference replacement: A survey and open questions, *Soc. Choice Welfare* 16, 373-394.
- [60] Y. Wang (1999). The additivity and dummy axioms in the discrete cost sharing model, *Econ. Letters* 64, 187-192.
- [61] S. Yang and B. Hajek (2005). Revenue and stability of a mechanism for efficient allocation of a divisible good, mimeo, University of Illinois at Urbana Champaign.
- [62] H.P. Young (1985). Producer incentives in cost allocation, *Econometrica* 53, 757-765.