
This paper applies principles from game theory to the problem of allocating the cost of a shared facility, such as a pipeline. The theory of cooperative games strongly suggests that no method exists for allocating costs that will achieve all major policy goals. We apply results from the theory of cooperative games and principles of cost allocation to assess some commonly adopted rules for allocating costs and defining unit charges. Most notably, the postage-stamp toll is found to fail a minimal set of commonly applied principles.

Cet article applique les principes tirés de la théorie des jeux au problème de la répartition des coûts d'une installation partagée telle qu'un pipeline. La théorie des jeux coopératifs suggère fortement qu'il n'existe pas de méthode de répartition des coûts qui puisse satisfaire tous les objectifs principaux en matière de politique. Nous appliquons les résultats tirés de la théorie des jeux coopératifs et des principes de répartition des coûts pour évaluer certaines règles d'usage adoptées pour répartir les coûts et définir les frais unitaires. En particulier, il ressort que le droit timbre-poste ne satisfait pas à un ensemble minimal de principes d'usage.

Cost-Allocation Principles for Pipeline Capacity and Usage

D.J. SALANT and G.C. WATKINS

I. Introduction

Transmission facilities, such as pipelines, lead to debates about cost sharing whenever there are multiple users of large segments. The cost-allocation literature strongly suggests that there exists no way of allocating pipeline costs which is immune to criticism. And a system of uniform rates (postage-stamp rates), for example, is no exception.¹ Our intent in this paper is two fold. First, to outline some of the main principles that most would agree a cost-allocation system should serve to satisfy the oft-cited statutory admonition of being "fair and reasonable." Second, to explain the implications of those principles.

To provide suitable focus initially we discuss postage-stamp systems in the context of a natural gas pipeline system and explain its pros and cons. Then we take a more analytical approach, but with no predetermined bias as to what constitutes the optimal way in which to allocate pipeline network costs among users. Instead, we work from first principles. Over the past decade or so there have been developments in economic techniques that apply notions of fairness and equity, as well as

D.J. Salant and G.C. Watkins are with the Law & Economics Consulting Group, Emeryville, California.

1/ Postage-stamp rate design has been applied by Nova Gas Transmission Limited in the Province of Alberta, Canada for most of that system's life.

efficiency, to cost allocation. There has been increased recognition of the need to look at the role that various fairness criteria play in allocating costs. Our paper makes it apparent how some current schemes such as postage-stamp rates can conflict with commonly accepted fairness standards.

We examine the so-called axiomatic approach in order analytically to examine alternative concepts of fairness, or "just and reasonable," for determining how to allocate costs. Our analysis is based on axiomatic social choice theory developed over the past twenty or so years, and in particular on axiomatic cost-allocation theory. Axiomatic cost-allocation approaches have been applied to water systems, airport landing fees, managerial accounting, flood control, navigation, and power systems. We apply this theory to identify a formula for allocating costs. We find that a postage-stamp rate generally fails to pass most commonly used standards for fairness and reasonableness, and could induce both inefficient bypass and inefficient resource development. Application of the axiomatic approach can provide some assurance that hidden implications of commonly proposed notions of fairness have not been overlooked.

The paper is organized as follows. Section II briefly discusses postage-stamp schemes. Basic cost-allocation and fairness principles are outlined in Section III. Additional fairness criteria are discussed in Section IV, including the nucleolus and the Shapley value. Section V discusses how the Shapley value can be used as a guide for cost allocation. Section VI addresses other equity and efficiency issues. Section VII is a summary.

II. Postage-Stamp Schemes

In North America, regulated pipeline tolls are normally set to yield a total revenue requirement. There is typically some latitude for the regulator in determining how these tolls are set. These may consist of fixed and variable charges, be distance related, or fixed within zones, or may be uniform throughout: the so-called postage-stamp system.

A postage-stamp system is one in which all

users pay the same amount per unit, or parcel, of capacity, independent of transport distances. This type of rate structure is most appropriate when: (a) there are high fixed connection costs, so that the total costs are not so distance-sensitive; (b) there is little variation in the distances among the different users' shipments; (c) there are large transaction costs associated with distance-related tolls² when users have similar average distances of haul; and (d) when system complexity and cost interdependence make cost causation nebulous.

However, a postage-stamp tariff is inherently inefficient if total costs are distance-sensitive and/or if there is a significant variation in the sources of demand.³ For instance, if one user wishes to use only a small part of the pipeline and many others use most of its entire length, the stand-alone cost of the one short-haul firm could be much less than $1/n^{\text{th}}$ of the pipeline cost, where there are n firms that use it in total. This situation would encourage a potential contributor to the network costs to incur the cost of building bypass facilities. Such incentives can persist even if these bypass facilities were more costly than the incremental cost of allowing the short-haul firm access to, and use of, the pipeline system. And it is here that bypass is inefficient.

Furthermore, even when the postage-stamp rates do not initially create incentives for inefficient bypass, circumstances can change, which could cause such incentives to emerge. Technology can change, new fields can come on line, and a host of other factors can alter demand patterns in such a way as to create incentives for inefficient bypass. Moreover, the rate structure can affect incentives to bring new areas on line in the first place. We introduce principles for cost allocation that take into account the possibility that conditions can change over time.

2/ This may be manifest in high administrative savings for the utility itself or for the users of the facility with a postage-stamp regime.

3/ Also, see the later discussion of the indirect impacts of postage-stamp rates on the efficiency of resource allocation.

III. Basic Cost-Allocation Principles

We start by considering two commonly accepted minimal properties that a cost allocation should satisfy: (a) the stand-alone cost test; and (b) the incremental cost test. We explain why some simple approaches, such as the postage-stamp system, fail to meet these two principles.

The *stand-alone cost* test has two parts. First, it requires that the cost share borne by each user not exceed that user's stand-alone costs. If the proposed cost-allocation rule satisfies the condition that no single pipeline user can do better on his or her own than under the proposed cost allocation, then it satisfies the *individual rationality condition*. The second part of the stand-alone cost test applies to *subsets* or *groups of users*; it requires that the cost allocation satisfy a *group rationality condition*. The group rationality condition requires that no *group* of users be able to self-supply for less than their combined costs under the proposed allocation rule. If an allocation fails the stand-alone cost test for *any* coalition, or group of users, then any such group would have an incentive to bypass the system and self-provide. Together, the individual and group rationality conditions constitute the *stand-alone cost test*. The stand-alone cost test is a condition for all parties to cooperate voluntarily and use the system. It also means that each user will find it individually rational to remain on the system and pay his assigned cost share.

The other minimal condition for fairness is the *incremental cost test*. This test is satisfied if no single group of users is subsidizing another. The incremental cost test also means that the allocation of costs to any group of users must be at least as large as the incremental costs of including that group on the system. Both the stand-alone and incremental cost tests are *equity*, or *fairness*, conditions.

The incremental cost test is equivalent to a stand-alone cost test whenever joint costs are fully allocated. When costs are fully allocated and the cost allocation fails a stand-alone test, it is necessarily the case that cross-subsidies exist, in the sense that one group's contribution to the total costs is less than the incremental

costs of serving it. To see this, suppose there are two groups of pipeline users, and that the costs allocated to the first group were to exceed its stand-alone costs. The allocation of the total costs to this group and the remaining group will sum to the total system costs, assuming all costs are allocated. Thus, the costs of the entire system less the stand-alone costs of serving only the first group will then exceed the costs allocated to the second group. In other words, if the costs allocated to one group exceed its stand-alone costs, the costs allocated to everyone else are less than the incremental costs of serving them, where these are represented by the system costs less the stand-alone costs of serving everyone not in the first group.

A seemingly minimal requirement for a cost allocation is that it be fair at least in the sense that it passes both stand-alone and incremental cost tests. Then it would provide incentives for all interested parties to cooperate, would not allow cross-subsidies to exist, and would allocate all the costs among all users. The set of all such cost allocations is called the *core*. This basic, if somewhat abstract, concept is helpful in limiting discussion of how costs should be allocated and can eliminate some allocations, such as postage-stamp rates, that might otherwise seem reasonable.

Consider a simple example adapted from Young (1994), in which the costs of serving firm A alone is \$11 million, firm B alone is \$7 million, and the costs of serving the two together is \$15 million — which provides savings of \$3 million over separate systems serving each firm. Such savings are precisely what a pipeline system is intended to offer users. It is not obvious what is the right way to allocate costs or cost savings in this situation. An equal division of costs (which is how postage-stamp rates are usually set up) would set the price at \$7.5 million each, and firm B would not wish to participate in a joint project because it would be better off on its own. Thus, an equal division of costs fails the stand-alone cost test. Further, suppose that firm A uses three times the capacity that firm B does, at least over the part of the system they both use. Then a cost allocation in proportion to capacity, such as would be the case with a purely demand-re-

lated toll, would result in a price to firm A of \$11.25 million and \$3.75 million to firm B. This is another instance in which what seems to be a sensible cost-sharing rule fails to be in the core, that is, it does not pass a stand-alone test, given A's stand-alone costs of \$11 million.

However, a number of cost-sharing rules will be in the "core" for this example. An equal division of the savings above their respective stand-alone costs will result in cost shares of \$9.5 million and \$5.5 million for A and B, respectively. Division of savings in proportion to demand will result in a cost allocation of \$8.75 million and \$6.25 million. Division of savings based on opportunity, or stand-alone, costs implies a cost allocation of \$9.17 million and \$5.83 million. All three of these allocations are in the core, since both firms have allocated costs below their respective stand-alone costs. More generally, the core includes all cost allocations which fall in a particular range of values. That is, there will typically be upper and lower bounds on each firm's cost share for any cost allocation in the core. Although the core can rule out some harmless sounding cost-sharing schemes, such as equal splits of costs or postage-stamp schemes, it does not identify a unique split. Notice too that the logic of the stand-alone criteria can also be used to characterize the potential problems with cost-sharing arrangements such as a postage-stamp scheme, which indeed does fail the crucial stand-alone test.

Aside from non-uniqueness, another difficulty with the core is that it could be empty: it is possible that no cost allocation will satisfy the core conditions. Suppose, for example, the costs for a stand-alone system for each of firms A, B, or C were \$6 million, the costs of serving any two firms was \$7 million, and the costs of serving all three were \$11 million. In this case, the core would be empty.⁴ Whether or not

4/ Here, the constraints for a cost allocation to be in the core cannot all hold instantaneously. In other words, the costs allocated to any pair of firms cannot exceed \$7 million, and there are three such constraints, which in aggregate imply that the total costs allocated to the three firms cannot exceed \$10.5 million. Moreover, the \$11 million total costs must be split among the three. These two conditions are contradictory. To put it another way, simultane-

ously supplying all three is obviously most economical, and would require that each user pay \$3.67 million on average (one may pay less than \$3.67 million, but then the other two will have to pay more than \$7.33 million, or (\$11 million - \$3.67 million), violating the stand-alone cost test. In any case, one pair of buyers will end up being assessed for more than their \$7 million stand-alone costs. So there will be a pair of buyers who will prefer to build their own system rather than paying their share of the total system costs. The core is empty, as it requires that the stand-alone test hold for all coalitions as well as individuals.

IV. Additional Fairness Criteria

There are a number of other fairness conditions that a cost-allocation mechanism should probably satisfy. Not all of them can always be satisfied simultaneously. Policy makers' choice of a formula for allocating costs will depend on which fairness criteria they judge to be the most important at the time. Here, we first present a heuristic discussion of the major standards that have been analyzed in the theoretical literature. We then explain, at least in the context of a theoretically ideal world with no uncertainty and no administrative or compliance costs, how these principles can nail down specific cost-sharing formulae.

One condition we shall want to impose on a cost-allocation rule is that it "work" in changing environments. That is, the principles laid out one day should not be revised the next due to a change in circumstances. In the case of Nova Gas Transmission Limited (NGTL), the Alberta Energy and Utilities Board has upheld the postage-stamp toll with the justification that the system should encourage gas development in remote areas of Alberta. By doing so the Board, at least implicitly, made the decision that it was worth sacrificing the stand-alone cost test for the sake of this other policy objective. Over time, with the development that has occurred, the justification for the cross-subsidy embodied in postage-stamp

ously supplying all three is obviously most economical, and would require that each user pay \$3.67 million on average (one may pay less than \$3.67 million, but then the other two will have to pay more than \$7.33 million, or (\$11 million - \$3.67 million), violating the stand-alone cost test. In any case, one pair of buyers will end up being assessed for more than their \$7 million stand-alone costs. So there will be a pair of buyers who will prefer to build their own system rather than paying their share of the total system costs. The core is empty, as it requires that the stand-alone test hold for all coalitions as well as individuals.

5/ One condition for the core to be non-empty is that the cost function be concave in the sense defined in Young (1994) and discussed below.

rates has much less force and now might discourage development elsewhere or encourage excess development in remote areas.

Moreover, it is not clear how strong are the merits of a system that creates incentives to develop facilities in regions that would otherwise not be economically viable. Basic economic principles imply that any subsidies embodied in the postage-stamp regime are not justifiable on grounds of economic efficiency. Even if remote regions were economically viable, the effect of a uniform postage-stamp system is effectively to tax production in low-cost, not-so-remote areas and subsidize production in remote, high-cost areas. Both the tax and the subsidy create deadweight losses. This is because regions with above-average costs produce at higher-than-optimal (*i.e.*, efficient) rates, while those with below-average costs produce at lower-than-optimal rates.⁶

There are a number of other fairness criteria that policy makers might wish to apply in allocating costs. Below we describe several which have been analyzed and discuss some of their implications. One fairness criterion that most would agree is desirable is that equals bear equal cost shares. So if two firms affect system costs in the same way, they should be allocated the same costs. In addition, this *symmetry* condition requires that the cost allocation be invariant to the labelling of the firms and to the order in which users are added to the system. One significant objection to imposing a symmetry requirement is that, in some cases, an asymmetric cost allocation will induce some to stay on the system and contribute to total costs in excess of stand-alone costs, whereas a symmetric scheme will lead to bypass. Thus, the symmetry condition can conflict with the stand-alone cost test.

Three other fairness and reasonability properties that cost-allocation rules should satisfy are a decomposition principle, a monotonicity principle, and consistency. The *decomposition principle* requires that each user bears an equal share of the costs of the compo-

nents it uses. It also implies that no one should have to contribute to portions of the system that they do not use at all. In other words, only those who use some components should have to pay for them. *Monotonicity* implies that as total costs increase, allocated costs should also increase, or at least not decrease. *Consistency* in cost allocation says that the principles used in determining cost shares for the entire set of users should apply equally to subsets of users.

In combination, the decomposition principle and symmetry have strong implications for cost allocation. They essentially nail down a unique allocation in which everyone benefiting from a component pays essentially the incremental costs of satisfying their demands.

The fairness criteria we have listed above satisfy the condition that they continue to apply as the environment changes. However, they do sometimes conflict, and different sets of criteria imply different cost-allocation rules. In what follows we try to outline what, in our view, are some of the more important criteria, and explain potential conflicts and their ramifications.

In particular, we consider two cost-allocation rules that have been well analyzed in the economics literature: the nucleolus and the Shapley value. These are two alternative views of what constitutes an ideal cost-sharing rule. Subsequently (Section V), we explain how these two ideals can be applied to determine pipeline rates.⁷

IV.1 The Nucleolus: Consistent, Symmetric, and Homogeneous Cost Allocations

The *nucleolus* is derived from a set of axioms. In particular, the nucleolus is the *unique* cost allocation that is: (a) symmetric, in that it treats equals equally and does not change when agents are re-labelled, or when the order in

6/ A technical appendix, which illustrates the economic losses associated with uniform tolls, is available from the authors on request.

7/ Under certain circumstances, setting rates using "Ramsey prices", in which rates are inversely proportional to the elasticity of demand for pipeline use, will be efficient. However, as noted by Young (1994) and by Lewis (1949), Ramsey prices are inherently inequitable since they penalize those with least resort to alternatives. This aspect also makes Ramsey prices politically unpalatable. They are not discussed in this paper.

which they are added to the system is changed; (b) passes through costs directly incurred by shippers; (c) is homogeneous, in that if all costs go up or down by some proportion, α , all users' cost allocations go up or down by the same proportion α ; and (d) is consistent for sub-groups of the entire set of pipeline users.

The nucleolus can be calculated by splitting the costs equally among the users of a common facility, or a portion thereof. It is essentially the cost allocation that is the mid-point of the core. The nucleolus also has the property that it maximizes the cost savings of the group of users that has the smallest cost savings among all possible groupings of facility users.

The notion here is that various individual or groups of users of a system may enjoy various degrees of savings in using it. For example, a large-scale user may obtain fewer economies of scale or scope in relation to the relevant stand-alone costs compared with those obtained by a small-scale user (at least on a per-unit basis). The nucleolus maximizes the savings enjoyed by those enjoying the least advantage from being in the system compared with the best alternative available for that grouping. The main problem with the nucleolus is that it is not *monotonic*. What this means is that the cost share of a user could fall even though he were using a component of the system whose costs have increased.⁸

IV.2 The Shapley Value: Symmetric, Additive, and Monotonic Cost Allocations with No Cross-Subsidies

The *Shapley value* yields another cost-allocation

⁸/ This problem could be overcome by the *per capita* (or *per user*) *nucleolus*, which is also the maximum of the series of cost savings for all possible groupings of users and which will be monotonic. However, it is not consistent. Consistency is an important criterion when, for example, in the case of a pipeline the set of receipt and delivery points is changing over time. Consistency requires that the cost allocation for any coalition not change when the cost-allocation problem is confined to one involving only those in the given group. Note, too, that neither the nucleolus nor the per capita nucleolus will be easy to measure. What would be desirable is a cost-allocation rule that is relatively easily computed and satisfies the principles of consistency, homogeneity, and symmetry.

rule that satisfies many desirable properties. Like the nucleolus, the Shapley value can be derived from a set of axioms. These axioms differ slightly from those that identify the nucleolus. The Shapley value has the additional property that it is also a fairly natural extension of a simple rule that applies in special circumstances. This simpler concept, which is the *decomposition principle* mentioned earlier, roughly speaking says that a firm that uses several pieces of the system should pay a suitable share of those pieces it uses.

To apply the decomposition principle, the stand-alone cost of serving any group of users must be decomposable into the costs of the components, or the cost elements, used by that group. If the cost function were decomposable, then the decomposition principle would merely allocate the costs of each component suitably among each component's users. So, for example, if there were three firms using a given pipeline link, the decomposition principle would allocate the costs of that link in proportion to the decomposed costs among those three firms. This allocation of costs should be based on both usage and each firm's fraction of the reserve capacity for that link. In other words, the allocation of costs should be based on those factors that contribute to costs.

The decomposition principle does yield outcomes that are in the core, that is, outcomes that satisfy stand-alone and incremental cost tests. But the principle can only be applied when costs can be decomposed into elements that are additive. However, the same type of idea can be extended to cases in which the cost function cannot be decomposed so readily. The Shapley value is the resulting cost allocation. The exact expression for the Shapley value is somewhat complicated, but it essentially states that each firm will contribute an equal proportion of the total costs allocated to each possible group it could join.

More precisely, consider the incremental costs of serving a given user when that user is added to a group of users. Now, suppose that system costs are calculated incrementally when adding users to a group one at a time in a random order. The Shapley value for a given user is just the average of the incremental costs

for that user among all possible ways in which the incremental costs can be calculated for him. Thus, the Shapley value is a cost allocation for each user that is based on a measure of each user's average incremental costs.

As discussed in Young (1994), the Shapley value has a number of desirable properties:

1. It is the unique cost-allocation rule that is: (a) symmetric; (b) additive;⁹ and (c) charges nothing to firms who do not contribute to costs.
2. It is also the unique cost allocation that is symmetric and strongly monotonic, that is, it allocates *all* users larger cost shares whenever the total costs of serving everyone increases.
3. The core of every case in which the cost function satisfies a concavity condition (that can be explained in terms of the number of nodes and the length of the links) is non-empty and contains the Shapley value.

The conditions that the cost-allocation rule is additive and charges nothing to users who do not impose costs on the system constitute, what some would view, an important fairness condition. Suppose, for example, that a system component (sub-system) were built exclusively to serve one small group of users. These fairness conditions essentially imply that no one outside that group would have to bear any of the costs of that sub-system. Symmetry, as we have discussed above, requires equal treatment of firms that contribute equally to costs and have equal usage. The concavity condition, can be also expressed as a *submodularity condition*, and essentially means that the costs of serving two groups plus the stand-alone costs of serving those who are in both is less than the sum of the stand-alone costs of each of the groups. When costs are concave, the incremental costs of adding users at new receipt points or delivery points will be decreasing.

The Shapley value has one significant drawback in that it need not be in the "core." In other words, the Shapley value need not satisfy the stand-alone cost test that we discussed

⁹/ A cost allocation is additive when, if two users or groups of users are combined, the cost allocation for these users is the sum of the individual user cost allocations.

above. However, the Shapley value has two advantages: (1) it always exists; (2) it identifies a unique cost allocation.

V. Implementation Issues

V.1 Implementation of the Shapley Value

The nucleolus and the Shapley value provide benchmarks for devising a toll system which best approximates, as much as is practical, fairness and reasonableness standards. Neither can be directly applied with ease. To use either of them requires that some possibly costly administrative procedures be set up to impute incremental costs for each shipper, receipt meter station, and delivery sales station.

To appreciate how a multi-zone system, in which tolls are based on the zones corresponding to pipeline receipt and delivery points, can approximate the theoretical benchmark of the Shapley value, it is useful first to describe the steps that would be needed to implement it. We focus our discussion on implementing or approximating the Shapley value, although most of it applies to the nucleolus as well.

In cases where costs can be decomposed, it is relatively straightforward to compute the Shapley value. It is possible that costs can be decomposed in an appropriate manner for many pipeline projects. In such cases, the use of the Shapley value or the decomposition principle would eliminate the need for debate about whether rates should be based on the average distance to the delivery or the receiving point. The Shapley value would impose costs on those firms that use the relevant components of the network. Debate might still occur as to how to measure incremental costs. However, the Shapley value would essentially evenly divide costs of components shared by multiple gas producers or shippers. And all incremental capacity costs would be directly allocated to those shippers on the basis of for whom that capacity was constructed.

V.2 Pitfalls in Implementing the Shapley Value

The Shapley value is a theoretical ideal. For large and complex pipeline systems it is likely

to be difficult to apply directly. Here, we describe some of the difficulties in applying the Shapley value, and provide some suggestions for surmounting them. One problem, which we have noted above, is that the core need not exist, and even if it does, the Shapley value need not be in it. In such instances, there would be groups of pipeline users who would wish to break away to avoid the cost allocations imputed to them by the Shapley value.

Circumstances change, and so at some points in time the Shapley value will be in the core and at others it will fall outside of it. Changes are always occurring in the potential demands placed on the system. It would be helpful if the cost-allocation rule were to remain viable for any likely scenario. The Shapley value, which always exists, can always be imposed; however, coalitions would, at times, have incentives to break away. How likely, or how often, this would occur is an empirical question.

Another problem in implementing the Shapley value, even in a simple case in which costs are decomposable, is in measuring the incremental costs attributable to each user. The appropriate way to decompose and attribute costs is likely to be in terms of the capacity planned for each user. To appreciate the ease, or difficulty, that would be encountered in attempting to impute costs based on the Shapley value, it is useful to examine the types of cost attribution required. This will, in part, help guide how best to implement a cost-sharing rule in practice.

Incremental costs are based on the planned capacity requirements that drive them. Direct costs — the costs of actually moving the commodity in the transmission system — based on the fairness principles embodied in the Shapley value would be directly passed through. Indeed the decomposition principle, in conjunction with the principle that shippers imposing no costs on the system (or separable portion thereof) are not allocated any costs, requires that direct costs be passed through. Thus, implementing the Shapley value would require separation of direct costs, which are passed directly through, from common or joint costs. And then the common capacity costs, or

the costs of building the capacity to meet projected demand, would be allocated among users.

Of course not all capacity need be used, and not all anticipated demand need be realized. Conversely, in other instances, more demand may be placed on the system than anticipated. It is also the case that there could be considerable differences in the percentage of anticipated throughput that would actually occur. This means that some users would effectively have reserved more capacity than they needed and others less. Unused capacity could then be traded on a capacity-release market among the potential users. Such a market would alleviate potential shortages and enhance the efficiency and utilization of the system.

On the administrative side, there is a question of how to go about measuring component costs and capacity costs. In particular, use of the Shapley value requires that incremental capacity be imputed for each user. It could be difficult to obtain such measures based solely on accounting data. Accounting data are not intended to report the calculations made in capacity planning. Such calculations would be needed to reconstruct fully the capacity costs for each cost element. In particular, it can be difficult to reconstruct precisely how capacity planning and investment decisions were based on projected demands and to decompose those plans on an incremental basis. The best that might be possible is to allocate costs proportionally to what were the initial projections of demand or requests for services. Of course, those data might not be available, and then actual usage averaged out over some appropriately lengthy period would probably be the most appropriate procedure.

V.3 Practical Solutions

The discussion above indicates that it would likely be difficult to allocate costs based on the Shapley value. This is not to suggest it would be impossible, but rather more that it would be less practical the greater the complexity of the network. The Shapley value is difficult to explain. However, for simpler systems, or com-

plex ones that have been aggregated in zones, the Shapley value essentially reduces to the cost allocation determined by the decomposition principle, which is reasonably straightforward. So, to a large extent the practicality of the Shapley value will depend on the degree of cost decomposability and on the extent to which costs can be aggregated.

There would likely be difficulties in applying the Shapley value or decomposition principle to a complex system arising from a large number of pipeline delivery and receipt nodes. One example is the NGTL system in Alberta. Here there are a plethora of receipt and delivery points with a great deal of common costs. In such a case, it would be more practical to aggregate various sets of users which are similar in some dimension, such as geographic proximity, and then to employ a weighted version of the Shapley value to the groups in partitioning the entire set of users. Such weighted versions of the Shapley value would continue to satisfy many equity conditions as well. This also leads to the notion of zonal charges.

Breaking up the system into zones, so that all users who share more or less equally in the use of the system contribute equally to its costs, can approximate tolls that would be determined by the Shapley value. In other words, the average distances for shippers who use many of the same facilities can be used to determine the tolls. So, if two shippers require use of transport facilities through some region, the average distance of transport, as well as the cost per mile or kilometre of the facilities, can be used to set tolls, or charges, for shipments in, through, or out of that zone.

VI. Other Equity and Efficiency Concerns

The cost-allocation rule used in setting tolls for a pipeline system has an effect on user incentives to participate, to bypass, and to invest in the development of new and existing fields. Here we discuss how these factors can affect the optimal design of a cost-allocation scheme.

VI.1 Incentive Compatibility for Shippers and the Pipeline

One specific problem that is typically a concern of regulatory agencies when allocating costs is the fact that the costs allocated to a firm might not exceed that firm's stand-alone costs (and therefore satisfy the stand-alone and incremental cost fairness criteria) and yet exceed the firm's willingness to pay. For example, a gas producer might prefer to shut down some wells rather than pay its allocated share of the costs of serving those sites, and yet be willing to pay the incremental costs for the pipeline to provide service to those wells. The regulatory authority will not generally know the gas producers' minimum or "choke" prices, nor will the authority know the incremental costs of providing service to each receipt point.

The notions of fairness and efficiency embodied in the incremental cost test imply that the cost allocation should not establish tolls in such a way that a user ever faces a cost allocation exceeding his willingness to pay when that willingness to pay is more than his incremental costs. Additionally, the system operator should face incentives to provide service to every gas producer whose willingness to pay exceeds the incremental costs of service. However, these fairness and efficiency goals typically cannot be fully realized in practice. A regulator will not know each shipper's "choke" price, that is, there is incomplete information. In addition, the pipeline company will not know that price either. Similarly, neither the regulator nor the shippers will know the pipeline company's costs.

The optimal tariff scheme will maximize fairness and efficiency goals subject to incentive compatibility and constraints, that is the pipeline and the shippers will respond to the tariff rule so as to maximize their own objectives (such as long-run profits) given their costs (which are known only to them). However, a toll system which meets these incentive-compatibility conditions imposes tolls, which in some cases will deter the pipeline from providing service to wells where willingness to pay exceeds incremental costs. This follows because in practice the regulatory au-

thority will not know the willingness to pay, choke prices, or incremental costs, and neither the pipeline nor the gas producers have much of an incentive to report these values accurately. So, rather than over-pay the pipeline, the regulator might, for instance, wish to allow some situations to arise in which some gas that should be shipped is not.¹⁰

There is a sizable literature on these types of incentive problems. Price-cap and incentive-regulation schemes are designed, in part, to provide a utility operator with the appropriate incentives under regulation to provide service to every user for which incremental costs are less than willingness to pay. Such pricing flexibility will typically lead to distance-sensitive tolls.

Additionally, some incentive schemes that can be implemented present each user with a menu of choices. The choices would effectively allow each user to reveal its valuations.¹¹ These mechanisms need not satisfy many of the fairness criteria discussed above. However, it is also the case that the participation constraint, that tolls not exceed choke prices because the cost allocation assigns them too large a cost share, might not be a pressing issue when transport costs are a relatively small share of the total costs of marketing gas.

VI.2 Static and Dynamic Efficiency

Another concern is that any cost-allocation result be as efficient an outcome as possible. Trade-offs between efficiency and equity can arise. In choosing between policy measures it is useful to keep both in mind, and certainly options that adversely affect both efficiency and equity should be avoided.

In terms of static efficiency, one of the main concerns is that gas be delivered to users at

minimum total costs — including extraction and shipping costs. A postage-stamp scheme, or any other cost-allocation scheme that is not distance-sensitive, will effectively cross-subsidize output from remote and more costly sites (as already noted). Moreover, non-distance-sensitive cost-allocation mechanisms, and any other cost-allocation mechanism that provides cross subsidies, have effects on investment incentives that could accentuate welfare losses from cross subsidies over time.

For example, cross-subsidization of remote sites at the expense of nearby ones could lead to increasingly larger output at the remote sites than would have been the case without subsidies. Without a cross subsidy, a firm might invest in new facilities in a location closer to delivery points than would be the case if a cross subsidy existed. This means that the social cost of extraction and delivery will be higher than optimal in the long run. The short-run effects of the postage-stamp scheme will simply tend to alter extraction rates between facilities, but could also cause some locations that should remain open to shut down.

VI.3 Complexity

Strict adherence to many of the principles discussed above can impose significant costs on both the regulatory agency and the pipeline. However, the principles do have practical value. For instance, as noted above price-cap schemes are, in some sense, simplified incentive-compatible mechanisms. Similarly, a simple system of zonal changes can be used to approximate the Shapley value or the nucleolus. The practical problem facing a regulatory agency is to balance theoretical performance with administrative and compliance costs arising from possible complexity.

VII. Summary

We have examined various ways in which a fair and reasonable pipeline cost-allocation scheme can be implemented. Uniform charges such as those reprinted by postage-stamp cost allocations will not usually satisfy most concepts of fairness and reasonableness. In addi-

10/ Other concepts of fairness, such as those embodied in the Shapley value, or a desire to subsidize development in some regions over others, can also conflict with the application of the incremental cost test.

11/ It has been shown that mechanisms can be constructed that are: (a) efficient; (b) incentive-compatible; (c) individually rational (*i.e.*, pass a stand-alone cost test); and (d) allocate costs exactly. See Young (1994).

	Fairness and Equity Criteria					
	Symmetry	SAC/IC	Decomposable	Determinable	Monotonic	Consistency
Postage Stamp	Y	N	N	Y	Y	N
Nucleolus	Y	Y	N	Y	N	Y
Shapley Value	Y	Y	Y	Y	Y	Y

Notes: SAC/IC = Stand-Alone/Incremental Cost Tests; Y = Yes; N = No.

tion, postage-stamp cost allocations can result in inefficient production patterns and in inefficient bypass.

We have argued that cost allocations should pass a stand-alone cost test and an incremental cost test. In other words, no one should pay costs in excess of their stand-alone costs and no one should pay less than their incremental costs. These criteria can rule out many obvious cost allocations, such as postage-stamp rates, but do not identify a unique outcome.

We then proposed that a cost allocation in which each party pays its proportional share of the parts of the network it uses would meet most of the criteria for fairness and reasonableness considered. In particular, such a cost scheme would be symmetric, in that it would treat equals the same, pass through direct costs, be consistent when the set of users and load patterns change, and be monotonic, in that no one's cost share could fall if total costs were to increase.

We explained how such a rule might be implemented. The main difficulty is in determining the appropriate manner in which to decompose the cost elements when there is no

direct contract between the parties at the receiving node and the delivery node. In such cases, those involved in determining the tolls at the two ends would need to identify which system components are being used to provide service to the parties involved. We argued that a system of zonal charges can approximate the ideal cost allocation, and can involve much lower administrative costs. Our results are summarized in the table (above).

One of the cost-allocation methodologies, the postage-stamp scheme, fails to respond to more than half the six fairness and equity criteria considered. The nucleolus responds favourably to four of these six criteria, while the Shapley value will in most cases address all six.

References

- Young, H. Peyton (1994) 'Cost Allocation,' Chap. 34 in Robert Aumann and Segiu Hart (eds.), *The Handbook of Game Theory, Volume II* (Amsterdam: North-Holland).
- Lewis, W. Arthur (1949) *Overhead Costs; Some Essays in Economic Analysis* (London: G. Allen & Unwin).