Approximating Linear Programs with

Piecewise Linear Summary Functions: Pseudo Data

with an Infinite Sample

Paul V. Preckel

Thomas W. Hertel*

Staff Paper #86-7

*The authors are assistant professors in the Department of Agricultural Economics, Purdue University, West Lafayette, Indiana 47907.

This research has been partially funded by USDA cooperative agreement #58-319V-5-00134.

I. Introduction

Ever since the pioneering work of Leontief, Koopmans and Dorfman, Samuelson and Solow, linear programming (LP) has played an important role in applied economic research. More recent advances in economic theory and operations research have greatly facilitated the formulation and solution of non-linear models. However, there are many problems which continue to lend themselves to linear analysis. Consider the widespread use of the linear programming approach for industrial process modeling. Recent examples include the Russell-Vaughan model of steel production (1971) and those of Thompson et al. (1977) for the petroleum refining, electric utility and chemical industries. These tend to be popular because of the fixed coefficient nature of specific manufacturing processes, with aggregate factor substitution arising from the use of activities with different relative factor intensities.

Another common use of the LP approach is in the formulation of national planning models where a considerable degree of disaggregation is desired. A good example is provided by the CARD (Center for Agriculture and Rural Development) model of U.S. agriculture. In one version (English et al.) the nation is divided into 105 producing regions, each with 10 land groups and detailed irrigation information. Twelve commodities are produced using 330 alternative crop rotations and a dozen different soil conservation alternatives. Transportation and marketing are also treated in this model. This detailed modeling of agriculture permits analysis of the link between broad national policies and natural resource use and production levels in particular producing areas.

Despite their relative computational simplicity, there are many instances where repeated solution of the LP may be either too time-consuming, too costly, or simply uninformative. This has given rise to a literature on methods for summarizing such models. For example, Griffin (1977a) uses pseudo data generated from an LP describing electricity power generation to estimate a translog, summary cost function. Use of the latter in the Wharton econometric model results in considerable computational savings. [He estimates (p. 113) that direct use of the LP would require it to be solved 100 times for each 20-year macroeconomic forecast.]

For those LP models used in policy analysis, the qualitative results provided by summary functions can prove extremely valuable. Policy analysts often do not have the luxury of waiting a day, or possibly a week, for "what if" questions to be answered by the large-scale, national planning models. Furthermore, they are often only interested in the qualitative results associated with the impact of a particular intervention. In such circumstances it is very desirable to have a set of elasticities summarizing the partial effects of various price and quantity changes on the full complement of endogenous variables in the LP.

The proliferation of studies since Griffin's original application of the pseudo data technique is further testimony to the interest in obtaining summary measures of LP responses [eg., Griffin (1977b, 1978), Kopp and Smith (1980, 1981) Smith and Vaughan (1979, 1980, 1981)]. However, widespread use of this approach has been hampered by doubts about reliability of the summary elasticities. Critics of this method (eg., Madala and Roberts) have pointed out that results may be quite sensitive to sample design. Griffin himself (1982), has expressed concern about the sensitivity of his results to the frequency and range over which the sample points are selected. This problem arises from the fact that the sampling is conducted independently of the underlying model structure. Depending on the location and nature of the LP's basis changes, a given sampling strategy may be effective in one case, but not in another. Similarly, a slight change in sample design may yield very different estimates.

In this paper, a new method for summarizing LP models which takes account of the underlying model structure is described. The selection of sample points is dynamic, and is based on information about the particular LP under consideration. This gives rise to an efficient and relatively robust algorithm, whereby each successive LP solution obtains a rapid reduction in the uncertainty about the true model structure. The power of this method is illustrated by the fact that it is equivalent to the pseudo data approach, when the latter is based on an <u>infinite</u> sample.

Section II introduces the summary function problem in an intuitive way, and illustrates fundamental distinctions between the pseudo data method and the proposed approach. The latter is dubbed "Linear Program Summary Functions" (LPSF). In the third section the LPSF algorithm is developed in detail. Section IV provides a comparison of approaches and sampling strategies based on a relatively simple, well-understood LP model. A final section provides the reader with a summary and the conclusions which can be drawn from this research.

II. Overview of the Summary Function Problem

The major problem associated with current techniques for summarizing the price responsiveness of LP models is perhaps best illustrated in a visual manner. Figure 1 depicts the supply function for soybeans from the farm level model discussed in Section IV. Each of the eleven steps refers to a different optimal basis corresponding to soybean prices which vary over the range between 75% and 125% of the base price. Since soybeans are used in rotation with other crops, as long as total acreage is fixed, there is little incentive to change crop mix for small price movements. This is reflected in the highly inelastic supply function for all price changes between -25% and +20% of the base. However, supply is quite responsive to soybean price changes greater than +20%.

How might one best summarize this supply relationship? The simplest approach involves computing an arc elasticity which evaluates quantity change over the entire range of prices. Since the base case is our point of reference, the percentage quantity and price changes are computed relative to their base levels. This yields an arc elasticity of supply equal to 0.343. Unfortunately this overstates the model's price responsiveness over most of the relevant range.

By contrast the pseudo data technique samples a number of points between the extremes of ± 25% and fits a curve through them. Once again the elasticity may be computed with reference to the base case. The results (see Section IV, for more details) may be shown to depend heavily on the number of points sampled.[1] For example, adding the base price to the two extremes results in a supply elasticity of 0.60. As more points are added this estimate becomes more reliable. Thus, for 5 prices it is 0.51 and for 11 evenly-spaced sample points along the soybean price axis the supply elasticity (evaluated at the base price) drops to 0.39.

As noted above, the sensitivity of estimated elasticities to sampling frequency has been a source of great concern. Without prior knowledge of the model's underlying structure it is impossible to determine a "good" sample design. Furthermore, in many cases successive sample points may do little to reduce uncertainty about the underlying supply relationship. For example, consider the case where the sample consists of eleven evenly spaced points in the range of ±25% (one for every 5% price change). This places 10 of the 11 points along the segments A-E, over which there is almost no supply response.

By the same token the segments F-J are not sampled at all. This is clearly an undesirable situation.

The Linear Program Summary Function (LPSF) technique outlined in the next section takes into account the fact that sample responses are derived from a linear programming model. This knowledge permits us to utilize techniques from sensitivity analysis in order to extract considerable information about the underlying model structure with each additional sample point. Furthermore, we are able to place bounds on the extent of our remaining ignorance. This gives rise to a dynamic sampling process which rapidly reduces this uncertainty.

The advantages of this approach are readily evident from Figure 1. Based on information from the 3 LP solutions at ± 25% and the base price, the LPSF algorithm constructs the segments A, C and K. Furthermore, bounds are placed on the maximum remaining error. Since this is greatest between C and K the next sample point comes between these two segments. (Details of this procedure are provided in the following section.) In particular, segment E is discovered next. The sampling sequence for the remaining segments is determined dynamically (using the same criteria) in the following order: H, B, F, D, I, G, J. After 11 solutions of the LP we have discovered a great deal about the "true" supply function. This is a considerable improvement over the pseudo data technique where the eleven solution case neglected key pieces of the piecewise-linear soybean supply function.

After fitting a smooth summary function to the true response surface, an elasticity may be computed. The resulting base-price supply elasticity is 0.21 which is an improvement over the alternative techniques. Of course none of these alternatives is as reliable as using the piecewise linear supply function directly. As long as the specified price perturbation is along a sampled direction, the LP response can simply be "read off" of the piecewise linear supply function. This gives the same result as would be obtained by re-solving the entire LP. If the price perturbation does not lie along a sampled direction, a differentiable summary function provides a logically consistent method for evaluating the approximate response.

III. Development of the Algorithm

In this section the Linear Program Summary Function algorithm is developed in detail. By way of introduction, as well as for the purposes of comparison, a brief review of the pseudo data technique in our notation is presented here.

Pseudo Data Revisited

Consider the following parametric linear program designed to determine the profit maximizing commodity bundle (X). Prices in the base case (P^0) and resource endowments (R) are given.

$$(1)\quad \max \pi = (P^0 + \alpha D_i)^T X$$

$$\text{subject to:}\quad A X \leq R,$$

$$X \geq 0.$$

The matrix A is comprised of technological coefficients. The vectors D_i ($i=1,\ldots,I$) denote the directions in which the objective function coefficients will be perturbed, and α is a scalar denoting the amount of the perturbation.[2] The pseudo data technique generates data for estimation of the associated profit function by choosing several directions D_i, and values for the perturbation parameter, α.[3] The j^{th} value for direction D_i will be denoted α_{ij}. The optimal objective function values, denoted $\pi^*(P,R)$ (where $P = (P^0+\alpha_{ij}D_i)$), are then recorded for use in the estimation problem.

As Madala and Roberts (1980) point out, given sufficient price variability in the pseudo data set, this restricted profit function may be estimated directly, using ordinary least squares. This problem may be stated as follows (output and resource levels remain constant and are therefore suppressed):

$$(2)\ \min \sum_{i=1}^{I} \sum_{j=1}^{J} [G(P^0 + \alpha_{ij} D_i | \theta) - \pi^*(P^0 + \alpha_{ij} D_i)]^2$$

where $G(P|\theta)$ is the estimated cost function, given the vector of parameters θ. A commonly employed functional form for $G(\cdot)$ is the translog:

$$(3)\ G(P|\theta) = \beta_0 + \sum_{k=1}^{N} \beta_k \ln(p_k) + 1/2 \sum_{m=1}^{N} \sum_{n=1}^{N} \beta_{mn} \ln(p_m) \ln(p_n)$$

This may be viewed as a second-order Taylor series approximation to the logarithm of the true profit function, π^*, in the neighborhood of the unit price vector. Once symmetry and homogeneity are imposed, this yields $1 + 2(N-1) + 1/2[(N-1)^2 - N - 1]$ parameters in the θ vector.

Rather than working with the profit function directly, it is common practice to apply Hotelling's lemma and estimate the associated net supply equations (Griffin, 1977a). In the case of the translog, this yields share equations of the form:

$$(4)\ \frac{\partial \ln G}{\partial \ln p_k} = S_k(P|\theta) = \beta_k + \sum_{m=1}^{N} \beta_{km} \ln p_m.$$

Thus problem (2) becomes one of minimizing the system sum of squared residuals associated with actual and fitted profit shares. Regardless of whether (3) or (4) -- or even both -- are estimated, the pseudo data methodology remains quite sensitive to the choice of sample values (α_{ij}'s). Furthermore, it fails to capitalize on our knowledge about the basic structure of the underlying process model.

Linear Program Summary Functions

The proposed alternative to pseudo data is the Linear Program Summary Function (LPSF). The LPSF is created by a two step process. The first step is to construct a piecewise-linear summary function for the LP. The second step is to estimate a differentiable summary function which is, by some measure, "close" to the piecewise-linear summary function. The most important difference between the summary function and pseudo data methods is that the new method takes into account the fact that the responses due to price changes are actually being generated by a linear program.

Consider, for a moment, the situation wherein we know the optimal objective value of the LP for every reasonable configuration of prices. In this case, the optimal value for each variable, X_i may be determined by computing the derivative of the optimal objective response function with respect to the appropriate price. Hence, there is no further information about the primal solution that can be obtained by solving the linear program. Linear program summary functions are approximations to this optimal objective response function.

What is known about the optimal objective response function as a function of prices? LP theory tells us that this function is convex and piecewise-linear. Also, the points where the optimal function is nondifferentiable correspond to points where basis changes occur. The response function forms a polytope in n+1 dimensional space (where n is the number of prices in the summary). Figure 2 displays the negative of the optimal objective response ($Z^* = -\pi^*$) as a function of two prices. Ideally, we would like to have an efficient method for generating this polytope. Any useful method for describing the polytope would have to store a quantity of information equivalent to a list of the extreme points of the facets.[4] Unfortunately, this seems to be impractical, since a facet of the polytope may potentially have the same number of extreme points as there are bases for the LP. Furthermore, the computation of each extreme point requires the solution of a linear system. Thus, a simpler problem must be defined.

Consider again how the pseudo data technique is used to generate data. First a base price, about which the approximation will be constructed, is chosen. Second, the response function is evaluated at various prices along line segments in price space that intersect the base price. This would seem to suggest a useful intermediate step between finding the entire response surface on the one hand and sampling only a few points on the other. Thus, the true response function is constructed over selected lines in price space that intersect the base price (P^0). Figure 3 displays the optimal response function with two lines imposed in price space (ab and cd). Their projection onto the optimal response function results in the piecewise linear arcs a'b' and c'd'. In price space these lines are parallel to the price axes. Thus, they move only one price at a time so that $D_1 = (1, 0)$ and $D_2 = (0, 1)$. The generalization to lines which are not parallel to the axes is straightforward and implies two non-zero elements in the D_i vector.

To understand the approximation technique, it is useful to consider a "slice" of the graph in Figure 2 taken parallel to the P_1 axis and through the base price. This captures that portion of the optimal objective response function given in Figure 4. The vertical axis in this figure corresponds to the level of the (negative) optimal objective response, $Z^*(P)$, where:

$$P = P^0 + \alpha_1 D_1$$

As before, P^0 denotes the base price about which the local approximation to the response function will be constructed, D_1 denotes the direction in which prices are being perturbed, and α_1 denotes the amount of the perturbation. In this subspace, the response function retains the properties of the full surface -- it is concave, piecewise linear and points where the function is nondifferentiable correspond to basis changes.

At this point it must be noted that even the task of determining $Z^*(P)$ above the lines ab and cd in Figure 3 may become impractical when the number of linear pieces is very large. Hence, a method for efficiently generating a good approximation to the response function using a minimum of LP evaluations is needed. (Note that by using techniques from sensitivity analysis each evaluation of the LP yields the slope, intercept and end points of one linear piece.) Since the idea is to create a local approximation, the function only needs to be constructed over a limited range of perturbations about the base price. In the context of Figure 3, this means that the true response function is only sampled between the lower and upper bounds given by L_1, and U_1, respectively.

The piecewise linear approximation algorithm commences by determining the linear pieces that contain the extreme values for α_i (call these "endpieces"). At any stage of the approximation process there is a collection of intervals over which the response function is known and another collection of intervals where it is unknown. For example, in Figure 5 segments containing P^0 and each of the endpieces have been found. The remaining 2 pieces are "unknown" since they have not yet been sampled. At this point the approximation may simply be "completed," or the sampling process may be continued.

Completion of the approximation is achieved in the following manner: over the intervals where the function is unknown, extend the linear pieces from the adjacent intervals (where the function is known) until they intersect. For example, this involves extending c'e' in Figure 5 out to g'. Similarly the segment g'f' is added. This simple procedure always results in a function which is defined over the entire interval and which shares the most important properties of the true response function (piecewise linearity and concavity). One measure of the error made by the completion process is the integral of the absolute value of the difference between the true and approximate functions over the range: $L_1 \leq \alpha_1 \leq U_1$. This error measure is bounded due to the concavity of the response function. A lower bound on the function itself may be computed by connecting the end points of the linear pieces of the intervals adjacent to an unknown interval (eg., e'f' in Figure 5). On the other hand, the completion described above provides an upper bound on the true, but unknown response function. Hence, the error in forming the completion can be no greater than the integral of the difference between these two functions. This integral corresponds to the sum of the areas of the shaded triangles in Figure 5. When this global error bound becomes sufficiently small, further sampling in that direction stops and the resulting completion will be used in place of the true response function.

In addition to permitting formulation of an intuitive stopping rule, the error bounds for the individual intervals may be used as the basis for a dynamic sampling process. In particular, the next sample value of the perturbation parameter, α_1, may be placed in that interval with the largest error bound. The precise value for α_1 is chosen so that half of the area of the error triangle lies on either side of it. In the context of Figure 5, this results in the next sample point along D_1 being determined by α_{new}. This method of dynamically selecting sample points results in a rapid reduction in approximation error, even when the error triangles are highly skewed.

In addition to the stopping rule based on global error, it may be desirable to exclude some unknown intervals from consideration as candidates for the next evaluation. For instance, if an interval is very short or the height of the error triangle is small, it may be desirable to ignore the interval. Similarly, if the change in slope between the adjacent known intervals is small there may be little gained by evaluating the LP in the unknown interval. Finally, for practical purposes it is useful to specify a maximum number of linear pieces that will be sought along each direction.

Once the piecewise linear response function has been constructed it may be used directly, or as with the pseudo data technique, it may be used as the input to an estimation routine. Unlike the pseudo data technique, the estimation routine for a differentiable summary function must be designed to accomodate an <u>infinite</u> number of data points. For purposes of illustration, consider the problem of fitting a differentiable summary function to the optimal objective response function in Figure 5. As noted above, given sufficient price variability, it is possible to fit the cost function directly, using a least squares criterion. In the continuous (infinite sample) case the analogous parameter estimation problem becomes:

$$(5) \qquad \min_{\theta} \sum_{i=1}^{I} \int_{L_i}^{U_i} [G(P^0+\alpha_i D_i|\theta) - \pi^*(P^0+\alpha_i D_i)]^2 \, d\alpha_i$$

Where θ denotes the parameters of the differentiable summary function, $G(\cdot)$, and $[L_i, U_i]$ denote the extreme values for the perturbations from the base price (along the i-th direction, D_i). (Note that this is a least squares problem using an infinite number of data points.) This problem is illustrated in one dimension with a quadratic summary function in Figure 6, where the objective is to minimize a measure of the shaded area. Just as the pseudo data estimation could be conducted with a system of behavioral equations (by appealing to Hotelling's lemma), so can (5) be reformulated to apply to such a system. Once again we assume the translog form for $G(\cdot)$ so that this system is expressed in shares rather than quantities.

For many of the desirable functional forms, the integrals in question may have no closed form expression (e.g. translog demand equations). Furthermore, even when closed form expressions do exist, the objective is nonlinear.

Fortunately both of these obstacles are easily surmounted. For the purposes of this study, the estimation problem was solved using the MINOS 5.0 nonlinear optimization package (Murtaugh and Saunders [1983]). The integration of the objective terms in (5) is handled numerically using routines that have been specialized for the optimization framework (see Kaylen and Preckel [1986]). Details on the performance of this method are discussed in later sections.

IV. Empirical Implementation

In order to illustrate use of the proposed LPSF algorithm, it has been applied to a relatively small, well-understood process model. This permits the comparison of a range of sampling issues, including questions of sampling frequency and direction. For the initial analysis, Griffin's approach -- perturbing one price at a time -- is employed to demonstrate that as sampling frequency grows, the difference between pseudo data estimates and those of LPSF diminishes. The second step involves perturbing relative prices simultaneously (diagonal directions). We find that this changes the results dramatically, and conclude that the hitherto unexplored issue of sampling direction is crucial to the construction of an accurate summary function.

A. The Process Model

The process model employed here is a modified version of the Purdue Crop Budget Model (B-9) (McKinzie, et al.). The B-9 is among the most extensively validated of all process models, having been used daily by extension and research staff, graduate and undergraduate students, as well as by thousands of midwest farmers over the course of its 15-year evolution. It is a linear programming formulation of a profit maximizing farm firm. The formulation utilizes highly detailed information including the farm's machinery working rates, available time for working in the field during different periods of the production year, and cultivation practices.

Timing of production activities is given particular attention in the B-9 model. Expected crop yields are generally acknowledged to decline as planting (and harvesting) of the crop are delayed. However, it is not economical to maintain the machinery necessary to plant (harvest) all of the crop at one time. This process model captures tradeoffs between the cost of larger, more expensive machinery sets and the benefits associated with improved yields due to timeliness of planting and harvesting. The latter effect serves to promote diversification among crop outputs. While corn is often the most profitable crop to be planted during late April and early May, soybeans may be the preferred alternative in late May. This occurs since soybean yields decline at a slower percentage rate than do corn yields as planting is delayed. In addition, there are significant economies from rotating corn and soybeans. Costs rise for corn grown continuously on the same land. Yields decline for both continuous corn and continuous soybean crops. The effect of including these complementarities from crop rotation is to give the product transformation curve for corn and soybeans more curvature in the region of equal acreages. Thus, holding other things constant, a greater change in relative prices is required to achieve a given amount of substitution between these crops.

B. Sampling Along the Axes

One of the major issues in summary function analysis pertains to the directions in price space (D_i's) along which the sampling should be conducted. As noted, Griffin perturbs a single price at a time, thus sampling along the axis directions. This provides sufficient variability to infer cross-price effects if it is the demand (share) equations which are being estimated. (Of course, if the profit function were to be estimated directly, some combination of simultaneous price movements would be required to estimate interaction terms.)

Table 1 presents a selection of gross supply elasticities from the translog profit functions (estimated in share equation form) for different data sets.[5] The first row corresponds to the estimates obtained by numerical integration designed to minimize the squared difference between the fitted translog share equations and those implied by the piecewise linear profit function with a maximum of 11 pieces being evaluated per direction. The next three rows A.2-A.4 of elasticities correspond to pseudo data estimation of share equations with 11, 5, and 3 points, respectively (along each axis).[6]

Focusing initially on the own-price elasticities of supply and demand, note that as the number of sample points increases from 3 to 5 to 11, the pseudo data estimates generally move towards the elasticities derived from the LPSF technique. This suspicion is confirmed by a more formal measure of distance between the full matrices of elasticities. The Frobenius norms[7] of the difference matrices, derived by subtracting each of the pseudo data elasticity matrices from the LPSF matrix, are given in the last column of Table 1. They indicate that this distance declines from 1.96 to 1.75 and then to 1.49 as the number of sample points increases.

The own-price elasticities in part A of Table 1 are roughly of the same order of magnitude, however, the situation for selected cross-price elasticities is somewhat different. In particular, the demand elasticity for machinery with respect to the price of soybeans is much greater for the LPSF estimates, as is the labor/corn price elasticity of demand. The sign of the wheat/soybeans cross price elasticity actually differs in sign. This is quite disturbing and leads us to question the reliability of the cross-price elasticities when sampling is solely along the axes.

C. Introducing Simultaneous Variation in Relative Prices

While the axis directions provide sufficient price variability for estimating cross-price effects when one works with share equations (instead of the profit function itself), there are many cases where the resulting estimates may be quite unreliable. Consider, for example, the following two functions:

$$f(P) = \begin{cases} 1 - P_1 & \text{if} \quad P_1 \geq P_2 \text{ and } P_1 \geq -P_2 \\ 1 + P_1 & \text{if} \quad P_1 \leq P_2 \text{ and } P_1 \leq -P_2 \\ 1 - P_2 & \text{if} \quad P_2 \geq P_1 \text{ and } P_2 \geq -P_1 \\ 1 + P_2 & \text{if} \quad P_2 \leq P_1 \text{ and } P_2 \leq -P_1 \end{cases}$$

$$g(P|\epsilon) = \begin{cases} f(P) & \text{if} \quad |P_1| \leq \epsilon \text{ or } |P_2| \leq \epsilon \\ 1+\epsilon-P_1-P_2 & \text{if} \quad P_1 \geq \epsilon \text{ and } P_2 \geq \epsilon \\ 1+\epsilon+P_1-P_2 & \text{if} \quad P_1 \leq -\epsilon \quad P_2 \geq \epsilon \\ 1+\epsilon+P_1+P_2 & \text{if} \quad P_1 \leq -\epsilon \quad P_2 \leq -\epsilon \\ 1+\epsilon-P_1+P_2 & \text{if} \quad P_1 \geq \epsilon \quad P_2 \leq -\epsilon \end{cases}$$

where the second function, g, is parameterized by ϵ, a very small positive number. Level sets for the functions f and g are displayed in Figure 7 (ϵ is chosen to be indistinguishable from zero in this graph). Letting the base price be the origin, it is clear that these functions and their gradients agree when only one price is perturbed at a time. However, given that they agree above the axes and that they are both concave, these functions are as far from each other as possible away from the axes.[8] While this is an artificial example, it illustrates the fact that cross price effects may only be reliably measured by considering diagonal directions which perturb several prices at once (e.g. the dashed lines in Figure 7). For the LP model introduced above, consideration of all possible combinations of price movements yields 63 diagonal directions.[9]

Selected elasticities from the LPSF based on introduction of the diagonal sample design are presented in the first line of Part B in Table 1 (line B.1). Several of the cross-price terms change dramatically, when compared with line 1 of Part A, the most notable case is that of the wheat/soybeans cross-price elasticity which changes from +0.01 to -0.41. Of course, the off-diagonal net supply elasticities are ultimately linked to their own-price counterparts via the homogeneity condition. Thus, the latter may also be affected, as is indeed evident from Table 1. However, it is reasuring to note the similarity in gross supply elasticities for the three crop outputs in this model.

D. <u>Efficiency Gains From Dynamic Sampling</u>

Having noted the importance of introducing non-axis directions into the sample design, it is interesting to turn to the issue of sampling frequency, given a specific set of directions. In this section the 63 diagonal directions serve as the basis for examining the early efficiency gains from additional LP evaluations.

The LPSF algorithm begins by evaluating the LP at base prices and at the extreme point in each sampling direction. LPSF estimates based on these three segments in each of the 63 directions are given in line B.3 of Table 1. The corresponding pseudo data estimates are given on line C.3 of the same table.

Each of these lines may be compared to the more accurate estimates in lines B.1 and C.1, respectively. The distance between the full matrices is given by the Frobenius norms in the final column of rows B.3 and C.3. Thus, the LPSF estimates with three evaluations are somewhat closer to their "true" LPSF counterparts (Frobenius norm = 0.67), than are the pseudo data estimates with three points to those with 21 points (Frobenius norm = 0.72). This is not surprising, given that the LPSF approach is based on three line segments whereas the latter estimates are based on three points. The longer the length of these 3 segments, the greater the advantage of the LPSF approach.

Of course there are additional advantages to the LPSF approach. In particular, the dynamic sampling procedure stands out. As noted above, additional LP evaluations are selected in a manner which brings about a rapid reduction in the possible approximation error, whereas the pseudo data approach is oblivious to accumulated information about model structure. (For this example problem, the reduction in the total error during the early stages of sampling is typically on the order of 60% per L.P. evaluation. During later stages - as the number of unknown intervals increases - the reduction is somewhat less; i.e. 20% per L.P. evaluation.) This distinction is reflected in the speed with which the distances between the elasticities in B.1 and B.3, on the one hand, and those in C.1 and C.3, on the other, are reduced as additional sample points are added. Lines B.2 and C.2 of Table 1 provide estimates of the LPSF and pseudo data elasticities when two more LP evaluations are permitted along each of the 64 directions. As expected, the efficient selection of additional sample points results in more rapid convergence using the LPSF algorithm. In particular, the two new evaluations reduce the Frobenius norm by 2/3's, as opposed to only one-half in the case of pseudo data.

A final point of interest in this empirical section arises from evaluation of the distance between the 7 and 63 direction summary functions. The Frobenius norm for this difference matrix under the LPSF approach is 1.91. The comparable number for pseudo data (11 observations per direction) is 0.94. These distances are generally larger than those separating elasticities based on shared methods and directions with fewer points per direction (Parts B and C of Table 1). The implication is that these summary function elasticities are more sensitive to sampling direction than they are to sampling frequency.

V. Summary and Conclusions

In this paper we have introduced a new method for summarizing the economic content of linear programming models. Dubbed "Linear Program Summary Functions" (LPSF), it provides considerably more reliable summary elasticities than the pseudo data method introduced by Griffin in 1977. In fact the LPSF approach may be viewed as pseudo data with an infinite sample.

The increased reliability of the LPSF methodology stems from its efficient, dynamic sampling of the underlying model structure. This sampling strategy is designed to maximize the reduction in approximation error between a constructed, piecewise linear profit function and the model's true underlying response function. At any point in time the sampling process may be halted and the constructed response function completed. An upper bound on the remaining error may be computed and provides a measure of the quality of this approximation. Of course the latter may be improved upon at some later date, if increased accuracy is desired. Fitting a differentiable summary function to the completed response function is often desirable and a method for doing so is also presented in this paper. This enabled a direct comparison with the pseudo data approach.

The difference between LPSF and pseudo data methods for summarizing LPs was further explored in Section IV, using a relatively simple model. As expected, increases in the number of pseudo data points yielded summary elasticities which were generally closer to their LPSF counter parts. The greater efficiency of the LPSF sampling procedure was also borne out in this example.

Perhaps the most important discovery in the context of the numerical example is the importance of diagonal directions in determining the size and sign of cross-price elasticities. While it is technically possible to estimate these second-order interaction effects with axis sampling only (via use of Hotelling's lemma), the results change significantly when simultaneous price variation is introduced. This problem was anticipated with a hypothetical example (Figure 7), but it is significant that this issue crops up in a commonly employed model. As a result, we hypothesize that introduction of additional sampling directions may be relatively more important than increasing the frequency of sampling along given (presumably axis) directions. This is a hypothesis which deserves further examination as a part of future research efforts in this area.

Footnotes

1. The results also depend on the direction of sampling, choice of functional form and estimation technique. The estimated elasticities in this section are based on axis sampling directions (perturbation of one price at a time) and least squares estimation of a system of translog share equations. See Sections III-IV (esp. footnote 3) for more details.

2. Griffin chooses the D_i to be unit vectors. This means that he is perturbing one price at a time.

3. Griffin (1977a) chooses values for α so that (with the directions chosen to be unit vectors) the individual prices are set to 50%, 80%, 90.5%, 110%, 125%, and 200% of the base values.

4. A facet of the polytope is defined by the intersection of the polytope and a plane in n+1 dimensional space where the intersection contains only boundary points of the polytope.

5. To make the axis and diagonal directions of comparable length, relative price changes of ±67% were evaluated for both axis and non-axis directions. Thus, these axis direction elasticities are different from those in the introduction, where ±25% relative price changes were employed.

6. Pseudo data points were obtained directly from the same picewise linear profit function which was used to fit the LPSF.

7. The Frobenius norm is defined as $[\sum_i \sum_j d_{ij}^2]^{1/2}$ where the d_{ij} represent elements of the difference matrix.

8. A small technical point - the functions are as far from each other as possible in the limit as ϵ approaches zero.

9. The 63 diagonal directions are appropriate for constructing a seven price summary function. These directions may be constructed by considering a seven-dimensional hypercube having sides of length two and baricenter at the origin. The coordinates of the corners of the hypercube are (±1, ±1, ±1, ±1, ±1, ±1), yielding 2^7 directions from the origin. Since we consider both positive and negative perturbations ($\alpha < 0$ as well as $\alpha \geq 0$), there are twice as many corners as there are distinct directions. Hence, there are $2^6 = 64$ distinct diagonal directions. One of these directions, (1, 1, 1, 1, 1, 1, 1), increases (or decreases) all prices simultaneously. Since the objective function is homogeneous in the seven prices, this amounts to a rescaling of the objective. Hence, that direction may be excluded from the sample design, leaving us with $2^6 - 1=63$ diagonal directions in the sample design. Note that it is unnecessary to include the axis directions in our sample design due to homogeneity of the L.P. in our seven prices. Perturbation along an axis direction, amounts to increasing (or decreasing) a single price relative to all others. Since base prices have all been scaled to one, perturbation along a direction with all but one component equal to -1 and the remaining component equal to +1 increases (or decreases) one price relative to all others. Hence, these directions are equivalent, and it is sufficient to sample only along the 63 directions described above.

References

English, Burton C., Klaus F. Alt, and Earl O. Heady. A Documentation of the Resources Conservation Act's Assessment Model of Regional Agricultural Production, Land and Water Use, and Soil Loss. CARD Report No. 107T, Iowa State University, Ames, Iowa, 1982.

Griffin, James M. "Long-run Production Modeling with Pseudo-data: Electric Power Generation," Bell J. Econ. 8(Spring 1977a):112-27.

______________. "The Econometrics of Joint Production: Another Approach," Rev. Econ. & Stat. 59(Nov. 1977b):389-97.

______________. "Joint Production Technology: The Case of Petrochemicals," Econometrica 46(1978):379-96.

______________. "Pseudo Data Estimation with Alternative Functional Forms," in V. K. Smith (ed.) Advances in Applied Microeconomics, Vol. II, JAI Press, 1982.

Kaylen, M.S. and P. V. Preckel. "MINTDF User's Guide," Station Bulletin (forthcoming), Department of Agricultural Economics, Purdue University, West Lafayette, Indiana, 1986.

Kopp, R. J., and V. Kerry Smith. "Measuring Factor Substitution with Experimental Evaluation," Bell J. of Economics,

______________. "Measuring the Prospects for Resource Substitution Under Input and Technology Aggregations," in Modeling and Measuring Natural Resource Substitution,

Madala, G. S. and B. Roberts. "Alternative Functional Forms and Errors of Pseudo Data Estimation," Rev. Econ. & Stat. 62(1980):323-27.

McKinzie, L., T. W. Hertel and P. V. Preckel. Incorporating Endogenous Input Choices and Rotational Corn-Soybeans Into an Indiana Cropping Model (B-29), Purdue University Ag. Experiment Station Bulletin No. 461, 1984.

Murtagh, B. A. and M. A. Saunders. "MINOS 5.0 User's Guide," Technical Report SOL 83-20, Department of Operations Research, Stanford University, Stanford, California, December 1983.

Russell, C. S. and W. J. Vaughan. Steel Production: Processes Products and Residuals, Baltimore: Johns Hopkins University Press, 1971.

Smith, V. K. and W. J. Vaughan. "Some Limitations of Long Run Production Modeling with Pseudo Data," J. Industrial Econ. 48(1979):201-207.

______________. "The Implications of Model Complexity for Environmental Management," J. Env. Econ. and Management 7(1980):184-208.

______________. "Strategic Details and Process Analysis Model for Environmental Management," Resources and Energy 3(1981):39-54.

Thompson, R. G., J. Callaway, et al., Environment and Energy in Petroleum Refining, Electric Power and Chemical Industries, Houston: Gulf Publishing, 1977.

Table 1. Selected Gross Elasticities of Supply and Demand

	Own-Price Elasticities*							Cross-Price Elasticities				Distance
	C.	S.	W.	L.	M.	S/D	O.	L./C.	W./S.	M./S.	M./W.	Frobenius Norm
Part A. 7 Axis: Directions												
A.1 LPSF: 11 pieces	1.70	.83	2.42	-.15	-.26	-.18	-.03	.86	.01	-.27	.05	0
A.2 PDATA: 11 points	1.52	.86	2.41	-.14	-.15	-.11	-.16	.53	-.27	-.11	.10	1.49
A.3 PDATA: 5 points	1.32	.78	2.64	-.18	-.15	-.07	-.06	.47	-.32	-.13	.20	1.75
A.4 PDATA: 3 points	1.37	.87	2.40	-.20	-.17	-.05	-.02	.45	-.39	-.12	.23	1.96
Part B. 63 Diagonal Directions												
B.1 LPSF: 11 pieces	1.66	.90	2.44	-.11	-.14	-.06	-.12	.32	-.41	-.45	-.01	0
B.2 LPSF: 5 pieces	1.71	.90	2.35	-.14	-.14	-.05	-.11	.31	-.36	-.43	-.03	0.22
B.3 LPSF: 3 pieces	1.83	.97	1.99	-.18	-.17	-.04	-.11	.34	-.22	-.46	-.09	0.67
Part C. 63 Diagonal Directions												
C.1 PDATA: 11 points	1.63	.87	2.36	-.14	-.18	-.09	-.07	.33	-.39	-.26	-.04	0
C.2 PDATA: 5 points	1.49	.77	2.32	-.15	-.16	-.10	-.01	.26	-.26	-.36	-.01	0.35
C.3 PDATA: 3 points	1.54	.84	1.90	-.20	-.20	-.12	-.01	.28	-.15	-.32	-.01	0.72

C. = Corn   S. = Soybeans   W. = Wheat   L. = Labor   M. = Machinery   S/D = Storage and Drying
O. = Other Inputs

Figure 1: Soybean Supply Response Along Axis Direction

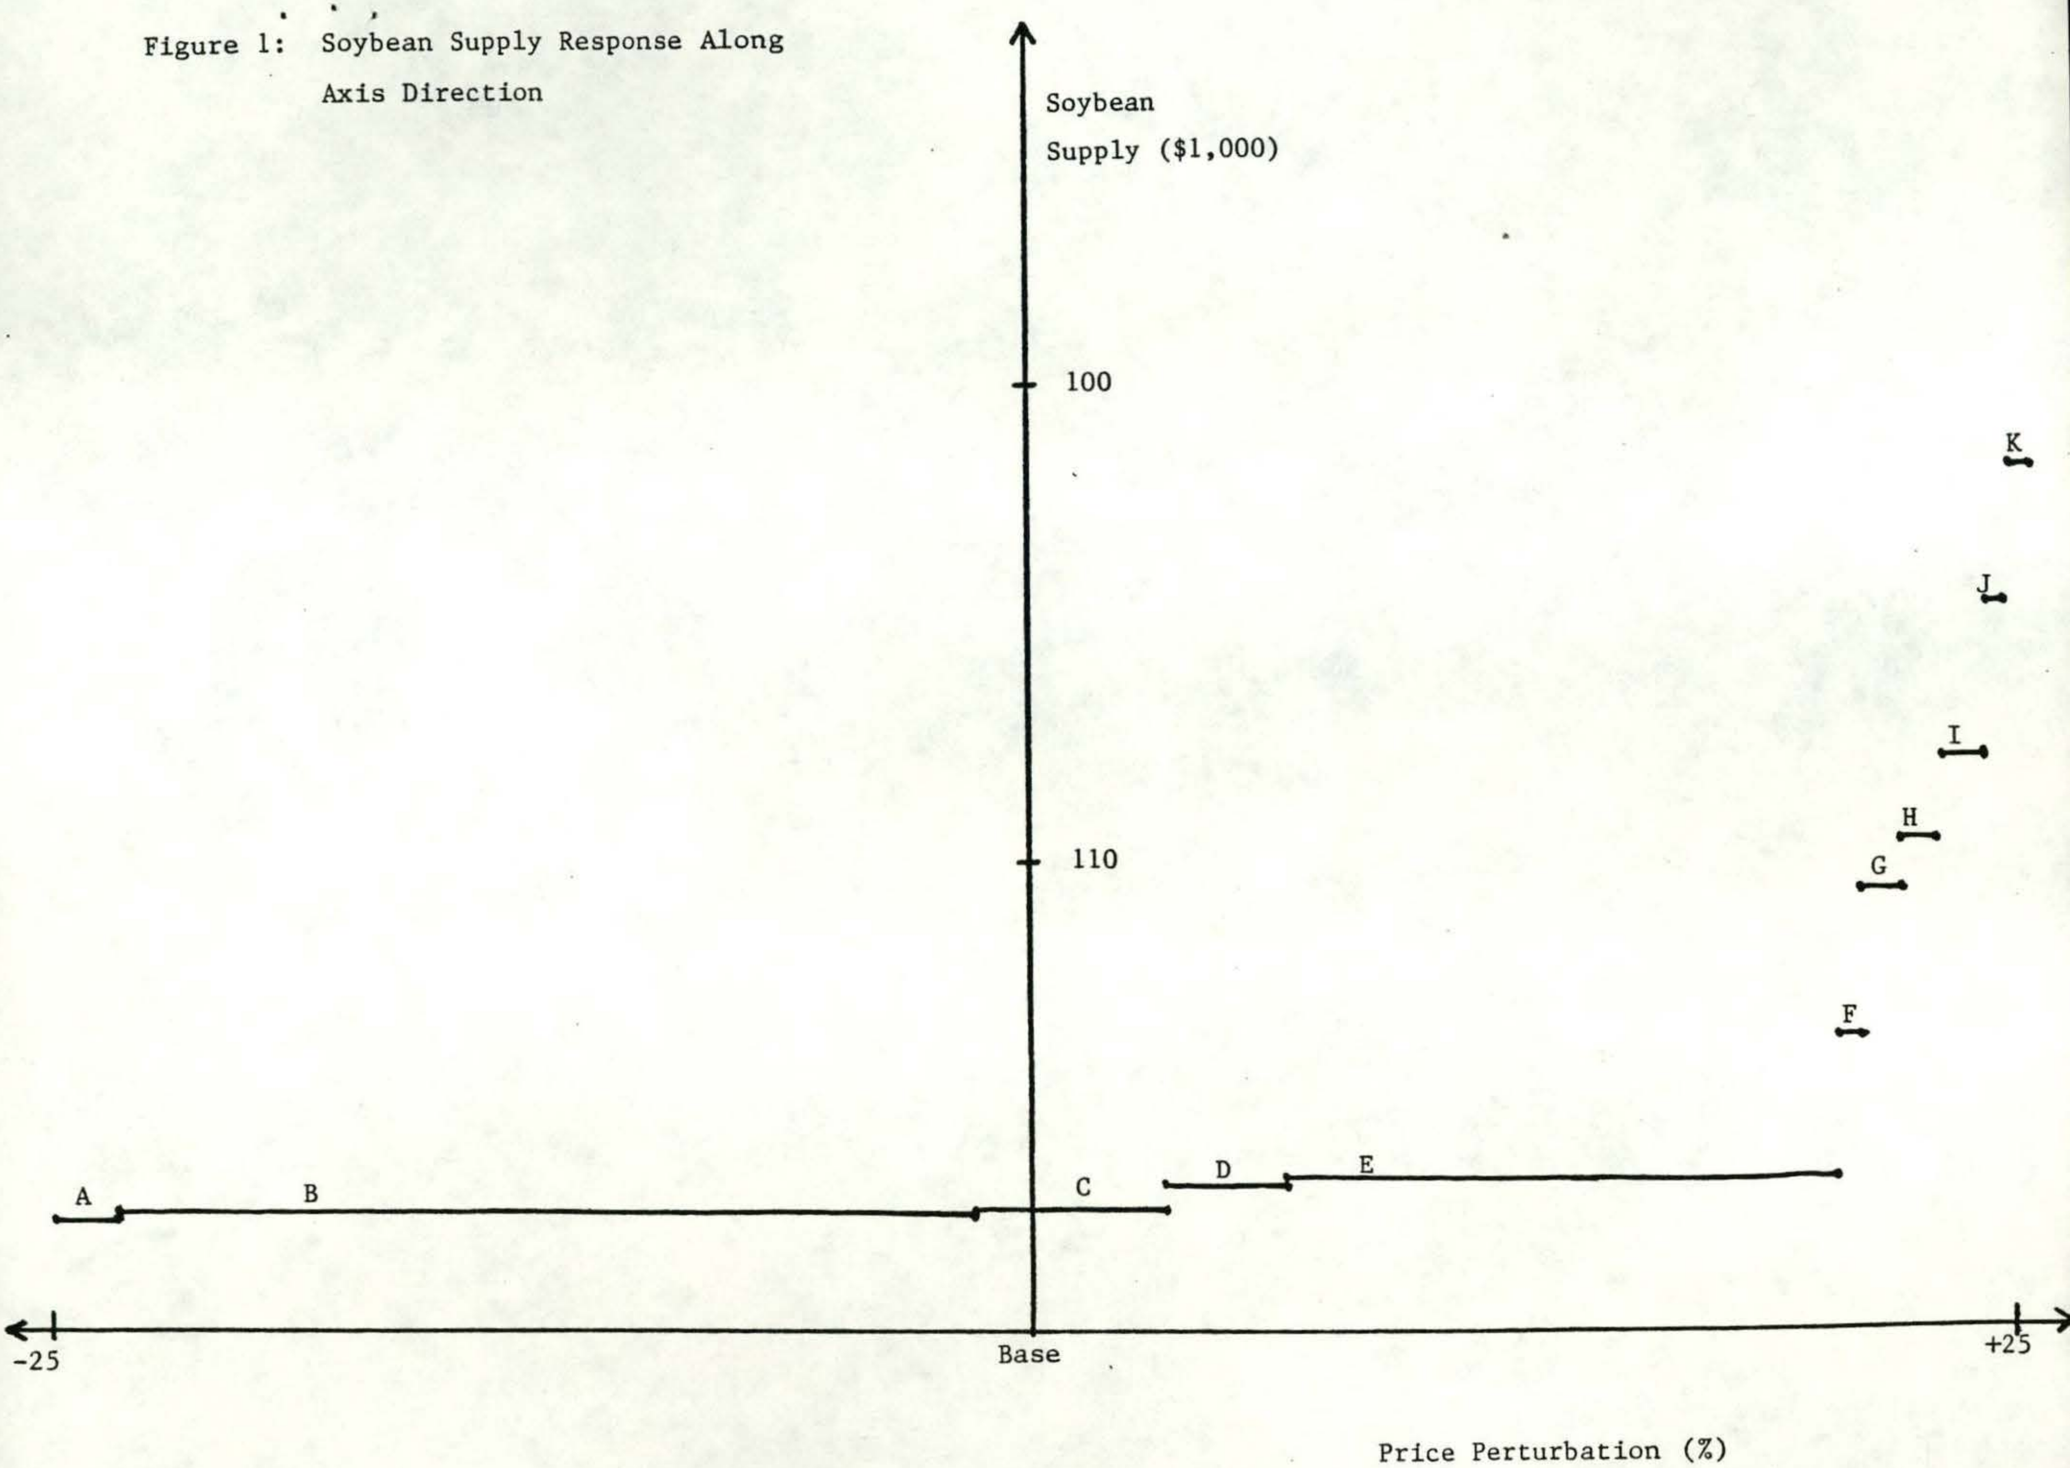

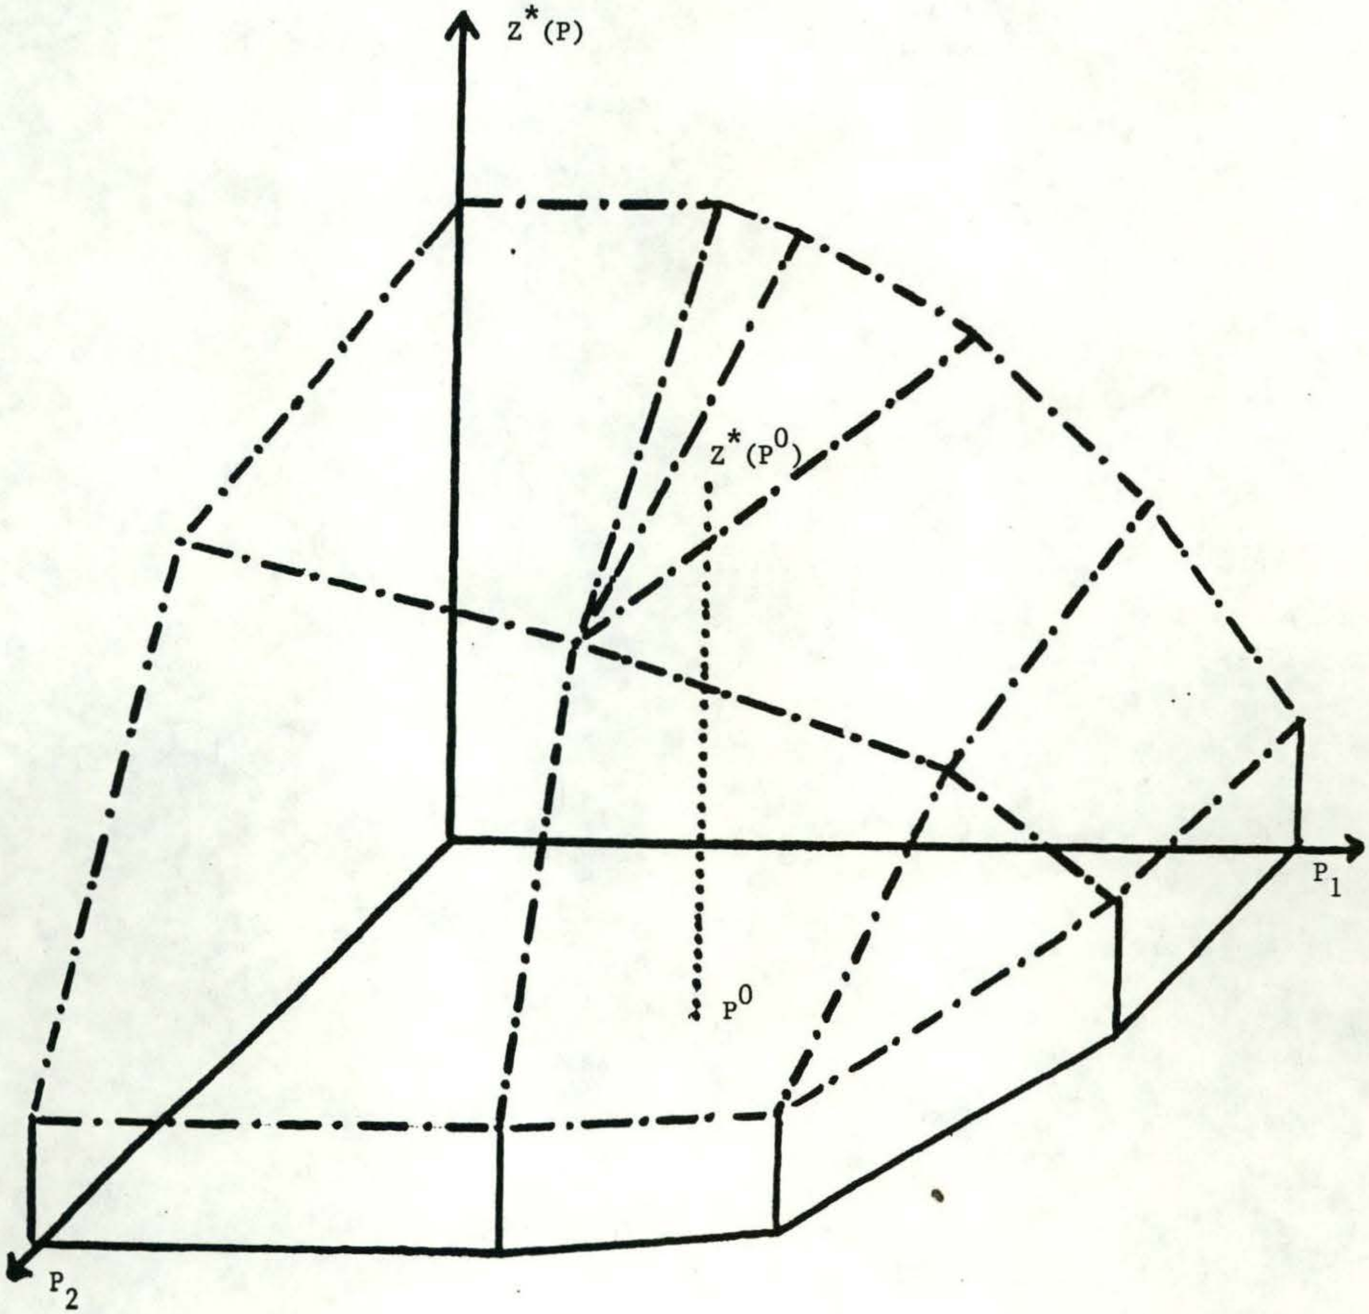

Figure 2: The Response Function

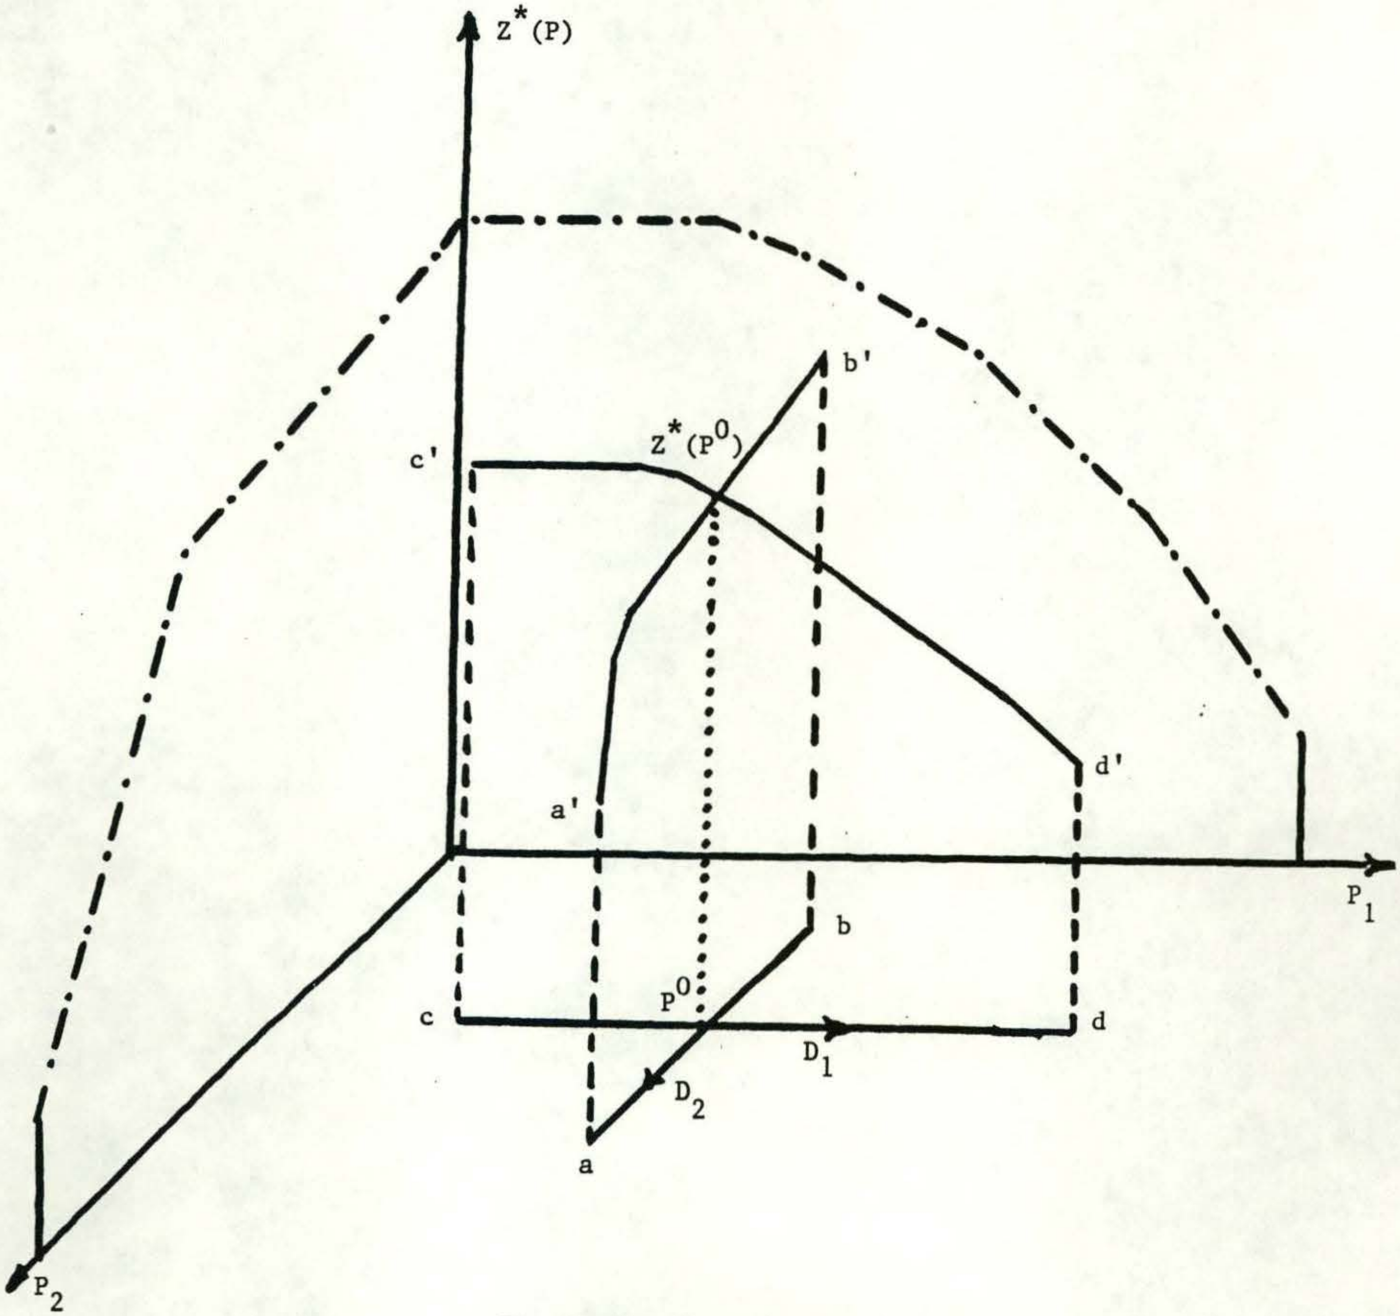

Figure 3: Sampling Directions in Price Space

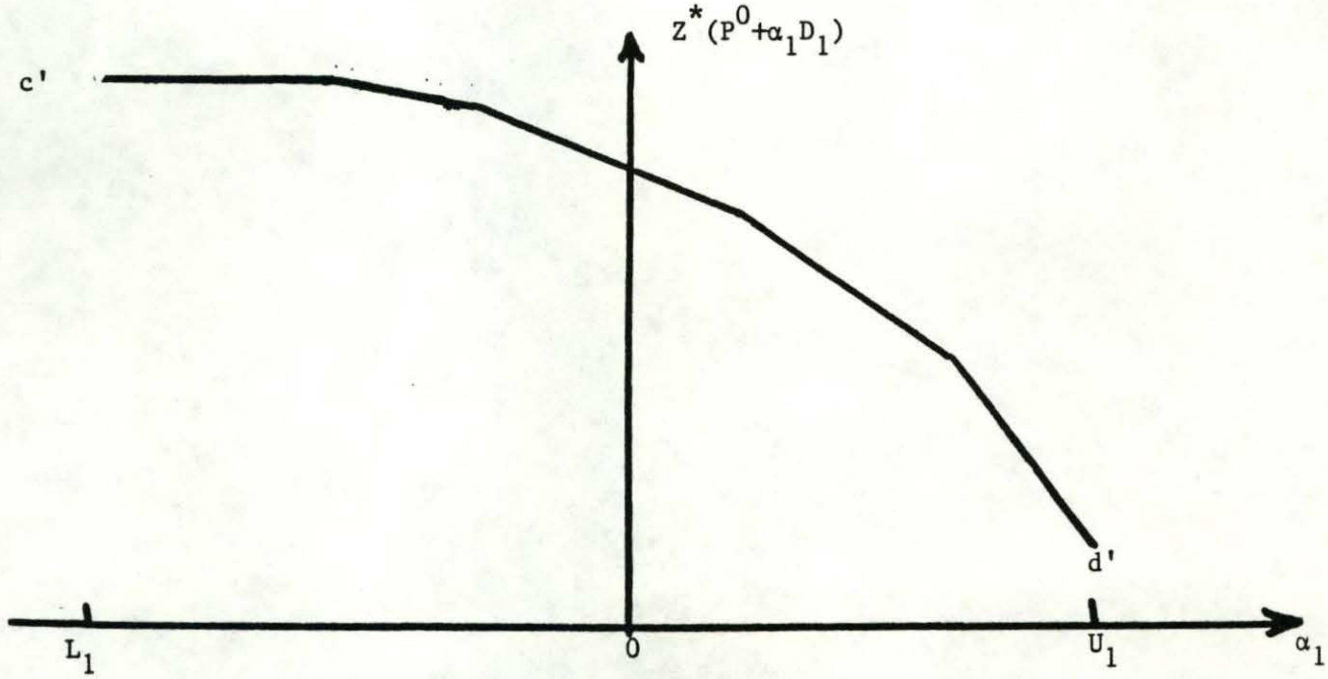

Figure 4: The Response Surface Above a Sampling Direction in Price Space

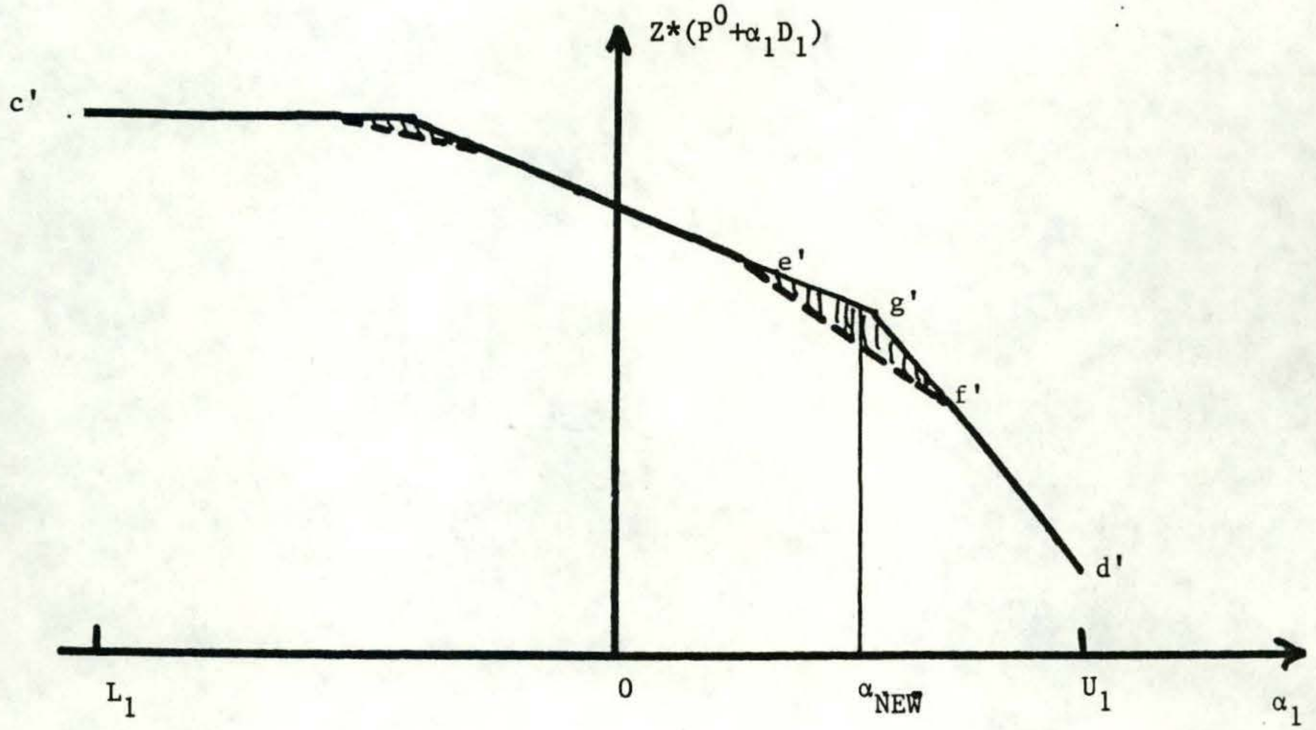

Figure 5: The Approximation Process

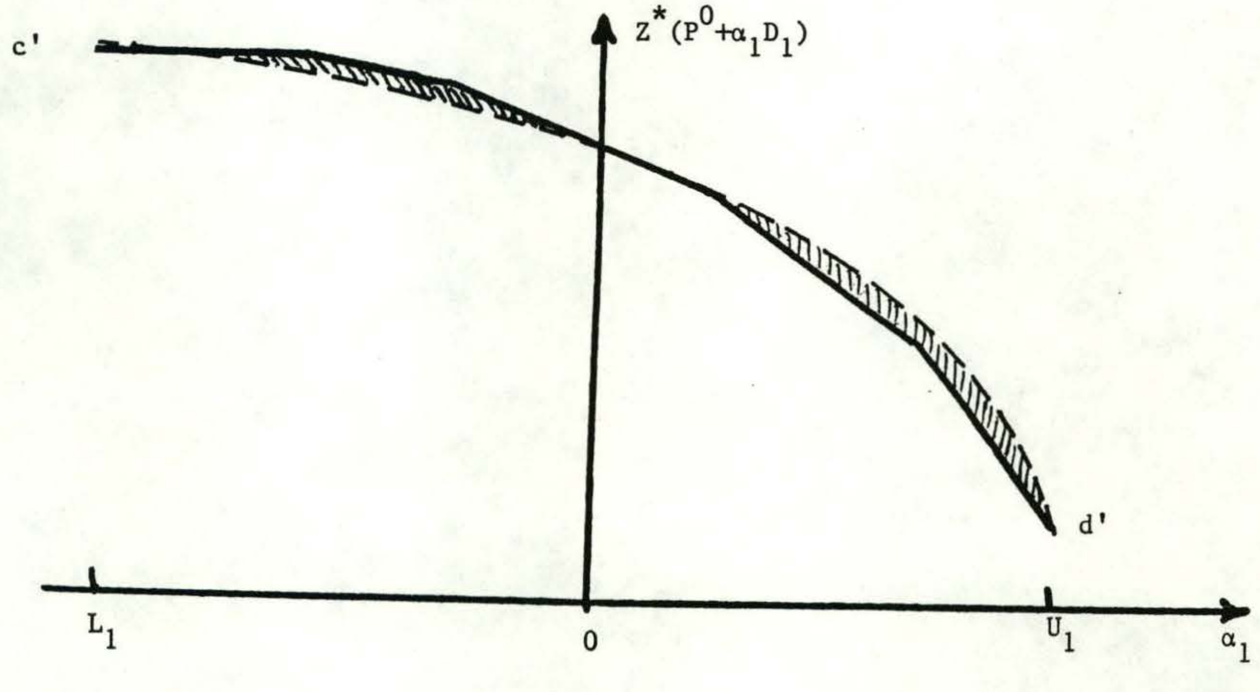

Figure 6: The Differentiable Approximation

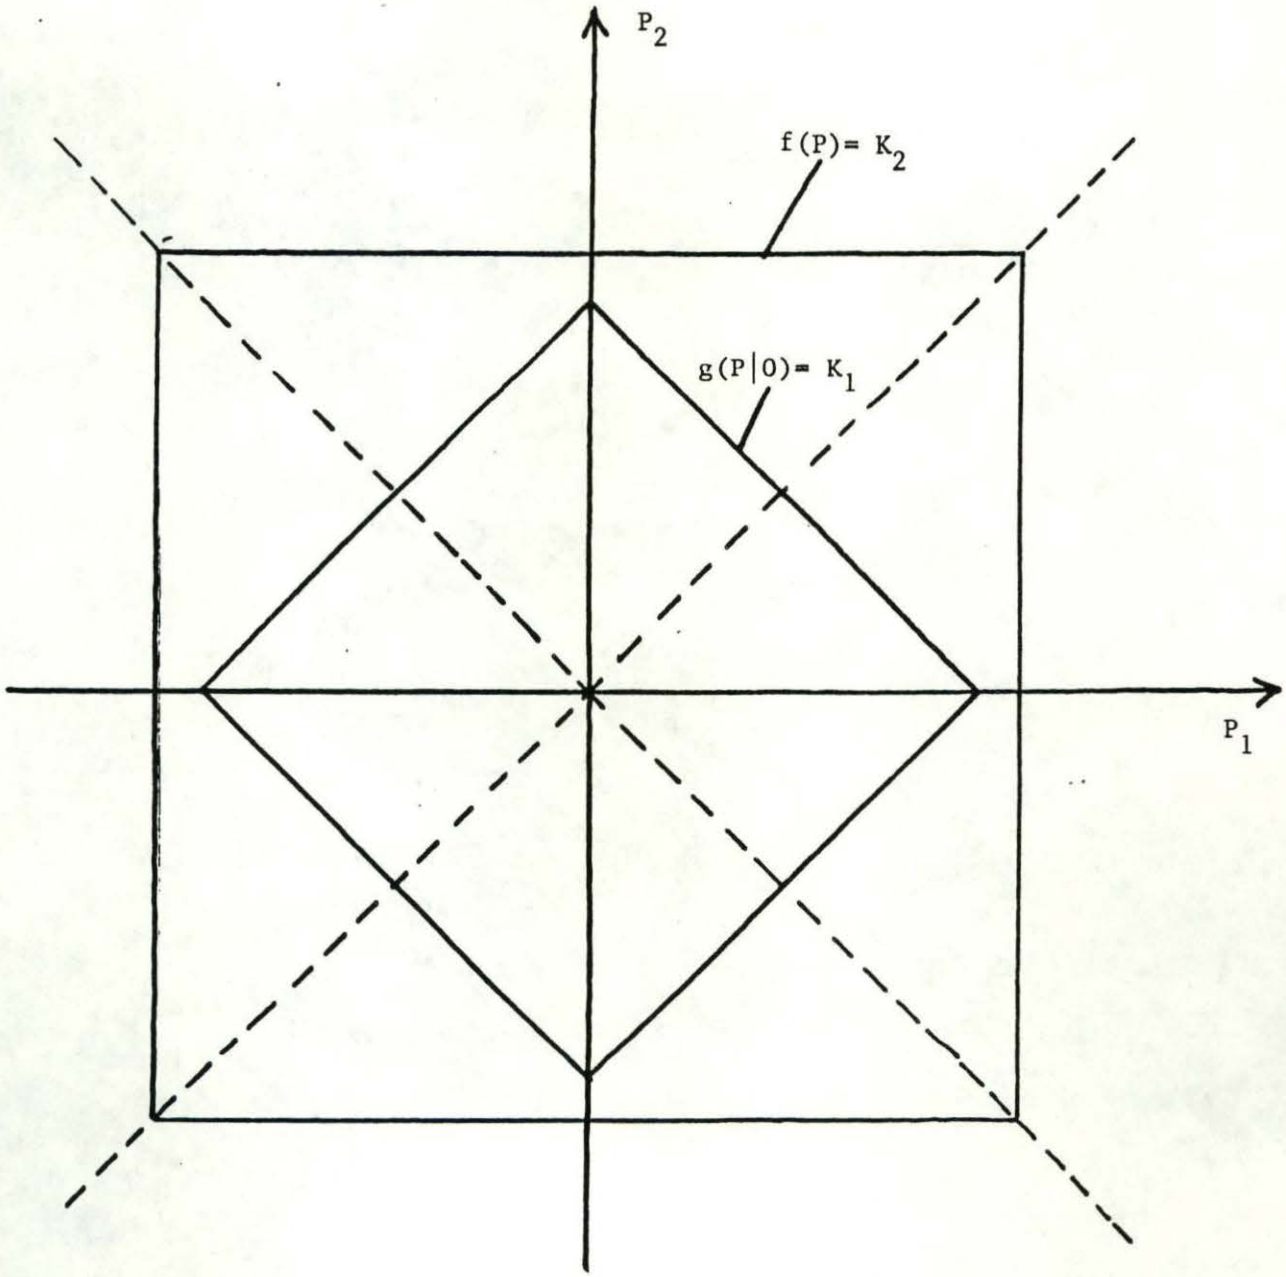

Figure 7: Level Sets for Two Response
Functions in Two-Price Space