American Economic Review Papers and Proceedings, May 1976.

Replaces Caletch
SSWP No. 106

An Experimental Market for Public Goods: The *PBS* Station Program Cooperative

By JOHN A. FEREJOHN AND ROGER G. NOLL*

Beginning with the 1974–75 season, the Public Broadcasting Service (*PBS*) undertook a three-year experiment to decentralize through a market process the selection of programs to be broadcast over the national noncommercial television network. The experimental market, called the Station Program Cooperative (*SPC*), enables each of the approximately 150 participating stations to "vote with dollars" for the programs they prefer, based on information about the subject matter, quality and costs of the alternatives.

The *SPC* had two purposes. The first was to determine whether a decentralized market mechanism for acquiring national programming can be established that, first, has a reasonable cost and, second, retains the basic features of a network. A network is a means of centralizing the acquisition, transmission and promotion of programming. Although the advantages of networking are several, probably the most important is that it saves substantial transactions cost in comparison to the two other alternatives: syndication (in which each station negotiates for nationally distributed programs with the owner of the rights to the program) and localism (in which each station produces its own programming). (See Bruce Owen, Jack Beebe and Willard Manning.)

The second purpose of the *SPC* is to determine its impact upon the types of programs that are broadcast. Centralized programming authorities such as the present commercial networks and the Corporation for Public Broadcasting, because they are few in number, are relatively easy targets for political pressure with respect to program content, may at least tacitly collude in programming decisions, and in any event, have limited experience concerning the nature of the demand for programs in all markets in the network system. (See Noll, Merton Peck and John McGowan, Ch. 8.) Because of the cost advantages of networking and the technical barriers to entry of new networks, in principle a network can produce an inefficient programming mix and still meet minimum standards of profitability or, in the case of the noncommercial sector, public acceptability. Concomitantly, station managers may be too unsophisticated about judging programming quality and sensing the structure of demand in the local market to select programming that, in benefit-cost terms, is "better" than that selected by a centralized decision maker. While direct evidence on which system works better is not possible to acquire, two useful questions can be answered, at least in principle: (1) Do the *SPC* and centralized decision making lead to the selection of different programs, and if so, is that difference systemmatic? and (2) Which system produces higher audience

* California Institute of Technology. Financial support for preparing this paper was provided by the National Science Foundation RANN program, grant #APR76-01920. We are indebted to Richard Beatty and *PBS* for providing us with remarkably complete data for all three *SPCs*. We are, of course, solely responsible for the contents of this paper.

ratings and/or greater total revenues (private contributions plus government appropriations) for the noncommercial television network? This paper does not address these issues. (See Natan Katzman and Noll.) That research is still underway.

Though unintentional, the *SPC* experiment has a third, more general use. The *SPC* is a test of a novel mechanism for public decision making about a particular class of public goods, specifically those for which binary (all or nothing) exclusion is possible.

The purpose of this paper is to make the connection between the *SPC* and the theory of public goods, and to draw some preliminary conclusions about the workability of this kind of artificial market. It is, therefore, addressed to the first and third issues.

A television broadcast is a public good in that the cost of transmitting the program within a given geographical area is independent of the number of television sets that are tuned to it. So, too, is the program a public good as far as each station is concerned. Excluding rents arising from differing popularities of programs, the costs of a program are independent of the number of stations over which it is broadcast. Even national network distribution costs are, in most instances, independent of the number of stations receiving the program. Nearly all programs distributed by networks are carried over the national telecommunications system. If the program is distributed from coast to coast, the transmission costs are only slightly dependent on the number of intermediate points at which the program is received and broadcast. As the distribution of television programs moves to satellite transmission, national interconnection costs will become totally independent of the number of stations in the network.

Nevertheless, exclusion of stations from use of a particular program is also essentially costless. In both the commercial and noncommercial systems, few stations accept ("clear") all network programs. One way to view the process by which stations make programming decisions is that each station has a preference ordering over all available programming possibilities and selects the most preferred feasible set of programs, given its budget and the costs of the alternatives. Since programs are public goods, these decisions have external effects; namely, if a program is selected by one more station, costs per station decline, and hence from the vantage point of other stations some programming alternatives involving this program that dominate a previously selected alternative may become feasible.

I. The Mechanics of the *SPC*

The heart of the *SPC* is a mechanism for calculating the price of each program for each station that broadcasts it. These prices are determined through a sequential process in which a series of trial price vectors are proposed, responded to, checked for acceptability with program producers, and, if they generate too little or too much revenue, updated.

The first stage of the *SPC* is the dissemination of program proposals. Program producers—primarily stations, but also some independents such as the Childrens Television Workshop ("Sesame Street")—provide descriptions of proposed programs (series or special events), prices, and in many cases pilots. The form of the final proposal is usually negotiated with *PBS*, which among other things, tries to keep program prices roughly in line with production costs. In each of the first three years of the *SPC*, about 250 programs were proposed.

The second stage of the *SPC* culls the list of proposed programs in the following manner. On the basis of the final pro-

posals, stations are asked to assign "priorities" to each program, which amounts to rating the programs on a scale of one to five (five is best). The sum of the ratings for each program is calculated, and the programs are rank-ordered according to their scores. Programs ranking in the top third, plus a few others which *PBS* believes deserve special consideration (their classifications show high variance, or additional information about the program will become available and might change scores), are then allowed to enter the market phase of the selection process. This process served to cull the list of contenders to 93 in the first year (*SPC* I), and 136 in the second (*SPC* II).

Next, the market phase of the *SPC* is begun. Its main feature is an algorithm for calculating the program prices facing each station. The formula for calculating the price P_{ij} of program j to station i is:

$$P_{ij} = C_j \left(.8 \frac{B_i}{\sum_{k \in Sj} B_k} + .2 \frac{N_i}{\sum_{k \in Sj} N_k} \right)$$

where

C_j = the cost of program j
B_k = the programming budget of station k
N_k = the population of the area served by station k
S_j = the set of indexes of stations that have agreed to purchase program j.

Obviously, prices and the identity of purchasing stations are interdependent. In the *SPC*, the initial price calculations are based upon the assumption that each program is selected by stations representing 80 percent of the combined budget and total audience of the entire public television system. In subsequent rounds, prices are calculated on the basis of the budgets and populations of the stations selecting the program in the previous round. Once the market has eliminated all but a few more programs than are likely to be purchased, the *SPC* manager announces that in subsequent rounds stations will be required to continue selecting a program that they selected in the previous round if the price does not increase. The market process continues until all programs either generate revenues acceptable to the producer of the program, receive no votes from other than the producing stations at prices based upon selections made in the previous round, or are withdrawn from consideration by the producer.[1] Then a final "late purchase" round takes place in which each station can select additional programs from the list of purchased programs. Stations eventually purchased 25 programs in *SPC* I and 38 in *SPC* II, each station buying about 20 in the first and 30 in the second.

The costs of programs are only upper bounds on the amount of revenue producers eventually receive from the *SPC*. At any phase of the market process, a producer may guarantee that the current price vector for his program will not rise —that is, that the producer will absorb the shortfall in revenues in relation to the listed total cost. Since many of the programs entered in the *SPC* will be produced in any event for local or regional audiences, any revenue the producer receives from the national market is a net gain. Similarly, coalitions of buyers are permitted to guarantee the stability of current prices by agreeing to make up the difference between revenues and costs.

II. The Economics of the *SPC*

The appropriate theoretical framework for analyzing the properties of the *SPC* is the extension of general equilibrium

[1] The purchase rounds are conducted through a nationwide computer network. Each station manager sits at his teletypewriter, receives messages about the progress of the *SPC* (including price updates), and each round enters his selections.

theory that includes pricing of public goods. Models of optimal choice of political jurisdiction, the theory of clubs without congestion, and Lindahl pricing of public goods are obvious examples. (See James Buchanan, Bryan Ellickson, Duncan Foley, Eric Lindahl, R. W. Rosenthal and Paul A. Samuelson). In a market with p programs, a particular pattern of program selection and rejection can be represented by a p-dimensional vector in which the elements are either zero or one according to whether a program was rejected or selected by the station. A family of 2^p such vectors describes all possible decisions by a station, and presumably a station is able to generate a transitive weak ordering over these vectors that reflects its programming preferences. The problem faced by a station is to pick the highest-ranking selection vector—e.g., a vertex on the n-dimensional unit square—that is feasible, given its budget and the prices of the programs.

Each program producer seeks to maximize the revenue it derives from a program, subject to two constraints: the revenue received through the *SPC* must cover the incremental cost of allowing the program to become national, and it must not exceed the maximum allowable costs negotiated with *PBS,* which are intended to be approximately total production costs. The minimum and maximum prices are identical for programs that will be produced only for national distribution unless the producer can find other sources of financial support, is willing to take a loss, or is a station that values the program idea sufficiently highly that it is willing to pay more than its share of the costs to guarantee its production.

The *SPC* is a process of generating for each station a price vector and a vector of selections and for programs a cost vector that constitute a stable equilibrium in the feasible production and budget sets. This algorithm has two features: an initialization procedure that generates the first trial set of prices and an updating procedure that changes prices until a stable equilibrium is found.

From a theoretical standpoint, it is a miracle that the process works. One source of difficulty is that the production set is not convex. The space in which programs are located is of an unknown number of dimensions of program characteristics. Past research has been notably unsuccessful in defining these dimensions in an empirically useful way (see Edward Greenberg and Harold Barnett), but as a working hypothesis a program can be thought of as a particular combination of production talents, production quality (e.g., the average number of "takes" per printed scene, the number of scenes filmed on location, etc.), and subject matter. Each program proposal in the *SPC* is, to all practical purposes, a single point in this space—not a ray, but a point—because essentially no variance is possible during the operation of the *SPC* in any of the relevant dimensions, including the number of units of the program to be sold. As a result, linear combinations of programs cannot be formed to arrive at locations in program-characteristics space that are between the program points. Consequently, even if at all factor prices one unit of a particular program would never be selected in market competition with one unit of each of the other programs, it may still be selected in the *SPC* since the combinations of other programs that strictly Pareto dominate it cannot generally be formed.

The first two *SPCs* had another potential problem arising from the initialization procedure. Because the initial price vector was calculated from grossly unrealistic assumptions about the number of stations that would select each program, the station selection vectors they called forth were

necessarily outside the system's feasible set. In *SPC* I second round prices averaged nearly four times first-round prices! This kind of initialization causes a massive reduction in the number of program selections per station in the first few rounds. If initial selection vectors are sufficiently diverse that few programs appear clearly out of the running in the initial stages, stations will not quickly concentrate votes on a relatively few popular offerings, and so for nearly all programs prices will continue to rise through the first few rounds. If preferences are sufficiently diverse and the number of programs large enough compared to the number of stations, the process might not converge to the feasible set. One can imagine a desperate round in which each station, facing astronomical prices, picks exactly one program, all stations pick a different program and no program can cover costs on the budget of a single station.

Theoretically, whether initialization and updating procedures for the *SPC* can be derived that reach a stable equilibrium is as yet unanswered. Our conjecture is that they can, particularly if the procedures constitute an improvement algorithm rather than a search for the feasible set from outside it.

The third *SPC* will try a new initialization procedure. Instead of calling out a selection vector and observing a budget-breaking consequence, each station will provide the *SPC* with a total expenditure limit and a preference ordering of the programs. The *SPC* will then calculate new prices based on these preferences and tell each station how far down its preference ordering it can go, given its budget constraint. For the first round, preference ordering will not be based on an explicit price formula as stations will draw their own inferences from program cost information. In subsequent rounds, the same price updating procedure used in the other *SPCs* will be used. The basic idea is first to find a point on the boundary of the feasible set, and then move along the boundary until an equilibrium is reached. The main difficulty with the procedure is finding a nontrivial initial feasible solution.

III. The Results of the First Two *SPCs*

The *SPC* results of interest here relate to the convergence properties of the system; we will not investigate the issues related to the types of programs the system selected.

In both of the first two *SPC* experiments, the market phase of the selection process was allowed to run for twelve rounds. *SPC* I had two-thirds as many programs from which to choose, and eventually purchased two-thirds as many programs, as *SPC* II.

By round 6 in *SPC* I and round 8 in *SPC* II, the programs that were purchased had all but been determined. In *SPC* I, the three programs still in the market during round 6 that were subsequently eliminated received support from 3, 5 and 8 stations, respectively, while the three lowest numbers of supporting stations for programs that were eventually purchased were 19, 21 and 34. Nineteen of the 25 programs that were eventually purchased were, by round 6, receiving support from over half of the stations (75).

In round 8 of *SPC* II, the six programs that were eventually dropped received support from 2, 9, 36, 40, 64 and 66 stations, respectively. Of the 38 programs that were eventually purchased, 31 were supported by more than half the stations in round 8. Five more received support from between 58 and 73 stations and two that were eventually purchased received only 40 and 42 votes in round 8. One of the latter was clearly a special case, consisting of 100 sixty-second messages promoting coming attractions. Thus, of

the ten conventional programs receiving support from between 36 and 73 stations, six were purchased. All six that were purchased had total costs below $15,000 per program-hour, which is inexpensive. Of the four that were rejected, hourly cost data are available only for two. One was a children's program that cost $57,000 per hour—more than triple the cost of "Sesame Street" or "The Electric Company"—and the other was a documentary series that cost more than $100,000 per hour. Thus, further classification of programs by cost apparently cleanly separates the purchased from the excluded programs.

In both markets, the first few rounds saw massive reductions in the number of programs selected per station and generally rising program prices. In *SPC* I, prices for programs remaining in the market rose abruptly at the start, and did not, on average, begin to fall until after the seventh round, while in *SPC* II, prices began to fall after round 3.[2] In both markets, the average price of purchased programs did not return to the average round 1 initialization price until the last round.

The first two *SPCs* appear to have reached an equilibrium, at least within the context of the operating rules. Although the twelve purchase rounds occurred over only a few days, late purchases were possible for several weeks. Two programs in *SPC* I and seven in *SPC* II increased the number of stations supporting them by 3 to 8 percent during the late purchases. In both markets about 55 percent of the programs made no additional sales during the late purchase, while 28 percent in *SPC* I and 26 percent in *SPC* II increased their number of supporters by less than 2 percent.

Outside the context of the *SPC* rules, the final list of programs purchased may not be an equilibrium. The key operating rule in the *SPC* appears to be the nonreincarnation rule: once eliminated, a program cannot reenter the market at a later stage. This eliminates the possibility of cyclic selections, but may also prevent the system from reaching an efficient equilibrium if one exists. About one-third of the programs that eventually were purchased received twice or more as many votes in the final selection round as they did at their ebb during the market. Many of these programs also had mediocre premarket rating scores, averaging below a ranking of three in the five point scale.

With the exception of buyer and seller guarantees as explained above, stations did not have an opportunity to price-discriminate against themselves in order to register an intense preference for a particular program. A group of stations with an intense preference for a particular program, A, but unable to guarantee its purchase, could not offer remaining stations a discounted price to tempt them away from another program, B. If the remaining stations had only a slight preference for B over A, the pricing rule could prevent a substitution of A for B at prices that would make the move a Pareto improvement.

IV. Conclusions

Because our research on the *SPC* has only begun, few conclusions can be drawn as yet. Most of the inferences we are prepared to draw relate to the additional research that is required to gain a better understanding of how the system operates, if it could be adapted to the acquisition of other public goods (national parks, commercial broadcasting, etc.), and whether it produces acceptable programming results at reasonable transactions cost.

[2] This statement is true for both the unweighted average price and the average of prices weighted by the number of stations selecting the associated program. Movements of the latter index the change in the average budget of a station, holding its selections at the mean of all stations.

An obvious candidate for further work is theoretical examination of the equilibrium properties of the *SPC*. Two types of questions are relevant: What normative properties, if any, do the existing *SPC* rules have, and would other initialization and updating procedures work better? Although of some interest is whether the system reaches an equilibrium that is Pareto optimal, for policy purposes the conclusion that it does not is not persuasive that the *SPC* should be abandoned. A centralized system, for reasons outlined briefly above, also is likely to be inefficient. Of prime practical interest are the consequences of alternative procedures for operating the system.

Laboratory experiments could provide a useful additional source of information about alternative structures of *SPC*. Very briefly, the idea of each experiment would be to ask a group of subjects to select an unstated number of items from a list of alternatives. Total costs for each alternative would be provided, as well as a payoff function for each group member. An initialization and updating procedure would be adopted, and the general mechanics of the *SPC* followed. By running the identical experiment several times, the uniqueness of the results of the system could be checked. By varying payoff functions while holding procedures fixed, the sensitivity of the system to intense preferences could be explored. And by varying procedures while holding payoffs constant, the performance of various initialization and updating methods —particularly of improvement algorithms and of other procedures—could be compared. The advantage of the experimental system, of course, is that it is far less expensive—in terms of operating costs, risks to something of significant value (the public television system), and the time between observations—than, through the years, manipulating the *SPC*.

Finally, the overriding conclusion of our investigations to date is that experimental markets for public goods are an interesting, potentially useful and almost completely unexplored focus for research. If this paper conveys to other social scientists our perception of the fascinating possibilities suggested by the *SPC*, its most important purpose will have been served.

REFERENCES

J. M. Buchanan, "An Economic Theory of Clubs," *Economica*, Feb. 1965, *32*, 1–14.

B. Ellickson, "A Generalization of the Pure Theory of Public Goods," *Amer. Econ. Rev.*, June 1973, *62*, 417–32.

D. Foley, "Lindahl's Solution and the Core of an Economy with Public Goods," *Econometrica*, Jan. 1970, *38*, 66–72.

E. Greenberg and H. J. Barnett, "TV Program Diversity—New Evidence and Old Theories," *Amer. Econ. Rev. Proc.*, May 1971, *61*, 89–93.

N. Katzman, *Public Broadcasting's Station Program Cooperative*, Aspen Program on Communications and Society, Palo Alto 1974.

E. Lindahl, "Just Taxation—A Positive Solution," in *Classics in the Theory of Public Finance* (R. A. Musgrave and A. T. Peacock, eds.), St. Martin 1964.

R. G. Noll, "Decentralization of Public Television," in *Conference on Communications on Communications Policy Research: Papers and Proceedings*, Office of Telecommunications Policy, Washington 1973.

———, M. J. Peck and J. J. McGowan, *Economic Aspects of Television Regulation*, Brookings Institution, Washington 1973.

B. M. Owen, J. H. Beebe and W. G. Manning, Jr., *Television Economics*, Lexington 1974.

R. W. Rosenthal, "External Economies and Cores," *J. Econ. Theory*, June 1971, *3*, 182–87.

P. A. Samuelson, "Pure Theory of Public Expenditure and Taxation," in *Public Economics* (J. Margolis and H. Guiton, eds.), St. Martin 1969.