Stabilizing Pakistan's Supply of Wheat: Issues in the Optimization of Storage and Trade Policies

THOMAS C. PINCKNEY*

INTRODUCTION

The variability of Pakistan's wheat production in recent years has highlighted the need for the country to develop a storage and trade strategy which takes into account the needs of producers, consumers, and the government. Production variability is a problem despite the fact that cereal production is considerably less unstable in Pakistan than in most other countries in the world, primarily because of the presence of irrigation works [14]. Nevertheless, the country has experienced three significant shortfalls in wheat production in the last nine years, and is presently faced with the different but no less severe problem of an exceptionally large crop. A strategy is needed for dealing with both production shortfalls and surpluses.

Last year at this conference, Cornelisse and Kuijpers [10] presented a useful paper which examined some of these issues. In particular, they note that analysis of a stocking strategy must take into account runs of several years, and should not examine only the stock level required to meet the possible shortfall in a single year. In other words, the analysis must be dynamic. Also, they point out that stocks are costly; any attempt to discover the government's desirable stock size must consider both the benefits and costs of holding the stock. With these points in mind, Cornelisse and Kuijpers proceed to estimate the best level of wheat stocks for the government to hold at the beginning of a marketing year. They do not examine to what degree stocks should be built up in a good year, nor to what degree stocks should be drawn down in a bad year.

Any attempt to determine the single, optimal level of opening stocks, however, ignores the cost of moving from the government's present position to that optimal level. If these costs are included in the analysis, differences in the state of the world in any one year will lead to different optimal closing stock levels.

*The author is Research Fellow at the International Food Policy Research Institute (IFPRI), Washington, D. C. This research is supported by USAID. The views expressed in the paper are the author's, and not necessarily those of IFPRI or USAID. The helpful comments of Alberto Valdes are gratefully acknowledged. The author alone is responsible for any errors or inaccuracies.

For example, consider two possible cases for Pakistan: in the first case, the government has opening stocks of two million tons, the world price of wheat is \$120 per ton, and production is 14 million tons. The optimal closing stock would be quite high since total supply is large and there is not much benefit to be gained from exporting. In the second case, however, opening stocks are 100,000 tons, the world price is \$200 per ton, and production is only 10 million tons. The optimal closing stock in this case would be considerably less than that in the first case. Clearly, there is no such thing as one optimal stock level.

What the government needs, then, is a strategy for answering the following question: Given the present situation — specifically, the stock on hand, the world price, production, and expectations about next year's values for these variables —, how many tons should be carried? How much should be imported or exported? In other words, the Cornelisse and Kuipers paper tried to discover optimal *opening* stocks for all years; this paper takes the situation in any one year as given, and then examines what *closing* stocks should be.¹ So the result of the analysis is not one optimal stock level, but a table with a list of every possible combination, opening stock, and world price on the left hand side, and the consequent best level of closing stock and trade on the right hand side.

This paper can only take preliminary steps in that direction, partly because the research is just beginning, and partly because no mathematical model can include all the objectives of the government or all the complexities of the wheat market. But there is an important role for modelling in policy design, and it is hoped that insights into the formulation of efficient storage, trade, and price policies will be gained both through this paper and from the ongoing research.

The question will be addressed in the following manner: first, the tools available to the government for responding to fluctuations in supply will be examined. This will be followed by a description and simulation of a commonly prescribed rule for responding to supply fluctuations. Next, the benefits of an optimizing approach will be delineated and the advantages and disadvantages of several optimizing algorithms discussed. Finally, a dynamic programming methodology which explicitly takes into account the competition between government objectives will be applied in a preliminary way to Pakistan.

THE GOVERNMENT'S TOOLS: STOCKS, TRADE, AND PRICE VARIABILITY

Consider a country where the government controls all foreign trade of a particular commodity, and in which the previous year was a normal production year

¹The problem will be simplified by making two assumptions: the expected world price next year is equal to the present world price; and the trend in mean production is known. Also, note that stocks here refer to inter-annual stocks. Seasonal storage is not included.

with the amount produced equal to the amount consumed at a price considered reasonable by the government. Assume for simplicity that there is a constant relationship between producer and consumer prices, so that only one price needs to be studied. There was no change in government stocks during the year, and no foreign trade. In the present year, however, production is one million tons less. What are the government's options?

If substitution between commodities is ignored, there are only three possibilities: (i) decrease stocks; (ii) import the commodity; and (iii) allow consumption to fall. In other words, production variability must be translated into either stock variability, trade variability, or price variability.² For many governments, none of these is desirable. Most governments want to limit consumption/price fluctuations for both welfare and political goals. Although decreasing stocks in any one year may not be a problem, the effect in subsequent years of carrying a low stock may be unsatisfactory. Also, a strategy which requires the government to hold stocks for more than a year or two can be prohibitively expensive. Highly unstable trade can be a problem for countries in which the cereal import bill is a large percentage of export revenues, although this does not appear to be a problem for Pakistan. Also, some countries seem to have an aversion to importing a stable food for reasons of pride.

Nevertheless, a country in this situation has no choice but to rely on these three mechanisms for countering a production shortfall. One type of policy for integrating the three techniques is a price band/buffer stock policy (for brevity, referred to as a "price band policy" below).³ While there are numerous variants of this type of policy, the simplest is as follows: the domestic price is allowed to fluctuate within a band; as long as the price is between the maximum and minimum, the government does not intervene. When the price hits the maximum, the government sells all the grain that is demanded at that price. Sales come from the government stock. When stocks reach a predetermined minimum, grain is imported, whatever the world price. A similar mechanism works for price decreases.

So this type of policy employs all three possible government instruments: consumption and stocks are allowed to fluctuate first, but the extent of the fluctuation in each case is controlled by a band. Trade is used last, but these fluctuations are not limited in size.

As long as the price band is limited to one which is symmetric with respect to a target price, the price band policy can be defined by only three parameters: the target price, the percentage price deviation allowed, and the difference between the maximum and minimum stock levels. This last parameter will be termed "stock variability". The purpose of stocks is to absorb shocks in present production and

²"Consumption" as used here could include changes in private stocks over which the government has no control.

³Some examples of similar but somewhat more complex policies can be found in articles by Reutinger, Bigman, and Abbott [1; 3 4; 5; 6; 7; 22; 23].

move grain from a surplus time-period to a deficit time-period. The degree to which stocks are able to accomplish this task is determined by the stock variability, not by the minimum or maximum stock level alone.

The target price here is taken to be the market clearing price in a year of normal production and no government intervention. This makes sense when such a price is between import and export parity, as it is for wheat in Pakistan at present.

Results of simulating such a policy for two price bands are presented in Table 1. Space does not permit a detailed analysis, but note that the only reason for holding stocks under a price band scheme such as this is to decrease imports.⁴ With price variability held constant, the minimum-cost solution is when no stocks are held. If minimizing of imports is not a separate objective apart from cost, the government would hold no stock.

Trade-offs between government objectives calculated from the price-band simulation will be accurate provided the underlying model is sufficiently realistic and *provided the price band/buffer stock schemes are the most efficient policies for meeting the three government objectives*. This consideration implies that a search of policy space is necessary in order to discover whether the price-band policies are efficient.

In a closed economy, Newbery and Stiglitz [18] have shown that price-band policies are quite inefficient. Since the government tools in this simple policy do not respond to the world price, it would seem that these policies may be even less efficient in an open economy. Thus, the trade-offs are likely to be inaccurate.

There are two alternatives for the analyst who wants to measure trade-offs accurately: (1) modify the price-band policies to make them more flexible to state parameters such as the world price; or (2) set up an optimizing model which chooses the best level of the government control variables for each possible state of the world, and then simulate the optimal policy.

Although optimizing is bound to be more difficult, there are four reasons for using an optimization method rather than simply simulating additional chosen policies.

(1) Since there is a potentially infinite number of possible policies, the choice of which ones to simulate will be to some extent *ad hoc*. There is no way of knowing if the best policy has been simulated.

(2) Optimization ensures that the trade-offs between objectives are measured accurately. When optimizing a multi-dimensional objective function, trade-offs between objectives are measured by varying the weights on the different objectives. If the weight on price stability is increased, for example, optimization ensures that the policy is adjusted in such a way that the cheapest method of increasing price stability is chosen.

⁴ For a more detailed analysis of a similar simulation for Kenya, see Pinckney [19].

Table 1

Simulation of Price-Band Scheme for Pakistani Wheat

('000 tons except where indicated)

Price Band (%)	Standard Deviation of Price (\$/mt)	Stock Variability	Average Annual Fiscal Cost (\$ mil)	Average Closing Stock	Average Annual Imports
2.86	3.4	0	11.1	0	178
2.86	3.4	300	12.7	153	127
2.86	3.4	600	15.2	306	97
2.86	3.4	900	17.9	455	79
2.86	3.4	1200	20.5	599	69
2.86	3.4	1500	22.6	735	64
6.77	7.5	0	5.1	0	127
6.77	7.5	300	7.8	154	87
6.77	7.5	600	11.4	303	63
6.77	7.5	900	14.2	449	51
6.77	7.5	1200	16.7	589	46
6.77	7.5	1500	18.4	711	44

Source: Author's calculations. For assumptions, see Appendix.

(3) Finding the optimal policy allows for measurement of the degree of sub-optimality of alternative policies. Once the optimal policy for a given set of weights is found, it can be simulated in the same way as a price-band scheme. This allows for a direct comparison between different policies. This is especially important for policy recommendations since administrative rules generally are easier to understand and implement than the policies which come out of optimization routines. Yet ease of communication and implementation is an objective of the government which is difficult, if not impossible, to build into the optimizing process. Once the optimal values of the objectives are known, the values of the objectives produced by simpler administrative rules can be compared, and a trade-off between complexity and cost computed.

(4) Differences between the way the optimal policies and the administrative rules respond to the state of the world can be studied, and the administrative rules adjusted in ways suggested by the optimal policies. This should make possible reductions in the complexity/cost trade-off, and the formulation of an administrative rule which is much more efficient than that which would have been tested without considering the optimal policies.

So we have now examined the government tools of stock, trade, and consumption variability, described a typical price-band/buffer-stock policy, and considered the advantages of optimization. Next, the available optimization algorithms will be discussed.

CHOICE OF OPTIMIZATION TECHNIQUE

There are at least four optimization techniques that have been used to analyse storage problems in the literature. These are David Bigman's method, David Eaton's programming method, stochastic control method [16; 17; 21], and dynamic programming.⁵ They will be considered in order.

Bigman [4] has developed a technique which finds the optimal price-band/buffer-stock policy. If the government has decided already that a price-band policy is to be implemented, this is a useful technique. Otherwise, however, there could be a more efficient policy. Thus, in the absence of government direction in the matter, this technique cannot yield the four benefits outlined above.

Eaton's method [11] requires that the problem be formulated with a quadratic or linear objective-function and linear constraints. Production must be characterized as a random variable which is not influenced by the endogenous government decision variables. Thirty or more sequences of, say, ten years each of possible future production are then produced by a random number generator. A series of multi-period quadratic programming problems is then formulated, one for each ten-year production sequence, with the level of production for each year entering as a constraint. The quadratic programming results give optimal levels of the control variables for each particular future production sequence. The control variables clearly will be different for each production sequence and are chosen by the model with full knowledge of future production. Eaton acknowledges that this is unrealistic since the model is making use of facts which would be unknown at the time decisions are made. However, he suggests that more general optimal policies can be gleaned from these particular implementations by regressing the levels of the control variables on the values of other control and state variables that are known at the time decisions would have to be made.

The problem with this technique is that the best specification of the regression equations is unknown. Given the results of the quadratic programming problems, it is difficult to move ahead. So Eaton's method is not a true optimization method since there is no way to guarantee that the optimal specification for policy structure has been tested in the regression analysis. Thus, it cannot yield the four benefits listed above.

⁵Bellman [2] is the classic presentation and still one of the best. In the area of optimal stockholding, Gustafson [13] was the pioneer. Gardner [12] and Burt, Koo, and Dudley [8] used dynamic programming in more complex models. Gardner's book and Plato and Gordon [20] are particularly useful for the practitioner.

Control methods leave no question about optimality if prior conditions are satisfied. Stochastic control solutions of real world problems, however, are in general notoriously difficult to compute.

Theil [24] has shown that stochastic control problems can be simplified under certain conditions. A stochastic control problem with a quadratic objective function and linear state transition equations has a certainty-equivalent solution; that is, the variance of the stochastic variable makes no difference in the choice of the optimal policy. If the problem can be stated in those terms, the solution can be found simply by filling in the values of some variables.

These conditions — a quadratic objective-function and linear state transition equations — can be satisfied for Pakistan if the demand function is assumed to be linear. But consider the following case. There are two countries, each with an average production of 10 million tons. In one, the standard deviation of production is 5 million tons, while in the other the standard deviation is only half a million tons. If certainty equivalence is correct, provided both countries have the same opening stock and the same production, and face the same world price, then they should carry the same stock despite the large difference in the standard deviations. This is both counter-intuitive and incorrect. The problem is that the derivation of the certainty-equivalence theorem assumes that there are no inequality constraints. But stocks, imports, and exports are all strictly non-negative variables. (Because of the difference between import and export parity, imports and exports must be costed differently and thus must remain distinct rather than being added together as trade.) Such non-negativity conditions make the certainty-equivalent formulation unworkable.⁶ Thus, certainty-equivalent methods are not appropriate for the analysis at hand.⁷

The final optimization technique under consideration is dynamic programming. Formally, it is incorrect to make a distinction between dynamic programming and control methods since they can be shown to be equivalent [15]. But, in general usage, the term "control methods" refers to the use of calculus techniques to solve a continuous optimization problem, while "dynamic programming" refers to the use of numerical methods to solve a discrete form of the problem.⁸ The terms will be used here in this way.

⁶Despite these problems, some researchers have applied certainty equivalent methods to the analysis of storage decisions. They invariably recommend stock levels which are sub-optimal. See Gardner [12, pp. 16-17] and Pinckney [19, pp. 27-29].

⁷Control methods are most useful and interesting when the systematic variations from period to period are significant, making the non-stochastic case non-trivial. In this case, certainty-equivalent methods can be used to approximate a solution, and the stochastic case solved by minimizing deviations from the certainty-equivalent case. This is the technique used by Kim, Goreux, and Kendrick [17] and Rausser and Hochman [21].

⁸Calculus techniques often are used in the solution algorithms of discrete dynamic-programming problems when choosing which of the finite number of control applications to consider next.

The backward recursion solution algorithm for dynamic programming will not be described here. It is best introduced through a simple example, such as that presented in Pinckney [19, pp. 30–32] or Gardner [12, pp. 5–14]. It is clear from these examples that dynamic programming can solve dynamic, stochastic problems similar to Pakistan's wheat supply stabilization problem. Dynamic programming has the additional advantage of allowing inequality constraints of any type on the variables.⁹ Also, unlike Eaton's method or certainty-equivalent techniques, a non-linear demand curve can be used.

Two charges have been brought against dynamic programming. These are the need for a simplistic objective function, and the "curse of dimensionality".

Simplistic objective functions have often been used, but this has been by choice; it is not required by the methodology. The solution algorithm used in the present study is simply a brute force search; this is slow, but it works and is effective in dealing with three objectives and rather complex costings. Importantly, all the programmes can run on IBM PCs and compatibles so that the costs are not excessive. Consequently, there is no need to avoid dynamic programming because of the need for a complex objective function.

The second problem with dynamic programming is the so-called "curse of dimensionality". The problem multiplies in difficulty with the number of states of the world. As an example of the effects of the curse, consider the model used in the analysis below. The state of the world is determined by three variables: world price of wheat, wheat production, and opening wheat-stocks. If the model is made more complex and realistic by adding world price, production, and opening stocks of rice with, say, 7, 9, and 10 discrete levels respectively, the problem would take 630 times as long to solve.¹⁰

The "curse" is not a major problem for a preliminary study such as the one presented here. It does limit the complexity of the model that can be solved with standard dynamic-programming. With complex models, either different solution techniques would have to be used or dynamic programming used to optimize a sub-section of the model and simulation used to test for the optimal values of the complete model.

Thus, dynamic programming is clearly the best choice of optimization technique for the problem at hand. All the four benefits of optimization mentioned

⁹In fact, a limitation of this approach is that as a discrete system, inequality constraints are required on every variable. Thus it is difficult at times to ensure that the results are not constrained by inequalities which are a function of the particular specification of the solution algorithm rather than of the theoretical model.

¹⁰Burt, Koo and Dudley [8] use a solution technique which is not a backward recursion. Their method has the advantage of decreasing the number of calculations and allowing other variables from a more complex econometric model to influence the expected value of future costs. The disadvantage is that the expected values are no longer exact, as they are when a backward-recursion algorithm is used.

above — selection across many types of policies, measuring trade-offs accurately, maximizing the value of the objective function for comparison with sub-optimal policies, and providing suggestions for modifications of administrative rules — are gained through the use of dynamic programming.

In order to take advantage of the last three benefits, however, it is necessary to take the analysis one step beyond the optimization algorithm itself. The optimal policies must be simulated in the same framework as the price-band schemes in order to compare measurements of trade-offs, to judge the sub-optimality of the price-band schemes, and to learn from the optimal policies in order to adjust the price-band schemes. Before the optimal policies can be simulated, however, they need to be made "continuous" rather than discrete. This is accomplished by linearly interpolating the discrete optimal policies across state variables. While it is unlikely that the resulting policy will be the optimal "continuous" policy, it will not be substantially different if the discrete optimization includes a large enough number of levels of each state variable.

At this point, the policy problem and one proposed solution, the price-band/buffer-stock policy, have been described. Also, an optimization technique has been chosen which will shed light on the degree of sub-optimality of the proposed solution, and provide insights into the nature of an optimal policy. Results of the optimization process are presented in the next section.

DYNAMIC-PROGRAMMING RESULTS

Two different types of objective functions are used here in order to highlight the differences between the normal dynamic-programming results and the ones which take explicit account of competing government objectives. The first objective function is the standard for dynamic-programming exercises: maximize present plus discounted expected future consumer/producer surplus less international trade losses and storage costs. It has been shown elsewhere that such an objective function produces the free-market solution [13, pp. 48–49]. The second type of objective function minimizes present plus discounted, expected future value of a cost function consisting of the sum of three different components. The components are:

- (i) minimize price variability, measured here as the sum of squared deviations from the normal production-year equilibrium price;
- (ii) minimize imports;
- (iii) minimize fiscal cost, including storage and net foreign and domestic trading losses.

The equations used for each component and key assumptions made in the model are listed in the Appendix.

The first two components are weighted differently in different runs of the model in order to see how the optimizing process adjusts when one objective is relatively more important than another. In some runs the weight on imports is zero in order to test the case where the only negative aspect of imports is the cost. As mentioned above, once the discrete optimal policy is discovered for each set of weights by the dynamic-programming algorithm, it is interpolated and simulated in the same way as the price-band schemes were simulated above.

The free-market solution which results from maximizing consumer/producer surplus basically relies on trade to stabilize prices. The average stock level is 12,000 tons. Stocks are only held when world prices are very low and total supply is very high. Consequently, price variability is quite high: the standard deviation of price is \$16.7 per ton around a mean of \$125. Imports average 108,000 tons annually, and profits average 7.5 million dollars annually. If there were a competitive private sector with freedom to buy, sell, import, and export wheat, it would produce this result on its own.¹¹ Consequently there would be no fiscal costs involved.

The second type of objective function allows for the possibility of government preferences for lower price-variability and lower import-levels. Table 2 presents results for six different combinations of objective-function weights which have been chosen so that they produce only two levels of price variability.

Table 2

Simulation of Optimal Policies for Pakistani Wheat

Price Weight	Weight on Imports	Standard Deviation of Price (\$/mt)	Average Annual Fiscal Cost (\$ mil)	Average Closing Stock ('000 t)	Average Annual Imports ('000 t)
282*	0	3.40	9.7	97	169
358*	60	3.40	10.8	208	120
420	120	3.40	13.1	316	92
102	0	7.50	1.9	70	135
120	40	7.50	2.4	119	99
140	80	7.50	3.7	180	76

Source: Author's calculations. For assumptions, see Appendix.

*Interpolated from two different results in order to produce a price variability of 3.40.

¹¹This assumes that storage costs, costs of importing and exporting, and the discount rate are the same for the private sector and the government.

The weights for the objective function were chosen so that the price-band simulations produce price-variability levels equal to the dynamic-programming results. This allows for direct comparison of the optimal and price-band policies.

The Fig. 1 presents these results graphically. The horizontal axis is fiscal cost, while the vertical axis is average annual imports. The four lines in the figure represent possible combinations of cost and imports with price variability held constant. The two solid lines graph the price-band policies presented in Table 1. The dotted lines represent the optimal policies from Table 2. The price variability is typed at the bottom of each line.

The tables and graph clearly show that the optimal policies are superior to the price-band policies. However, the difference is much greater when price variability is 7.5 than when it is 3.4. Thus, the complexity/cost trade-off declines with decreasing price variability. Similar results were found for Kenya [19].

Also, note how each line becomes less steep as it moves down. This indicates that decreasing of imports from the highest levels is relatively cheap, but becomes more and more expensive as imports decline. For the optimal policies at a price variability of 7.5, the average annual cost is about 23,000 dollars per thousand tons of decrease for the first 49,000-tons decline, but rises to over 80,000 dollars per thousand tons for the next 28,000-tons decline.

Of greatest importance, however, is the large decrease in cost in moving from price variability of 3.4 to 7.5. Each dollar of price variability costs 2.1 million dollars annually when imports are 135 million tons; the figure rises to 2.4 million dollars when imports are 100 million tons.¹² If minimizing imports is not a separate objective, the trade-off is 1.9 million dollars annually per dollar of price variability.

The losses in the multi-objective policy can also be compared to the profits of the free-market solution. The difference represents the efficiency loss suffered in order to decrease price variability. If import levels are ignored, the cost to the economy of decreasing price variability from the free-market solution of 16.7 dollars per ton to 7.5 is 9.5 million dollars annually, or somewhat higher than 1 million dollars annually per dollar decrease in price variability.¹³

Finally, note in Table 2 that average optimal stock levels are quite low by Pakistan's historical standards. With price variability of 7.5 and no added weight on imports, average stocks are only 70,000 tons. Even when imports are penalized an additional \$120 per ton and the price variability is low at 3.4, average optimal stock levels are only 316,000 tons. It should be recalled, however, that there is no such thing as the optimal stock level. The dynamic-programming algorithm produces a different optimal closing stock for each combination of world price, opening stock, and production. The figure shown here is an average from the simulation of the

¹²The optimal results are interpolated linearly in order to hold imports constant.

¹³The model does not include any supply response to lowered price variability. Bigman [4] argues that this response is important.

Fig. 1. Trade-off Curves for Pakistan Wheat

optimal policy. For low world-prices, high opening-stocks, and high production, optimal stocks are considerably higher than the average (although never higher than 1.1 million tons). For high world-prices, low opening-stocks, and low production, optimal carryovers are often zero. Indeed, when the world price is 170 or higher, the optimal carryover is zero regardless of production and opening stock for the objective-function weights considered here.

In addition, it should be noted that stock levels in the optimization do not include government seasonal storage. The stocks modelled here are inter-annual stocks, that is those on hand in April just prior to the harvest.

Considerably more could be done with the results at hand. The optimal policy could be studied in order to see why it is more efficient. These lessons could then be applied in the selection of a relatively simple policy that would approach the optimal policy in efficiency. This has been done for Kenya [19], but is beyond the scope of the present paper.

In addition to extending the results, the model itself could be made more realistic and interesting by: (i) including demand fluctuations, (ii) breaking down the country into regions, (iii) dividing production into subsistence and marketed surplus, (iv) including a private storage equation, (v) including other crops, (vi) including more than one season, (vii) including a supply response to the level and variability of price, (viii) allowing supply and demand to diverge during each ten-year cycle, and (ix) allowing the relationship between producer and consumer prices to vary.

CONCLUSIONS

This paper has demonstrated the possibilities of using an optimizing algorithm for an analysis of Pakistan's trade, storage, and price policies. Dynamic programming has been shown to be useful for developing a strategy for contending with fluctuations in wheat production. These preliminary results suggest the following five points:

- (1) The cost to the government of holding prices constant or within a narrow band is quite high; much money can be saved by allowing fairly small increases in price variability;
- (2) Optimal policies are considerably more efficient than the price-band policies modelled here, although the relative difference decreases with lower price-variability;
- (3) Imports can be lowered considerably at small cost, but the cost rises the more the imports are decreased;
- (4) Average optimal stock levels are low, ranging from 70,000 to 316,000 tons for the objective-function weights considered here.
- (5) The free-market solution is characterized by relatively high price variability, but low costs.

It should be borne in mind that all of the numerical results should be considered indicative rather than prescriptive. There are numerous limitations of the simple model analysed here. Nevertheless, these five general conclusions appear to be robust and are expected to be confirmed by ongoing research.

Appendix

ASSUMPTIONS FOR MODELS

For both the simulation and the optimization programmes, the following assumptions are used. These numbers are only illustrative.

Normal-year equilibrium price: \$125 per metric ton

Own price demand elasticity for wheat: -0.3

Discount rate: 7 percent

Normal-year wheat production less seed, feed, and waste: 12.5 million mt.

Standard deviation of wheat production: 5 percent of the mean.

Shadow foreign exchange rate = nominal foreign-exchange rate.

Food-aid provided to make up 50 percent of the difference between 12 million metric ton and (production plus opening stocks).

Foreign and domestic transport costs totalling \$40 per ton must be paid on food aid.

Export parity is the world price minus \$50 per metric ton.

Import parity is the world price plus \$40 per metric ton.

The base world-price of wheat is \$130 per metric ton.

The world price of wheat moves in a random walk with a standard deviation of \$20 per metric ton.

The world price of wheat does not fall below \$30/mt.

Storage costs are \$25 per mt annually.

For the dynamic-programming algorithm, there are only seven possible world prices for wheat: 70, 90, 110, 130, 150, 170, and 190 dollars per metric ton. Consequently, the price cannot behave exactly as a random walk. As long as the price is not at one of the two extremes, however, the difference between the present price and the expected price is not great.

The components of the weighted objective function are as follows:

Price variability: $a \cdot (P_t - 125)^2$

Imports: $b \cdot M_t$

Fiscal cost: $P_t \cdot NP_t + M_t \cdot MPAR_t + X_t \cdot XPAR_t + S_t \cdot 25 + AID_t \cdot 40$

where P_t is price, a and b are weights, M_t is imports excluding food aid, NP_t is net government purchases, $MPAR_t$ is import parity, X_t is exports, $XPAR_t$ is export parity, S_t is closing stocks, and AID_t is food aid.

By considering only one price, the assumption is that the relationship between producer and consumer prices (i) is constant, (ii) covers the cost of marketing, and (iii) the government and the private sector are equally efficient at marketing.

The computer programmes used for this paper are available from the author on IBM-PC disks. They are written in Microsoft Fortran 3.31.

REFERENCES

1. Abbott, Philip C. "A Methodology for Evaluating National Grain Stockpiling Strategies". Rome. FAO. June 1985.
2. Bellman, *Dynamic Programming*. Princeton: Princeton University Press. 1957.
3. Bigman, David. *Coping With Hunger: Toward a System of Food Security and Price Stabilization*. Cambridge, Massachusetts: Ballinger Publishing Company. 1982.
4. Bigman, David. *Food Policies and Food Security Under Instability: Modeling and Analysis*. Lexington, Massachusetts: Lexington Books. 1985.
5. Bigman, David. "Supply Response and Efficiency Gains from Storage Operations: A Simulation Analysis". *European Review of Agricultural Economics*. Vol. 9. 1982. pp. 205-220.
6. Bigman, David, and Shlomo Reutlinger. "Food Price and Supply Stabilization: National Buffer Stock and Trade Policies". *American Journal of Agricultural Economics*. Vol. 61. November 1979. pp. 657-667.
7. Bigman, David, and Shlomo Reutlinger. "National and International Policies Toward Food Security and Price Stability". *American Economic Review*. Vol. 69. May 1979. pp. 159-163.
8. Burt, Oscar R., Won W. Koo, and Norman J. Dudley. "Optimal Stochastic Control of U. S. Wheat Stocks and Exports". *American Journal of Agricultural Economics*. Vol. 62. May 1980. pp. 172-187.
9. Cochrane, Willard W. "Some Non-Conformist Thoughts on Welfare Economics and Commodity Stabilization Policy". *American Journal of Agricultural Economics*. Vol. 62. August 1980. pp. 508-511.
10. Cornelisse, Peter A., and Bart Kuijpers. "On the Optimal Size of a Buffer Stock - The Case of Wheat in Pakistan". *Pakistan Development Review*. Vol. XXIV. Autumn-Winter 1985. pp. 335-348.
11. Eaton, David J. *A Systems Analysis of Grain Reserves*. Technical Bulletin No. 1611. Washington, D. C.: Economics, Statistics, and Cooperatives Service, United States Department of Agriculture. January 1980.
12. Gardner, Bruce L. *Optimal Stockpiling of Grain*. Lexington, Massachusetts: Lexington Books. 1979.
13. Gustafson, Robert L. *Carryover Levels for Grains: A Method for Determining Amounts that are Optimal Under Specified Conditions*. Technical Bulletin No. 1178. Washington, D. C.: United States Department of Agriculture. October 1958.
14. Hazell, Peter B. R. "An Overview of Workshop Issues". Presented at IFPRI/DSE Workshop on Sources of Increased Variability in Cereal Yields, Feldafing, Germany. November 1985.

15. Intriligator, Michael D. *Mathematical Optimization and Economic Theory*. Englewood Cliffs, N. J.: Prentice-Hall. 1971.
16. Kendrick, David A. *Stochastic Control for Economic Models*. New York: McGraw-Hill. 1981.
17. Kim, H. K., L. M. Goreux, and D. A. Kendrick. "Feedback Control Rule for Cocoa Market Stabilization". In W. C. Labys (ed.), *Quantitative Models of Commodity Markets*. Cambridge, Massachusetts: Ballinger Publishing Company. 1975. pp. 233-263.
18. Newbery, David M. G., and Joseph E. Stiglitz. *The Theory of Commodity Price Stabilization: A Study in the Economics of Risk*. Oxford: Clarendon Press. 1981.
19. Pinckney, Thomas C. "Production Instability and Food Security in Kenya: Measuring Trade-offs Between Government Objectives". Dissertation. Food Research Institute, Stanford University. March 1986.
20. Plato, Gerald, and Douglas Gordon. "Dynamic Programming and the Economics of Optimal Grain Storage". *Agricultural Economics Research*. Vol. 35. January 1983. pp. 10-22.
21. Rausser, Gordon C., and Eithan Hochman. *Dynamic Agricultural Systems: Economic Prediction and Control*. New York: North-Holland Publishing Company. 1979.
22. Reutlinger, Shlomo. "Policies for Food Security in Food-Importing Countries". In A. H. Chisholm and R. Tyers (eds.), *Food Security: Theory, Policy, and Perspectives from Asia and the Pacific Rim*. Lexington, Mass.: Lexington Books. 1983. pp. 21-44.
23. Reutlinger, Shlomo, and David Bigman. "Feasibility, Effectiveness, and Costs of Food Security Alternatives in Developing Countries". In Alberto Valdes (ed.), *Food Security for Developing Countries*. Boulder, Colorado: Westview Press. 1981. pp. 185-212.
24. Theil, Henri. "A Note on Certainty Equivalence in Dynamic Planning". *Econometrica*. Vol. 25. April 1957. pp. 346-349.

Comments on "Stabilizing Pakistan's Supply of Wheat: Issues in the Optimization of Storage and Trade Policies"

There is a perpetual problem, in the underdeveloped countries of relatively large shortfalls and surpluses in various commodities. This is particularly true in the agricultural sector where production is extremely uncertain. It is important to have a policy of maintaining stocks to buffer these variations, particularly for a major crop. In Pakistan the major crop is wheat. As such, the paper under discussion addresses a topic of the utmost importance to Pakistan's economy. The question is: "How much stock of wheat should be maintained in Pakistan?" The answer to it will depend on the situation. Thus, what is required is a prescription for determining the stock in any given situation. Of course, there can not be a completely general prescription, but a slightly more limited prescription could be looked for.

The paper under review was supposed to supersede the earlier work of Cornelisse and Kuijpers*. The objection made to that work appears to be that it was static while a dynamic approach was required. So far so good. The work of Cornelisse and Kuijpers could naturally have been extended to a comparative static analysis, but a dynamic analysis is surely better. My problem with this paper is that it is not clear that the dynamic analysis taken is valid.

An *ad hoc* algorithm for a numerical scheme for optimization is chosen. No economic justification or analysis is presented for it. The objective function is not given, nor is it explained what the economic objective is. The mathematical analysis involved and the computational scheme used are suppressed. As such we are asked to take on faith the numbers provided as giving the optimal stock of wheat at the end of the planning period.

Another problem is that the dynamic analysis is not apparent. Merely putting a suffix "t" on a symbol representing a quantity does not make the analysis dynamic. In fact, this is what could be used to develop a comparative static analysis. For a dynamic analysis an economically sound objective function must be set up and the Pontryogin maximum principle applied. The Hamiltonian constructed will then have to satisfy Hamilton's equations. It is quite conceivable that the equations would need to be solved numerically. What has been called, in this paper, "dynamic

*Peter A. Cornelisse and Bart Kuijpers. "On the Optimal Size of Buffer Stock - The Case of Wheat in Pakistan". *Pakistan Development Review*. Vol. XXIV, Nos. 3&4. Autumn-Winter 1985.

programming” is just the algorithm normally used to solve those equations. In standard optimization theory the entire procedure is called dynamic programming.

A further problem is the assumption made that a stochastic process will correspond to a deterministic process. It would do so, *over sufficiently long periods*, provided the objective function is linear or quadratic. However, as pointed out just now, the nature of the objective function has *not* been derived but simply assumed. There is no reason, then, to expect the behaviour of the meaningful objective-function to have any relation to the function which is minimized. Even if it did have such a relation, we may not have sufficiently long times. The averaged-out stock may be positive but the stochastic variable may have had to attain negative values in the mean time. In other words, the economy may have had to collapse long before the averaging effect could work. It is obvious that the problem of uncertainties (which brings in the stochastic variables) can not be simply “wished away” by assuming a quadratic or linear objective function.

Another point which worries me is that the numbers seem to have been chosen arbitrarily for illustrative purposes, but are not claimed to be realistic numbers. In any non-linear system, the nature of the solution can change drastically with a change of initial values by even a small amount. The stability of the solution is not guaranteed. Nowhere has a test of the stability of the solution been mentioned. From the way it is presented, it could have been any country instead of “Pakistan” and any stockable commodity instead of “wheat”.

To conclude, I am sure that the author has available with him all the details of mathematical analysis, economic reasoning, computational methods, etc., which I have been complaining about. I just wish that he had presented them in his paper so that we could all follow his reasoning.