OPTIMAL COMMODITY STOCK-PILING RULES

By DAVID M. G. NEWBERY *and* JOSEPH E. STIGLITZ

1. Introduction

IN May, 1976, a special meeting of UNCTAD passed a resolution calling for the establishment of an Integrated Programme for Commodities, which was intended, amongst other objectives, to stabilize the prices of the ten core commodities[1] identified by UNCTAD as suitable for stockpiling. Four years later, the Brandt commission endorsed these objectives, and called for adequate resources to enable the Common Fund 'to encourage and finance effective International Commodity Agreements (ICAs) which would stabilize prices at remunerative levels; to finance national stockpiling outside ICAs; and to facilitate the carrying out of Second Window activities such as storage, processing, marketing, productivity improvement and diversification.' (Brandt, 1980, pp. 158–9).

In response to this political call for commodity price stabilization, economists have not been slow to provide theoretical and empirical analyses of the desirability of such price stabilization. (See in particular, Newbery and Stiglitz, 1981a, and its references, or the extensive references in Wright, 1979). With a few notable exceptions the resulting literature is theoretically unsatisfactory as it ignores the obvious organising principle of competitive economic theory.

There are two strategies that economists can follow in addressing this appeal for commodity price stabilization. The broad primrose path can be caricatured as follows. First, is is true that it is desirable to stabilize prices in a market subject to random shocks? The question requires a criterion of desirability, and the usual choice is the maximization of average Marshallian surplus. Given this criterion, the answer is that, in specific models, perfect price stabilization would be desirable if it were feasible and costless to achieve. Many writers stop there, but some press on to ask how the gains of stabilization would be distributed between producers and consumers under various model specifications. (See e.g., the survey by Turnovsky, 1978). A very few observe that perfect price stability of the kind sought is infeasible, and/or, infinitely costly. Once it is recognised that storage is costly (if only because interest rates are positive) the question is naturally recast as one of finding the optimum degree of price stability. Of course, those involved in formulating Commodity Agreements recognise this problem at a rather earlier stage, and typically adopt a rather pragmatic approach. Since the objective is to stabilize prices, then what is needed is a specification of the limits within which prices are to be stabilized, and a decision as to the desired level of stocks to hold and of finance required for further purchases.

---

[1] The ten core commodities are sugar, coffee, cocoa, tea, cotton, jute, sisal, rubber, copper and tin. See UNCTAD documents, series TD/B/C.1.

The result is a *band width rule*, which the Agency will defend by selling commodities from stock when the price moves above the upper intervention price, and buying them for storage when the price falls below the lower intervention price. Since this is the way exchange rates are pegged, it is a natural rule to choose when the objective is to stabilize prices. The specification of the band width and the required stocks can be investigated analytically (Edwards and Hallwood, 1981) or by stochastic simulation (Behrman, 1978).

The other approach is to ask what economic theory has to say. Commodity markets are usually held up as paradigms of competitive markets. We know from competitive theory that, in the absence of externalities, if the market structure is complete and competitive, then the equilibrium is Pareto efficient. What case can be made for market intervention in the form of some Commodity Agreement? One obvious objection to the market equilibrium is that it is inequitable, and this theme carries much force in the Brandt report. If so, then the same competitive theory argues that lump sum transfers (aid, in short) are preferable to price distortions. This is a convenient way of side-stepping equity considerations which we shall provisionally accept for the moment, whilst recognising that the alternatives may not be politically feasible. We discuss what to do in such cases below in Section 4.

The second objection to the market equilibrium is that even if commodity markets are competitive (as seems to be the case), they are not complete. Although primary commodities provide most of the small number of examples of futures markets, these markets extend less than two years ahead, whilst the required insurance markets are, for the most part, simply absent. These objections carry considerable force, though remarkably few economists have addressed the question of what market incompleteness implies for policy intervention. Newbery and Stiglitz (1982) show that in general a competitive market equilibrium is not even constrained Pareto efficient; that is, the Government could make everyone better off by setting constant ad valorem taxes (or subsidies) plus lump sum subsidies (or taxes) even if it were not able to establish currently absent markets. However, they also show that if agents are risk-neutral and hold rational expectations, then the competitive equilibrium will be efficient. The reason is obvious. If agents are risk-neutral, and hold common (objective) beliefs about the economy, then they would not wish to trade on risk markets even if they existed. Similarly, if they hold common (objective) beliefs about future prices, they would not wish to trade on future markets, and in both cases such markets would be redundant.

Now, it is interesting to note that almost without exception, the models built to analyse the benefits of price stabilization assume risk neutrality, for otherwise their criterion of maximizing expected consumer plus producer surplus is inappropriate.[2] The economic theorist can thus immediately

---

[2] One notable exception is the paper by Turnovsky, Shalit and Schmitz (1980), which follows a similar approach to Newbery and Stiglitz (1979).

deduce what the optimum price stabilization rule is in these models—it is to reproduce the storage decisions which risk neutral competitive speculators would make if their price forecasts were rational (unbiased). It is also interesting to note that this result, now over twenty years old, was established, not by appeal to the theorems of competitive equilibrium, but by interpreting the solution to the underlying stochastic dynamic programming problem. Gustafson's (1958) seminal monograph derived the solution by adopting existing methods of optimal inventory analysis, but he was careful to point out that optimal amounts to be stored 'are exactly the same as the amounts that would be stored in the aggregate by private firms in a so called "idealized" free market'. (Gustafson, 1958, p. 32). This is later explicitly described as one in which 'firms seek to maximize discounted expected profit' (ibid, p. 49) i.e. are risk neutral.

In our own work on commodity price stabilization (Newbery and Stiglitz, 1977, 1981) we have proceeded as follows. We first establish the conditions under which a competitive market economy is constrained Pareto efficient in the absence of a complete set of risk and futures markets. If these conditions are satisfied, then the optimum stockpiling rule is just the competitive speculative storage equilibrium. The next step is to identify the biases present in the competitive equilibrium when these conditions are not satisfied, and either find the policy interventions required to restore efficiency, or find the desired degree of price stabilization. We would argue that there are several decisive advantages of this approach over the alternative of investigating particular price stabilization rules in specific, usually very special, models.[3]

The first advantage is that competitive theory is firmly based on individual utility, and hence does not prejudge the sense in which price stability is desirable. The second is that is immediately draws attention to potential market failures, and particularly the absence of risk and futures markets, which are typically ignored in simple modelling approaches. It thus raises questions which can then be addressed directly. Finally, it raises the question of decentralising the efficient allocation. In the present context this means asking whether the proposed stabilization should be undertaken by the market or by an agency, and, in the latter case, what the impact of the actions would be on private speculative storage, and whether indeed the actions of private speculators will jeopardize the desired outcome.

The present paper has two main objectives. The first is to characterise the competitive stockpiling rule, which is the optimum rule if agents are risk neutral. Once this has been done, it becomes possible to explore the impact of stabilization on the market, to measure the degree of stabilization achieved, and the benefits which accrue. It is also possible to compare the efficient rule with alternatives such as the band width rule, and measure the losses associated with following such inefficient rules. Our main contribution

[3] Some of the senses in which these models are special and misleading are listed in Newbery and Stiglitz (1979).

here is to identify a special case in which the optimum rule can be solved analytically, and to show how this solution is modified in more general cases.

The second objective is to demonstrate a method for measuring the bias in the competitive rule when producers are risk averse. We also suggest policy interventions which can then improve allocative efficiency, and show how to measure their costs and benefits.

2. Competitive stockpiling rules

The basic result which serves as a benchmark for our analysis is the following proposition.

Proposition

If agents are risk neutral and hold common objective beliefs about the price distribution over states of the world, and if they have access to perfect capital markets, then the rational expectations competitive equilibrium is Pareto efficient.

Proof

Consider the artificial economy identical to the one under consideration except that it possesses a full set of Arrow-Debreu contingent and futures markets. Since agents are risk-neutral and hold common beliefs, there are no gains from trade on any risk market, and such markets are thus redundant. Likewise, given a perfect capital market which permits inter-temporal wealth transfer, futures markets in goods are redundant. The artificial economy thus achieves the same allocation of goods as the reference economy since trade is confined to the same set of markets. Since the artificial economy is Pareto efficient, so is the reference economy.

Remarks

In general, markets serve two functions—to provide price information and to permit trade. Given shared objective information (the implicit assumption behind rational expectations) this first role is redundant. Risk-neutrality ensures risk market redundancy, so together the assumptions allow one to dispense with risk markets. The assumption of shared common beliefs in the absence of such markets is clearly a strong one. The assumption of risk neutrality is also strong, since it requires price risk neutrality as well as the more familiar income risk neutrality. It can be defended in partial equilibrium analysis if consumers spend a small proportion of their income on the (single) risky commodity, but in a General Equilibrium model is a very stringent condition (see Newbery and Stiglitz, 1982; Stiglitz, 1969).

We shall first examine the implications of the competitive stocking rule in a general model and then derive an analytical solution for a special simple case. Throughout we consider only supply disturbances, largely because the

welfare analysis of demand disturbances depends sensitively on the source of the disturbance and would require a lengthy examination of alternative cases. (For further discussion, see Newbery and Stiglitz, 1981a, Ch. 8.) The notation and structure of the model are as follows:

(i) There is a stock S_{t-1} carried forward from the previous year.

(ii) To this is added a random harvest, $\tilde{h}_t$, so that at date t the amount available for consumption, C_t, and for storage, S_t, is the total supply, x_t:

$$x_t = h_t + S_{t-1} = C_t + S_t. \tag{1}$$

We assume that there are no losses in storage (though these are easy to handle—see Samuelson, 1971), and that weather and other random factors are serially uncorrelated. In the simple case planned production does not vary from year to year, in which case harvests will be serially uncorrelated. This is a reasonable assumption for tree crops like cocoa and coffee where production decisions cannot be much altered during the life of the tree. For annual crops like wheat the assumption is less reasonable, for high current storage, resulting from a high current harvest, will depress expected future prices, and induce a supply reduction. In such cases harvests would be negatively serially correlated. In principle it is straightforward to model this supply response, which will require specifying planned production as a function of current expectations. The usual formulation makes planned production a function of the action certainty equivalent price, which depends on the expected future price. In the special linear case discussed in Appendix 2 the certainty equivalent price is equal to the expected price, but in general they will differ. (See e.g., Newbery and Stiglitz, 1981a, Ch. 5–6).

Finally, we assume that demand is stationary and non-stochastic, so that the market clearing price, p_t, depends only on current consumption, C_t:

$$p_t = p(C_t). \tag{2}$$

The competitive storage rule can be found by maximising expected social welfare (the approach taken by Gustafson, 1958, and Samuelson, 1971) or derived directly from the competitive arbitrage conditions. It is easy to demonstrate the equivalence of these two approaches (see Newbery and Stiglitz, 1981a) so we shall follow the second, more transparent approach.

If the annual storage costs excluding interest is k per unit at the margin, then a speculator who buys after the harvest at price p_t and sells after the next harvest at price p_{t+1} will have made a marginal profit (in money terms at date $t+1$) of

$$p_{t+1} - (p_t + k)(1 + r)$$

per unit, where r is the rate of interest. If speculators are risk neutral then they will store nothing if expected profit is negative (i.e., the current price is too high) but will otherwise continue to store until they have driven expected marginal profit down to zero. These two cases can be combined in

the fundamental arbitrage equation

$$\left.\begin{array}{r} p_t + k \geqslant \beta E p_{t+1} \\ S_t \geqslant 0 \end{array}\right\} \quad \begin{array}{l} \text{complimentarily (i.e., if one equation has a strict} \\ \text{inequality, the other must have an equality.)} \end{array} \quad (3)$$

where $\beta = 1/(1+r)$ is the discount factor. Our objective is to find a storage rule which satisfies this fundamental arbitrage equation. Since price demands only on consumption from (2), and since from (1)

$$C_t = x_t - S_t \qquad (4)$$

it follows that we are looking for a stock rule which is a function only of total supply, x_t:

$$S_t = f(x_t) \geqslant 0. \qquad (5)$$

This ensures that consumption, and hence price, will only depend on supply, x. Contrast this with the band width rule, where the stabilization authority defends upper and lower intervention prices. In this case the carry over depends typically on both initial stocks and the current harvest, as do consumption and prices. Clearly, the band width rule does not arbitrage prices in the sense of equation (3), and hence is not an efficient method of stabilizing consumption, which, given our assumptions about risk neutrality, is the basic reason for stabilizing price. Later we shall see just how inefficient this rule can be.

To return to the problem of finding the competitive stock rule, we seek a function $f(x)$ which solves equation (3), which, given (2), (4) and (5), can be written

$$\left.\begin{array}{r} p\{x_t - f(x_t)\} + k \geqslant \beta E p[\tilde{h}_{t+1} + f(x_t) - f\{\tilde{h}_{t+1} + f(x_t)\}] \\ f(x_t) \geqslant 0 \end{array}\right\} \quad \text{complementarily} \quad (6)$$

If planned production depends on the expected price then, with additive risk

$$\tilde{h}_{t+1} = q_t + \tilde{u}_t, \qquad E\tilde{u}_t = 0$$
$$q_t = q(E p_{t+1})$$

which would initially appear to further complicate the functional relationships. In Appendix 2 this complication is shown to be more apparent than real, because the arbitrage equation (3) allows this production dependency to be simplified. The real problem in solving equation (6) arises in the term $f\{h + f(x)\}$ and most of the ingenuity in finding solution algorithms lies in dealing with this term. One key feature of the rule is, however, immediate.

(i) The optimum storage rule is non-linear. This follows because stocks must be non-negative, so $f(x) = 0$ below some critical value of x_0, for which

$$p(x_0) + k = \beta E p\{\tilde{h} - f(\tilde{h})\} \qquad (7)$$

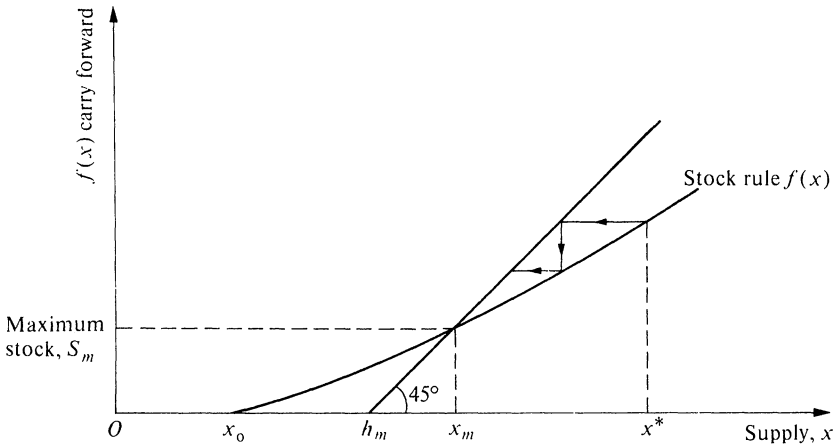The other characteristics of the optimum storage rule are not quite so

FIG. 1. Storage Rule

obvious, but should become clear from the arguments given in Appendix 1.

(ii) The stock function $f(x)$ is continuous and monotonically increasing.

(iii) The derivative $f'(x)$ is less than unity. This follows because current consumption is an increasing function of supply, i.e. $dC/dx > 0$, which, since $C = x - f(x)$, implies the result.

(iv) In a stationary world with bounded harvests stocks are also bounded. If the maximum possible harvest is h_m, then there is a unique number x_m such that

$$x_m = h_m + f(x_m), \qquad f(x) < x - h_m \quad \text{for} \quad x > x_m(> h_m). \tag{8}$$

This follows from results (ii) and (iii), which together imply that the stock function looks as in Fig. 1, with x_m defined by the unique intersection of $f(x)$ with the line $x - h_m$.

If by some unforeseen event supply ever rises above some level x_m, say to x^*, the stock level must steadily decrease, for even a sequence of bumper harvests h_m will lead to a successive decrease in supply and hence stocks, as shown by the arrowed line in Fig. 1.

3. Solving for the competitive stock rule

At least four different methods have been employed for finding the storage rule. The first two employ brute force and require the number crunching abilities of a computer. If supply is insensitive to price, then a fairly simple approach is available. Specify a time horizon (say, five years), and a zero terminal stock, and then calculate recursively the carryforward in year t given the probability distribution of supplies (and carryforward) in year $t+1$. This is the method described in detail in Gustafson, (1958), Gardner (1979), and adopted by Goreux (1978).

The second method has been developed by Wright and Williams (1981) to deal with supply responses, and is considerably more complicated. It involves iteratively refining a fourth order polynomial approximation to the relationship between the expected price and current carryover, $\psi(S_t) \cong Ep_{t+1}(S_t)$. The aim is to find by successive approximation for any given value of S_t a consistent set of planned production and expected price. Then the relationship between current supply and storage can be deduced from the arbitrage relationship. The effects of the storage rule are found by Monte Carlo simulation (10,000 trials). Clearly, such a method requires considerable programming ability. Both brute force methods require extensive sensitivity analysis to identify the contributions of the various parameters (discount rate, storage cost, demand and supply elasticity, and the magnitude of the underlying variability).

The third method is more suitable for the case of no supply response and is reported in Newbery and Stiglitz (1981b). The method is to expand $f(x)$ as a Taylor series in $(x - x_0)$, and solve by equating coefficients. This method works well if the distribution of the harvest $\bar{h}$ has a simple form, either discrete, or rectangular. It allows a reasonable approximation to be calculated in a few minutes using a pocket calculator.

The final method appears to be the best simple method available, and appears to be surprisingly accurate and quick. It was developed and tested by Gustafson (1958). The technique is to replace the random variable $\bar{h}$ on the right hand side of equation (6) with its expected value to find the form of $f(x)$, then using this approximation, to calculate x_0 from equation (7). Gustafson's calculations show that the general shape of $f(x)$ is approximated quite accurately by this assumption, but its horizontal location does depend on the variability of output. The first step finds the shape, the second locates the position.

Appendix 1 demonstrates that if the demand schedule is linear, then the approximation to the optimum stock rule, which we shall call $\phi(x)$, is piecewise linear, with successive segments having a steeper slope. Moreover, these slope coefficients and the length of the successive segments are given by simple formulae, and their numerical solution is immediate. They appear capable of replicating Gustafson's results to an accuracy of 1% over the whole range. If the demand schedule is non-linear, it is simple to take this into account in computing the linear approximation. For example, Appendix 1 shows that if the demand schedule has constant elasticity, the effect is to reduce the slope of the storage function, i.e., reduce the marginal propensity to store. Finally, Appendix 2 shows how to extend this method to allow for price responsive production, and again derives simple algebraic formulae. In this case the effect is to raise the marginal propensity to store, by an amount which increases with the ratio of the supply to the demand elasticity. Again, the method appears quite accurate at replicating the numerical results of Wright and Williams (1981).

The main reason for computing the optimum storage rule is to compare its

impact with either no storage or with the impact of some alternative rule, such as the band width rule. If the function has been computed numerically, the logical method is to use Monte Carlo simulation. If, however, the function can be reasonably approximated by a piecewise linear rule, then it is possible to proceed algebraically, as Newbery and Stiglitz (1981a, Ch. 30) show. However, the present purpose is to obtain a qualitative feel for the impact of the optimum storage rule, and to that end we seek a simple model specification for which the storage function is particularly simple. Fortunately, this is possible. First, assume demand is linear:

$$p(C) = \bar{p}\left\{1 - \frac{1}{\varepsilon}(C - \bar{C})\right\} \tag{9}$$

where ε is the elasticity of demand at mean consumption, $\bar{C}$, and price, $\bar{p}$. Choose units so that $\bar{C} = 1$, and assume that production is insensitive to price, so that $E\tilde{h} = \bar{C} = 1$. Equation (6) now becomes

$$\left.\begin{array}{c}\bar{p}(1-\beta)\left(1 + \dfrac{1}{\varepsilon}\right) + k \geq \dfrac{\bar{p}}{\varepsilon}\{x - f(x) - \beta E[h + f(x) - f\{h + f(x)\}]\} \\[2mm] x \geq x_0. \end{array}\right\} \begin{array}{l}\text{comple-}\\\text{mentary}\\\text{inequalities}\end{array} \tag{10}$$

where x_0 is defined by equation (7).

Next, suppose the marginal propensity to store is constant above x_0, so that the storage rule becomes

$$f(x) = \alpha(x - x_0), \qquad x_0 \leq x \leq x_m, \tag{11}$$

where x_m is defined by equation (8). Substitute (11) into (10) and rearrange to give

$$(1 + \beta)f(x) = x - a + \beta Ef\{\tilde{h} + \alpha(x - x_0)\}, \qquad x_0 \leq x \leq x_m \tag{12}$$

where a is a constant:

$$a = 1 + \varepsilon(1 - \beta + k/\bar{p}) = 1 + \varepsilon c, \qquad c \simeq r + k/\bar{p}. \tag{13}$$

The problem lies in the behaviour of the term $\tilde{h} + \alpha(x - x_0)$, which may be greater or less than x_0, the point at which $f(x)$ is non-linear. Define

$$\text{Prob}\{\tilde{h} \geq x_0 - \alpha(x - x_0)\} = \pi(x) \tag{14}$$

$$E\tilde{h} \mid \tilde{h} \geq x_0 - \alpha(x - x_0) = H(x) \tag{15}$$

then

$$(1 + \beta)f(x) = x - a + \alpha\beta\pi(x)\{H(x) + \alpha(x - x_0) - x_0\}, \qquad x_0 \leq x \leq x_m. \tag{16}$$

The stock rule can only be linear over $[x_0, x_m]$ if $\pi(x)$ is independent of x over this range; and $H(x)$ is constant or linear in x. This is equivalent to requiring

$$\text{Prob}\{\tilde{h} \mid (1 + \alpha)x_0 - \alpha x_m \leq \tilde{h} \leq x_0\} = 0. \tag{17}$$

For example, suppose the harvest has a two point distribution

$$h = \begin{cases} 1+u & \text{Prob } \rho \\ 1-\gamma u & \text{Prob } 1-\rho \end{cases} \qquad \gamma = \frac{\rho}{1-\rho}$$

so that $Eh = 1$, $\text{Var } h = \gamma u^2 \equiv \sigma^2$, the squared coefficient of variation.

To ensure a linear stock rule we require that if the current harvest is low $(h = 1 - \gamma u)$, then no matter how large was the carryover, current storage must be zero, and if the current harvest is high $(h = 1 + u)$, then even with zero carryover, some stocking must occur. In that case $\pi(x) = \rho$, $H(x) = 1 + u$, and equation (16) can be solved for x by finding ths value of x at which the RHS is zero:

$$x_0 = \frac{a - \alpha\beta\rho(1+u)}{1 - \alpha\beta\rho}, \tag{18}$$

The marginal propensity to store, α, is found by equating coefficients of x.

$$\alpha = \frac{1 + \beta - \sqrt{(1+\beta)^2 - 4\beta\rho}}{2\beta\rho}. \tag{19}$$

The maximum stock is found from equation (8).

$$x_m = \frac{1 + u - \alpha x_0}{1 - \alpha} \tag{20}$$

$$S_m = \alpha(x_m - x_0) = \frac{\alpha}{1 - \alpha}(1 + u - x_0) \tag{21}$$

The condition of positive carryovers if and only if the current harvest is good is equivalent to

$$1 - \gamma u + S_m < x_0 < 1 + u \tag{22}$$

which is identical to equation (17). Given specific values of β, ε, c, u, and ρ, α and x_0 can be found and checked to see if they satisfy the constraints of equation (22).

For example, if $\varepsilon = 0.8$, $\beta = 0.95$, $c = 10\%$, (so that $a = 1.08$) $\rho = \frac{2}{3}$ (so $\gamma = 2$), and $u = 0.15$ (so that $\sigma = 21.2\%$), then the parameters of the stock rule are

$$\alpha = 0.6501, \qquad x_0 = 1.031, \qquad x_m = 1.371, \qquad S_m = 0.221,$$

so that the conditions of equation (22) are satisfied. The optimum stock rule is

$$f(x) = 0.65(x - 1.031) \qquad x \geqslant 1.031 \tag{23}$$

It can be shown that condition (22) cannot be satisfied for a *symmetric* two point distribution (i.e. with $\rho = 0.5$), though the example we have chosen is empirically interesting, as it corresponds to the prospect of a periodic disastrous harvest (one year in three). However, it is easy to find a symmetric four point distribution for which the optimum rule is linear (above x_0).

For example, if the harvest can take values $1 \pm \frac{1}{2}u$, $1 \pm u$ with equal probability then if $a = 1.05$, $\beta = 0.95$, $u = 0.10$ (so that $\sigma = 7.9\%$) then $\alpha = 0.6$, $x_0 = 1.04$, $x_m = 1.19$, $S_m = 0.09$, and there will be positive carryovers only in the two good states.

3.1. *The impact of the optimum stocking rule*

The stocking rule increases current supply (by last year's carryover, possibly zero) and, in good years, reduces current sales (by the amount of the carryover). Both effects are readily described for the two state world just described. The probability distribution of last year's carryover is easily found, since it merely depends on the probability of a run of n previous good years. Let $S(n)$ be the size of the carryover in such cases, for which the probability is $(1-\rho)\rho^n$. Then

$$S(n) = \alpha(1 + u + S(n-1) - x_0), \qquad S(0) = 0,$$

$$S(n) = \alpha(1 + u - x_0) \cdot \left(\frac{1-\alpha^n}{1-\alpha}\right) \quad \text{with probability } (1-\rho)\rho^n, \qquad (24)$$

which clearly converges to the maximum value given in equation (21). The average stock carried will be

$$\bar{S} = \sum_0^\infty S(n)(1-\rho)\rho^n = \frac{\alpha\rho(1 + u - x_0)}{1 - \alpha\rho}. \qquad (25)$$

For example, for the parameters which gave equation (23), $\bar{S} = 9.1\%$ of average harvest. This large average stock size is very much a property of the shape of the probability distribution, which, in our example, is very disperse. Newbery and Stiglitz (1981a, Ch. 30) show that for the typical UNCTAD core commodity, if outputs are normally distributed, average stocks are more likely to be 5% of average harvest.

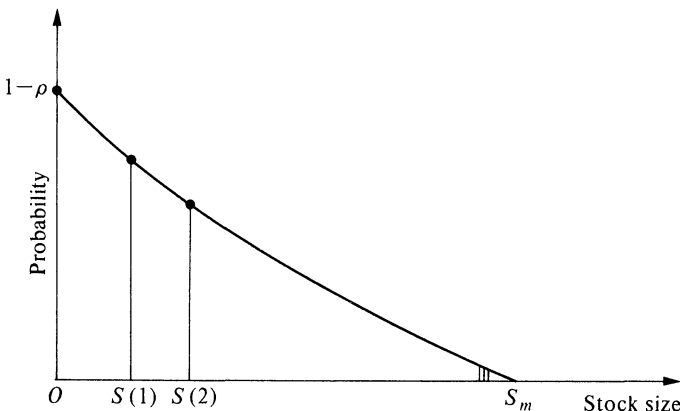The probability density of stock sizes is shown in Fig. 2, and will have the

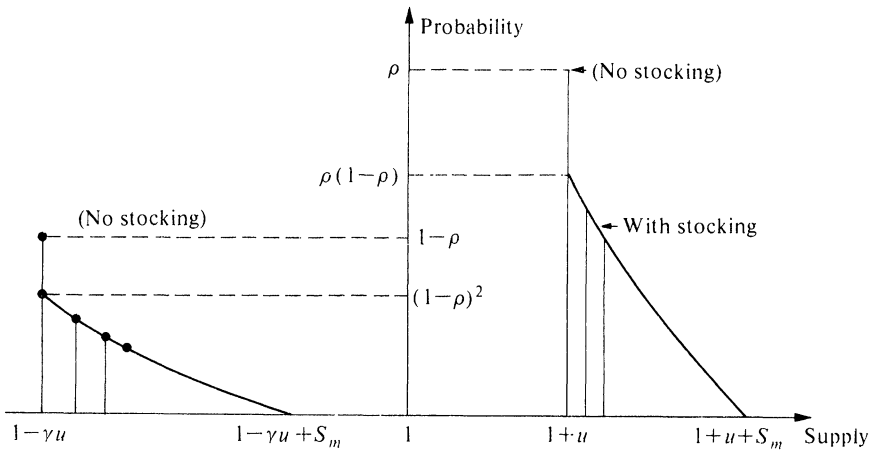FIG. 2. Probability density of stock size

FIG. 3. Probability density of supply

same general declining shape for all distributions of harvest, with a fairly high probability of zero stocks, and a rapidly decreasing probability of high stocks.

The variance of stock size can be computed in the same way as the average stock, and is

$$\text{Var } \tilde{S} = \left(\frac{1-\rho}{\rho}\right)\left(\frac{1}{1-\alpha^2\rho}\right)\bar{S}^2 \tag{26}$$

which, given the parameters of equation (23) is 0.00576, $CV = 83\%$.

The probability density of total supply is given by the sum of two independent random variables of current harvest and last year's carryover, and will look like Fig. 3.

The relationship between supply and price is given by equations (4), (5), (9) and (11):

$$p(x) = \begin{cases} \bar{p}\left\{\left(1+\dfrac{1}{\varepsilon}\right) - \dfrac{1}{\varepsilon}\,x\right\} & x < x_0 \\[3mm] \bar{p}\left[\left(1+\dfrac{1}{\varepsilon}\right) - \dfrac{1}{\varepsilon}\{(1-\alpha)x + \alpha x_0\}\right] & x \geqslant x_0. \end{cases}$$

The non-linearity in the supply-price relationship occurs at the point at which storage first occurs, and its effect is to increase the elasticity there by a factor $1/(1-\alpha)$, or 2.86.

The effect of storage on the effective schedule facing farmers is to shift the schedule to the left by the amount of last year's carryover, S, and to introduce a kink at the point $x_0 - S$, as shown in Fig. 4.

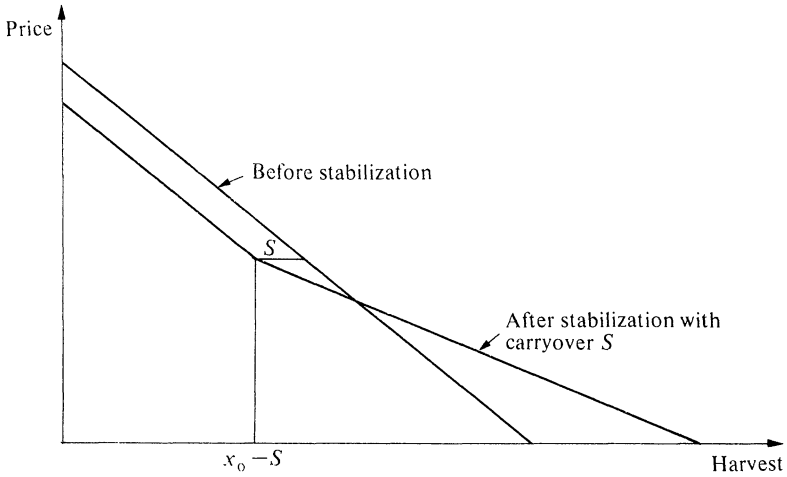It is interesting to note from (27) that the average price is unchanged by

FIG. 4. Effective demand schedule facing farmers

price stabilization in this linear example. The effect of introducing price stabilization is to raise farmer's average profits by

$$\Delta E\tilde{p}\tilde{h} = \frac{u\bar{p}\bar{S}}{\varepsilon} \tag{28}$$

whilst the effect on average social welfare, ignoring storage costs, is found from the utility function which corresponds to the demand schedule:

$$U(C) = \bar{p}\left\{\left(1 + \frac{1}{\varepsilon}\right)C - \frac{1}{2\varepsilon}\,C^2\right\}, \qquad \frac{dU}{dC} = p(C).$$

$$\Delta EU(C) = \bar{p}\left(-\frac{1}{2\varepsilon}\,\Delta EC^2\right) = \frac{\bar{p}}{2\varepsilon}\,E(C_u - C_s)(C_u + C_s) \tag{29}$$

where C_u is unstabilized consumption, C_s is stabilized consumption, and

$$C_s = C_u + \tilde{S} - f(\tilde{h} + \tilde{S}).$$

Solving:

$$\frac{\varepsilon}{\bar{p}}\,\Delta EU(C) = E\{-\tilde{S} + f(\tilde{h} + \tilde{S})\}\{\tilde{h} + \tfrac{1}{2}\tilde{S} - \tfrac{1}{2}f(\tilde{h} + \tilde{S})\}$$

$$= -\bar{S} - \tfrac{1}{2}ES^2 + \rho\alpha E(b + \tilde{S})\{1 + u + \tilde{S} - \tfrac{1}{2}\alpha(b + \tilde{S})\}, \qquad b \equiv 1 + u - x_0.$$

Note that, from (25), $\alpha\rho(b + \bar{S}) = \bar{S}$, so

$$\Delta EU(C) = \frac{\bar{p}}{\varepsilon}\{\bar{S}(u - \tfrac{1}{2}\alpha b\{1 - \rho(2 - \alpha)\}) - \tfrac{1}{2}(1 - \alpha\rho(2 - \alpha))ES^2\}. \tag{30}$$

This can be further simplified and/or solved using equation (26). For the

numerical example of equation (23) the gross benefit amounts to 1.3% of riskless expenditure.[4]

It is interesting to compare this with the gross benefits of perfect price stabilization—the conventional, but misleading benchmark. In this case equation (29) gives

$$\Delta EU(C) = \frac{\bar{p}}{2\varepsilon} \frac{\rho}{1-\rho} u^2$$

or 2.8% of riskless expenditure. However, we have so far ignored the costs or stabilization, which are considerable. The average annual cost of storage is $c\bar{p}\bar{S}$, or, in the present case, 0.9% of riskless expenditure and thus 70% of gross benefits, reducing net benefits to 0.4 of 1% of riskless consumer expenditure. Storage costs are a high proportion of the gross benefits of optimal stabilization and, obviously, a higher proportion of any greater degree of price stabilization. To assume costless price stabilization is misleading indeed.

3.2. The benefits of price stabilization

In general it is difficult to calculate the benefits of competitive price stabilization directly, and in any case, the method used for the linear stocking rule only gives the average net benefits, not the present value of introducing such a rule in any given year, with a known current supply. Instead it is necessary to use dynamic programming arguments, which are set out more fully and justified in Gustafson (1958), Samuelson (1971) or Newbery and Stiglitz (1977, Appendix F).

This approach has been applied in Newbery and Stiglitz (1981a, pp. 428–430) and gives particularly simple formulae in the case of the linear stocking rule. Thus, the present discounted benefit of introducing the storage rule when current supply is x is

$$B(x) = \frac{\alpha\bar{p}}{2\varepsilon} \left[\frac{\rho\beta}{1-\beta} (1+u-x_0)^2 + \{\min(x-x_0), 0\}^2 \right] \tag{31}$$

Obviously, it is more beneficial to start stockpiling in a year of high current supply, x, though if x is only slightly above x_0 the second term will be very small. If we ignore this term, and compute the benefit for the parameters of equation (23), the PDV is 7.287%, equivalent to an annual benefit of 0.384 of 1% of riskless consumer expenditures. (Compare this with the figure of 0.387 of 1%, computed above from equation (30).) Moreover, equation (31) is considerably more elegant and simple to evaluate than (30), showing the power of dynamic programming arguments.

3.3. Comparisons with the bandwidth rule

The bandwidth rule defines upper and lower intervention prices which are to be defended as far as possible by selling from or adding to stocks. Once

---

[4] We have chosen units so that average consumption is 1, as is average price, $\bar{p}$, so, in the absence of any instability (setting $u = 0$), consumer expenditure would also be unity. Benefits and costs are thus most easily measured as fractions of riskless expenditure. In the presence of fluctuations, with no stabilization average consumer's expenditure is $1 - (1/\varepsilon)\sigma^2$, or 0.94.

the intervention prices have been specified, then, assuming the demand schedule and the distribution of harvests are known, the size of the buffer stock follows a known stochastic process and can be analysed. The effect on the price distribution and the utility of consumption can also be determined, as can the cost of stock holding.

For example, in the case of the two point distribution of harvests, with harvest of $1 + u$ in good years (2 years out of 3) and $1 - 2u$ in bad years (1 year in 3) for which the optimum stock rule was given in equation (23), the obvious bandwidth rule is to place λu units into store in good years and remove $2\lambda u$ units in bad years. The buffer stock will now follow a generalised random walk, increasing by one step (of length λu) with probability 2/3, and being reduced by 2 steps with probability 1/3. The origin constitutes an elastic barrier, as would the maximum storage capacity. (If there were no such upper limit, the maximum stock size ever reached would increase without limit.) It is relatively straightforward to compute the average stock size as a function of the maximal capacity, and also compute the degree of price and consumption stability achieved.

Thus if the maximum capacity is $3\lambda u$ units, the average stock will be $2\lambda u$ units, and the distribution of consumption will be

$$
\begin{array}{lcl}
& 1 + u & \text{Prob } 24/75 \\
& 1 + (1 - \lambda)u & 26/75 \\
C = & 1 - 2(1 - \lambda)u & 16/75 \\
& 1 - u & 6/75 \\
& 1 - 2u & 3/75
\end{array}
$$

(The method of calculating the distribution of consumption is set out in Newbery and Stiglitz, 1981, Ch. 29.) Thus if $\lambda = 1/3$ the variance of consumption is reduced to 0.025 or by 45% of its unstabilized value, with average stock of 10% of average harvest. This bandwidth rule thus achieve essentially the same degree of price stabilization as the optimum rule (which achieved a reduction of 46%), with slightly higher average stocks. Its annual average net benefit is, however, only 0.3 of 1% of riskless expenditure or only 71% of that achieved by the optimum rule. If λ were chosen to be 0.5, the variance of prices would fall by 57%, but net benefits would be negative. In general as Newbery and Stiglitz (1981, §30.4.2) have shown, the best bandwidth rule achieves appreciably lower benefits than the optimum rule, whilst a carelessly chosen bandwidth rule will achieve the same price stabilization as the optimum rule at appreciably higher cost.

4. The bias in the competitive stock rule

The competitive stock rule is optimal if we ignore risk benefits and the distribution of income. Now it could be argued that the distribution of income should be ignored, since price stabilization is likely to be an inefficient method of influencing it, but in international negotiations alternative more efficient methods may not be available. We therefore ask, under

the competitive stock rule, would producers benefit from an increase in price stabilization (achieved, for example, by a subsidy on storage costs)?

This question can be answered by applying the methods developed in Newbery and Stiglitz (1979b) (and also set out in Newbery and Stiglitz, 1981a, Ch. 18). We ask, what will be the effect of storing an extra δQ units at date 0 when the harvest is high and the price low, to be sold next year when the harvest is (probably) smaller and the price higher. If Q_t is supply in year t, and C_t is consumption then, since the transfer only affects income in years 0 and 1, the change in the farmer's present discounted utility is solely the result of changes in p_0 and p_1:

$$\delta U = -U'(y_0)Q_0 \frac{\mathrm{d}p_0}{\mathrm{d}C_0} \delta Q + \beta EU'(y_1)Q_1 \frac{\mathrm{d}p_1}{\mathrm{d}C_1} \delta Q. \tag{32}$$

(Transferring δQ from year 0 lowers consumption then, which affects price, and hence income, $y_0 = p_0 Q_0$. The opposite effects occur in year 1. c.f. Newbery and Stiglitz, 1979, eq. 5, p. 804.)

If the demand schedule is linear, then

$$p(C) = \bar{p}\left\{1 - \frac{1}{\bar{\varepsilon}}(C - \bar{C})\right\}, \frac{\mathrm{d}p}{\mathrm{d}C} = -\frac{\bar{p}}{\bar{\varepsilon}},$$

where $\bar{\varepsilon}$ is the elasticity of demand at the mean price, $\bar{p}$. Substituting this in (32) gives

$$\delta U = \frac{\bar{p}}{\bar{\varepsilon}}\{U'(y_0)Q_0 - \beta EU'(y_1)Q_1\} \delta Q.$$

If the farmer is risk neutral, then U' will be constant, and his expected revenue (and utility) will increase with more storage, since Q_0 represents a large current harvest, above the expected future harvest, EQ_1. This result continues to hold for risk averse farmers provided $U'(y)$ does not increase too rapidly as Q falls and p rises. Provided income rises with price (i.e., provided demand is inelastic) the marginal utility $U'(y_1)$ will be lower than $U'(y_0)$, and producers continue to gain from more storage. Further analysis would identify the critical values of risk aversion and demand elasticity beyond which producers would prefer less stabilization. The only difference between the present case and that discussed in Newbery and Stiglitz (1979) is that there we ignored the effect of discounting, which, in the linear case, strengthens the case for stabilization.

If, on the other hand, the consumer demand schedule has constant elasticity, ε, then (32) can be rewritten as

$$\delta U = \frac{1}{\varepsilon}\left\{U'(y_0)\frac{p_0 Q_0}{C_0} - \beta EU'(y_1)\left(\frac{p_1 Q_1}{C_1}\right)\right\} \delta Q$$

If stocks are being transferred from date 0 to date 1 for consumption then, $Q_0 > C_0$, and $Q_1 < C_1$. For risk neutral producers, with constant $U'(y)$, the

change in revenue is then

$$\delta Y < \frac{\delta Q}{\varepsilon}(p_0 - \beta E p_1) = -\frac{k\,\delta Q}{\varepsilon} < 0$$

from the arbitrage equation (3). Hence a producer facing a constant elasticity of demand receives less revenue on average as the amount of storage increases. This result continues to hold for risk averse producers provided $U'(y)$ does not fall too far as price rises. This will be ensured if demand is elastic.

Changes in storage will affect consumers as well, but it seems reasonable to ignore the impact of risk on consumers, on the grounds that for the core commodities specified by UNCTAD consumers typically spend less than 1% of their income on these goods, and hence their marginal utility of income will be insensitive to price fluctuations. If transfers from (richer) consumers to (poorer) producers are considered socially desirable, then we can summarize the results of this section as follows:

(i) With linear consumer demand, inelastic over the relevant range, competitive price stabilization should be subsidized or supplemented.

(ii) With constant elastic consumer demand of elasticity greater than unity, competitive price stabilization should be taxed or reduced.

The determination of the optimal degree of intervention can be found by calculating the increased cost of further subsidization and the increased benefit in reduced producer income risk and increased transfers (weighted by the difference in producer and consumer weights).

4.1. Alternative stabilization policies and the problem of speculative attack

Once it is appreciated that producer's income risk is an important determinant of the attractiveness of price stabilization, the obvious question to ask is whether there is some more direct method of providing essentially income insurance without the conventional problems of moral hazard and adverse selection. This problem is discussed in Newbery and Stiglitz (1981a, Ch. 20) where it is shown that if so, the potential benefits are appreciable. We suggested a rule where producers and consumers face different prices. The consumers face market clearing prices whilst producers are paid a price

$$\hat{p}_t = \frac{\bar{Y}}{Q_t}$$

where $\bar{Y}$ is trend revenue for the region and Q_t is the actual production at date t. This scheme stabilizes individual producer's income to the extent that his output is correlated with regional output, but it apparently avoids the usual insurance problems of adverse selection and moral hazard. The first and most obvious problem is that the consumer price would typically differ from the producer price, and hence provide strong incentives for black markets to develop. The second problem is that it might be vulnerable to private storage by producers.

Assume, for instance, that there were only two states of nature, a good harvest with output $1+\sigma$ and a bad harvest with output $1-\sigma$. If there is a large harvest this year, the expected (proportional) increase in price is approximately

$$\frac{\frac{1}{2}\left(\frac{1}{1+\sigma}+\frac{1}{1-\sigma}\right)-\frac{1}{1+\sigma}}{\frac{1}{1+\sigma}}=\frac{1}{1-\sigma}-1=\frac{\sigma}{1-\sigma}$$

which, for large σ will exceed the rate of interest and private storage costs by a considerable amount. This provides a strong incentive for private storage (speculation). Parenthetically, the same problem of the vulnerability of buffer stock policies to speculative attack arises with the bandwidth rule.

Thus, the extent to which the buffer stock agency is constrained by private speculation in determining producer prices (assuming that it can separate that policy from the policy for consumer prices) is determined by the magnitude of the storage (interest) costs which they face, and in particular, by the extent to which they exceed those facing the buffer stock agency (net of any subsidy which the agency might receive).

4.2. The costs of increased price stabilization

The competitive storage rule is optimal for reducing fluctuations in aggregate consumption, but is not efficient if producers (or consumers) are risk averse. The effect of additional (or reduced) price stabilization on producers' income risk are reasonably easy to calculate, using, for example, the techniques developed in Newbery and Stiglitz (1979b) and illustrated above in Section 3. Here we are concerned to measure the costs of increasing (or reducing) price stabilization above (or below) the competitive level. The logical way to achieve this is to subsidize (or tax) storage (*not* the discount rate, since we wish the optimal inter-temporal pattern of stabilization to be retained—unless there is reason to believe that the market rate of discount differs from the efficient rate, in which case the remedy is obvious). If storage costs are reduced by a proportional subsidy, τ, the effect will be to lower x_0, raise $\bar{S}$, and lower the variance of prices and consumption. Here we examine the magnitude of these effects, using the two-state model developed in Section 3.1 for which exact formulae are available.

First, note that a is reduced, because from (13) a now becomes

$$a = 1 + \varepsilon\{1-\beta+(1-\tau)k/\bar{p}\}$$

Changes in a have no effect on the slope coefficient, α, but do affect x_0, for, from (18)

$$\frac{\mathrm{d}x_0}{\mathrm{d}\tau}=\frac{1}{1-\alpha\beta\rho}\frac{\mathrm{d}a}{\mathrm{d}\tau}=\frac{-\varepsilon k/\bar{p}}{1-\alpha\beta\rho}\,.$$

The effect on the average amount stored is given by (25)

$$\frac{\mathrm{d}\bar{S}}{\mathrm{d}x_0} = \frac{-\alpha\rho}{1-\alpha\rho}$$

whilst the effect on the annual average net benefits of stabilization, V_0 (ignoring distribution and risk) is given by differentiating the annualized form of equation (31).

$$V_0 \equiv \frac{1-\beta}{\beta} B(x_0)$$

$$\frac{\mathrm{d}V_0}{\mathrm{d}x_0} = -\frac{\bar{p}}{\varepsilon} \rho\alpha(1+u-x_0).$$

However, V_0 is calculated on an estimated storage cost which is $\tau k\bar{S}$ on average too low, so the annual *social* net benefit, Z, is changed by

$$\frac{\mathrm{d}Z}{\mathrm{d}\tau} = \frac{\mathrm{d}}{\mathrm{d}\tau}(V_0 - \tau k\bar{S}) = \frac{\mathrm{d}V_0}{\mathrm{d}x_0}\frac{\mathrm{d}x_0}{\mathrm{d}\tau} - k\bar{S}\left(1+\tau\frac{\mathrm{d}\bar{S}}{\mathrm{d}k}\right)$$

$$= \frac{k\alpha\rho(1+u-x_0)}{1-\alpha\beta\rho} - k\bar{S}\left(1+\frac{\alpha\rho}{1-\alpha\rho}\frac{\varepsilon k\tau}{1-\alpha\beta\rho}\right)$$

which, from (25), can be rewritten as

$$\frac{\mathrm{d}Z}{\mathrm{d}\tau} = -\frac{\alpha\rho k\bar{S}}{1-\alpha\beta\rho}\left\{1-\beta+\frac{\varepsilon\tau k}{1-\alpha\rho}\right\}$$

With the parameters given in equation (23) this has a value at $\tau = 0$ of -0.017 of 1%. The effect on average stock size is more dramatic though, with $\mathrm{d}\bar{S}/\mathrm{d}\tau = 5.2\%$ (of average output). Thus a 50% subsidy on direct storage costs (a reduction in the annual cost of storage of $2\frac{1}{2}\%$ of the average value of stock) would incur extra costs of 0.014 of 1% (of riskless consumer expenditure) and increase average stocks from 9.1% (of average output) to 11.7%.

With a continuous probability density of harvests the effect would be greater, as the probability of storage, as well as the amount then stored, would also rise, in contrast to the present two-state case.

5. Conclusions

The objective of the paper has been to characterize the optimal storage rule and hence the best way of stabilizing prices on competitive markets. We showed that if agents were risk neutral and held rational expectations, then the competitive arbitrage rule would be efficient. The first part of the paper showed how to characterize this competitive storage rule, and found a specification in which the rule could be solved analytically. This analytical solution was used to explore the degree of price stabilization, the costs and benefits of efficient stabilization compared with the alternative bandwidth

rule, and the effect of storage on the net demand schedule facing the farmer. In the simple case the storage rule is to store a constant fraction of the excess of current supply (stocks plus harvest) over a fixed amount. The appendices show how to derive successively more accurate approximations to the competitive storage rule in more general cases. The approximate rule will be piecewise linear for a linear demand schedule, with successive, steeper segments reflecting the probability that the stock will be held for one further year. If, however, the demand schedule has constant elasticity, instead of being linear as in the special case, then the marginal propensity to store is lowered, whilst if supply is responsive to the changes in the expected future price induced by storage, then the marginal propensity to store is increased.

The probability distribution of stocks appears quite skew, with a rapidly falling probability of high stocks, and a relatively small average stockpile (typically less than 10% of average harvest). The net gains from the optimal storage scheme, though small, are significantly greater than those obtained from the alternative storage policies (e.g. maintaining prices within a bandwidth) often proposed. It should be observed, however, that the optimal buffer stock scheme leaves considerable remaining price variability, typically reducing the variance of prices by about one-half.

The second part of the paper explored the bias in the competitive rule when agents are risk averse and distributional issues are important. Techniques which the authors developed elsewhere allow the direction of the bias to be identified, and this bias can in principle then be corrected by taxing or subsidizing private storage. The costs of such taxes and subsidies can also be readily calculated, given the form of the storage rule.

The main conclusion is that the optimal storage rule is quite different from the rules typically adopted by stabilization authorities, and yet can be approximated by a simple form. We would urge future researchers to employ this simple form (or more accurate approximations) in studying questions of price stabilization.

APPENDIX 1: APPROXIMATION TECHNIQUES FOR FINDING THE OPTIMAL STOCK RULE

Gustafson (1958) found that the optimal stock rule, $f(x)$, had a form which was relatively insensitive to the various parameters, and in particular, to the degree of harvest variability. Different coefficients of variation of harvest essentially displaced the function horizontally. He therefore suggested a short cut for obtaining an approximate solution—first calculate an approximate function $\phi(x)$ assuming no variability in harvest, and then use this approximation to estimate an accurate value of x_0 using equation (7).

First Step: The certainty equivalent approximation

This 'certainty equivalence' method is useful in characterizing and explaining the general form of the stocking rule, as well as providing a good approximation. The approximate function $\phi(x)$

satisfies

$$p\{x - \phi(x)\} + k = \beta p[1 + \phi(x) - \phi\{1 + \phi(x)\}]$$
$$x \geq x_0 \tag{A1}$$

where the random harvest $\tilde{h}$ has been replaced by its certainty equivalent of unity. The key to solving the functional equation (A1) is to note that if x is only slightly above x_0, that all the carryover $\phi(x)$ will be consumed next year. At some point, x_1, this year's carryover is sufficiently large that it will last two years, and at a higher level, x_2, it will last three years, and so on. Let $\phi_n(x)$ be the stocking rule relevant to a supply x which will be sufficient for n years of carryover, then (A1) can be rewritten

$$p\{x - \phi_n(x)\} + k = \beta p[1 + \phi_n(x) - \phi_n\{1 + \phi_{n-1}(x)\}] \tag{A2}$$

Consider the first segment of $\phi(x)$, $\phi_1(x)$, for which $x_0 < x < x_1$, and stock will all be sold next year. Since $\phi_0(x)$ is by definition 0 (no carryover), (A2) simplifies to

$$p\{x - \phi_1(x)\} + k = \beta p\{1 + \phi_1(x)\}, \qquad x_0 < x < x_1. \tag{A3}$$

The point x_0 is the point at which $\phi_1(x) = 0$, and is found from

$$p(x_0) + k = \beta p(1)$$

whilst x_1 is the point at which next year's supply would reach x_0, so that any higher initial supply would be carried two years. Hence x_1 satisfies

$$x_0 = 1 + \phi_1(x_1) \tag{A4}$$

The function $\phi_1(x)$ is implicitly defined by (A3), which is readily solved if $p(C)$ is linear. Here, our main interest is in the general form of $\phi(x)$, as exemplified by its slope. Differentiate (A3) and rearrange to give

$$\phi_1' = \frac{1}{1 + \dfrac{\beta p_1'}{p_0'}}, \qquad p_i' = \frac{dp}{dC}\bigg|_{C = C_i}$$

Clearly, if $p(C)$ is linear, then ϕ_1' has a constant slope over (x_0, x_1). In general, however, p_1'/p_0' is likely to be nearly constant, for two related reasons. First, arbitrage moves prices close together. With perfect arbitrage, $p_0 = p_1$ and clearly the slopes would be the same. With low carrying costs $(k + rp_0)$ the difference in the slopes will be small, and hence the ratio will be close to unity. Second, the prices will differ by an essentially constant absolute and relative amount $(k + rp_0)$, and so one might expect the slopes at the two points to bear a roughly constant relationship to each other.

Let $\alpha_1 = \phi_1'$, then

$$\alpha_1 \simeq \frac{1}{1 + \beta}$$

with exact equality for linear demand function. The same argument applies to each segment, and the slope, α_i, is found by differentiating (A2):

$$\phi_n' = \frac{1}{1 + \beta p_1'(1 - \phi_{n-1}')/p_0'} \qquad x_n < x < x_{n+1} \tag{A5}$$

$$\alpha_n \simeq \{1 + \beta(1 - \alpha_{n-1})\}^{-1} \tag{A6}$$

whilst x_n satisfies (cf A4)

$$1 + \phi_n(x_n) = x_{n-1} \tag{A7}$$

These two recursive relations can be readily solved to find the piecewise linear approximation to the certainty equivalent rule:

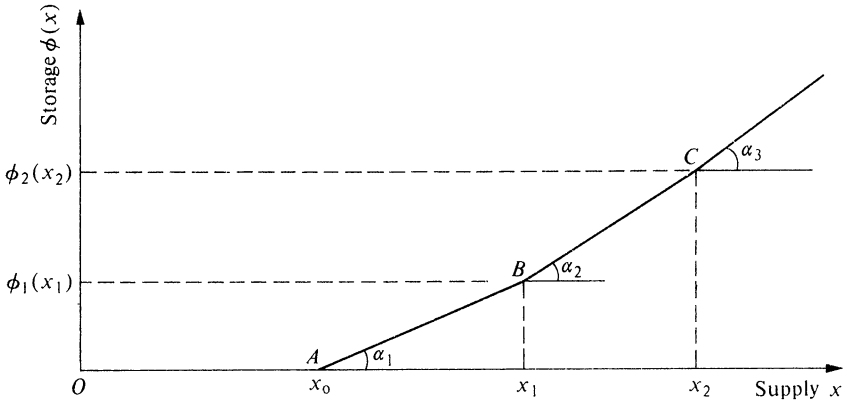$$\alpha_n = \frac{1 - \beta^n}{1 - \beta^{n+1}} \tag{A8}$$

FIG. A1: Piecewise linear approximation to storage rule

whilst

$$\phi_n(x) = \alpha_n(x - x_{n-1}) + \phi_{n-1}(x_{n-1}). \tag{A9}$$

Equations (A7) and (A9) imply

$$1 + \alpha_n(x_n - x_{n-1}) + \phi_{n-1}(x_{n-1}) = x_{n-1}$$

or

$$\alpha_n(x_n - x_{n-1}) = (x_{n-1} - x_{n-2}), \; x_{-1} \equiv 1. \tag{A10}$$

The coefficients α_n are immediately soluble from (A7), and the lengths of the segments are then solved from (A10). Fig. A1 illustrates.

Equation (A5) gives immediately one key characteristic of the storage rule—its slope is monotonically increasing but always strictly less than unity—which is equivalent to the conclusion that current consumption, $C = x - \phi(x)$, is an increasing function of supply, x.

With a non-linear convex demand schedule, since $p_1 > p_0$ for storage, $p_1'/p_0' > 1$ and the slope coefficients α_n will be somewhat reduced, the effect being similar to a rise in the discount factor, β, in equation (A8). For a demand schedule with constant elasticity ε the effect is one of replacing β in (A8) by

$$\beta' = \beta\{1 + (1 + \varepsilon)(r + k/p_0)\}$$

For $\varepsilon = 1$, $r + k/p_0 = 8\%$, $\beta' = 1.1$, and α_1 falls from 0.513 to 0.476. Less marginal storage is done, for with a convex demand schedule, current prices rise more rapidly with storage than with a linear demand schedule. (This argument refers to *marginal* storage—the location of x_0, which affects the amount of total storage, will be affected by the shape of the demand schedule as well.)

Second Step: Allowing for random harvests

Once the piecewise linear approximation has been computed, the next step in finding the approximate storage function is to recompute x_0 more accurately as the solution to

$$x_0 = a - \beta E\phi(h)$$

$$x_0 \cong a - \frac{\beta}{1 + \beta} \int_a^{h_m} (h - a) \, dF(h) \tag{A11}$$

where $F(h)$ is the distribution function of h, and f has been approximated by the first line segment, ϕ_1. For example, if h is normally distributed as $N(1, \sigma^2)$ with distribution function Φ

then the second term of equation (A11) is approximately

$$\psi(a) = \frac{\beta}{1+\beta}\left[\frac{\sigma}{\sqrt{2\pi}}\exp\left\{-\frac{1}{2}\left(\frac{a-1}{\sigma}\right)^2\right\} - (a-1)\left\{1-\Phi\left(\frac{a-1}{\sigma}\right)\right\}\right] \tag{A12}$$

APPENDIX 2: THE EFFECT OF SUPPLY RESPONSE ON THE STOCK RULE

It is relatively straightforward to study the effect of allowing planned future production to depend on the expected future price in a linear model with additive risk and risk neutral farmers. (Newbery and Stiglitz, 1981, Ch. 5–6 show that the action certainty equivalent price is then equal to the expected price.) Let demand be linear with elasticity ε at the mean price, and choose units so that in the absence of any stockpiling, average consumption is unity, as is average price. Then

$$p(C) = 1 + \frac{1}{\varepsilon} - \frac{1}{\varepsilon}C \tag{A13}$$

If planned production likewise has an elasticity η at the mean price, planned production will be

$$q = 1 - \eta + \eta Ep \tag{A14}$$

The arbitrage equation (3) implies that

$$q_{t+1} = 1 - \eta + \frac{\eta}{\beta}(p_t + k) \qquad x \geq x_0$$

$$= q_0 \qquad x < x_0$$

where q_0 is planned production if the current supply is too low to affect the future price. If current supply is x, then

$$C_t = x - f(x), \qquad C_{t+1} = q_{t+1} + \bar{u} + f(x) - S_{t+1}$$

$$S_{t+1} = f\{q_{t+1} + \bar{u} + f(x)\}, \qquad E\bar{u} = 0$$

where $q_{t+1} + \bar{u}$ is the actual harvest. The arbitrage equation can now be written, for $x \geq x_0$

$$p_t + k = \beta Ep(C_{t+1}) = \beta\left(1+\frac{1}{\varepsilon}\right) - \frac{\beta}{\varepsilon}\left[1 - \eta + \frac{\eta}{\beta}(p_t + k) + f(x) - ES_{t+1}\right]$$

Collecting terms; this becomes

$$\frac{\varepsilon}{\beta}(1+\eta/\varepsilon)(p_t + k - \beta) + f(x) = ES_{t+1}$$

Substitute for $p_t = p(C_t)$, and define

$$a - 1 + \varepsilon(1 - \beta + k) \simeq 1 + \varepsilon(r + k)$$

Then

$$f(x)\left\{1 + \frac{1}{\beta}(1+\eta/\varepsilon)\right\} = \frac{1}{\beta}(1+\eta/\varepsilon)(x - a) + ES_{t+1}$$

or

$$f(x) = \lambda(x-a) + \mu Ef\left[q_0 + \frac{\eta x_0}{\beta\varepsilon} + \bar{u} + \left(1+\frac{\eta}{\beta\varepsilon}\right)f(x) - \eta x/\beta\varepsilon\right] \tag{A15}$$

where

$$\lambda = \left(1 + \frac{\varepsilon\beta}{\varepsilon+\eta}\right)^{-1}, \qquad \mu = \frac{\beta}{1+\beta+\eta/\varepsilon}.$$

This can be solved using the approximation techniques set out in Appendix 1. Thus, setting $\bar{u} = 0$, x_0 solves

$$0 = \lambda(x_0 - a) + \mu f(q_0)$$

Now q_0 solves

$$q_0 = 1 - \eta + \eta E \left[1 + \frac{1}{\varepsilon} - \frac{1}{\varepsilon} \{ q_0 + \tilde{u} - f(q_0 + \tilde{u}) \} \right]$$

If instead of taking expectations we replace $\tilde{u}$ by its certainty equivalent value 0 this expression reduces to

$$q_0 = 1, \qquad f(q_0) = 0, \qquad x_0 = a. \tag{A16}$$

Once again, $f(x)$ can be approximated by the piecewise linear certainty equivalent function $\phi(x)$ with successive segments having slope α_n, found by differentiating (A3):

$$\alpha_n = \lambda + \mu \alpha_{n-1} \left[\left(1 + \frac{\eta}{\beta \varepsilon} \right) \alpha_n - \eta / \beta \varepsilon \right]$$

If $\qquad\qquad n/\varepsilon \equiv \delta < 1$ (for stability) then

$$\alpha_n = \left\{ 1 + \frac{\beta(1 - \alpha_{n-1})}{1 + \delta(1 - \alpha_{n-1})} \right\}^{-1} \tag{A17}$$

which reduces to the same formula as given in Appendix 1 if supply is inelastic, $\eta = 0$, and $\delta = 0$. The points of slope change, x_n, satisfy

$$q_0 + \left(1 + \frac{\delta}{\beta} \right) \phi_n(x) - \frac{\delta}{\beta} (x_n - 1) = x_{n-1}, q_0 = 1.$$

This can be rearranged to give (after factoring out $1 + \delta/\beta$)

$$1 + \alpha_n (x_n - x_{n-1}) + \phi_{n-1}(x_{n-1}) = x_{n-1}$$

which yields the same result as (A10).

The effect of supply response

If we restrict attention to the certainty equivalence rules, the critical point at which stocking occurs is the same—$x_0 = a$. Moreover, the points x_n satisfy the same equation (A10). However, equation (A17) shows that a positive value for γ (and δ) increases the slope coefficients α_n and hence results in more storage (to offset the negative serial correlation in supply).

Another way to interpret this result is that since planned production will fluctuate in a way uncorrelated with the current disturbance, the total variability in production will be larger the larger is the supply response, η. More production variability implies more storage, to achieve a given degree of price stability.

World Bank, Washington, D.C.,
Princeton University, New Jersey.

REFERENCES

ADAMS, F. G. and KLEIN, S. A. (1978), *Stabilizing World Commodity Markets*, Lexington, Heath-Lexington.

BEHRMAN, J. R. (1978), 'International Commodity Agreements: An Evaluation of the UNCTAD Integrated Programme for Commodities', Ch. 16 in Adams and Klein (1978).

BRANDT, W. (1980), *North-South: A Programme for Survival*, Pan Books, London.

EDWARDS, R. and HALLWOOD, C. P. (1981), 'The determination of optimum buffer stock intervention rules', *Quarterly Journal of Economics*, Feb., pp. 151–166.

GARDNER, B. L. (1979), *Optimal Stockpiling of Grain*, Lexington, D.C. Heath.

GOREUX, L. M. (1978), 'Optimal Rule of Buffer Stock Intervention', IMF Research Dept. mimeo DM/78/7, Washington, D.C.

GUSTAFSON, R. L. (1958), *Carryover Levels for Grains*, U.S. Dept of Agriculture, Technical Bulletin No. 1178.

NEWBERY, D. M. G. and STIGLITZ, J. E. (1977), *The Economic Impact of Price Stabilization*, Report presented to U.S. Agency for International Development, Washington D.C.

NEWBERY, D. M. G. and STIGLITZ, J. E. (1979), 'The Theory of Commodity Price Stabilization Rules: Welfare Impacts and Supply Responses', *Economic Journal*, No. 89, December, pp. 799–817.

NEWBERY, D. M. G. and STIGLITZ, J. E. (1981a), *The Theory of Commodity Price Stabilization*, Oxford University Press.

NEWBERY, D. M. G. and STIGLITZ, J. E. (1981b) 'Optimal Commodity Stockpiling Rules', Economic Theory Discussion Paper No. 40, Cambridge, England.

NEWBERY, D. M. G. and STIGLITZ, J. E. (1982), 'The Choice of Techniques and the Optimality of Market Equilibrium with Rational Expectations', *Journal of Political Economy*, 90, April, pp. 223–246.

SAMUELSON, P. A. (1971) 'Stochastic Speculative Price', *Proc. Nat. Acad. Sci.*, 68, Feb. pp. 335–37. Reprinted in *Collected Scientific Papers*, Vol. III.

STIGLITZ, J. E. (1969), 'Behaviour towards risk with many commodities', *Econometrica*, 37, pp. 660–67.

TURNOVSKY, S. J. (1978). 'The Distribution of Welfare Gains from Price Stabilization: A Survey of some Theoretical Issues', in Adams and Klein (1978).

TURNOVSKY, S. J., SHALIT, H. and SCHMITZ, A. (1980), 'Consumer's Surplus, Price Instability, and Consumer Welfare', *Econometrica*, Vol. 48, No. 1 (Jan.) pp. 135–152.

UNCTAD (1974), 'An Integrated Programme for Commodities', TD/B/C. 1/166, Trade and Development Board, Committee on Commodities, United Nations Conference on Trade and Development, Geneva.

WRIGHT, B. D. (1979), 'The Effects of Real Production Stabilization: A Welfare Analysis under Rational Behavior', *Journal of Political Economy*, 87, Oct., pp. 1011–33.

WRIGHT, B. D. and WILLIAMS, J. C. (1981) 'The Economic Theory of Commodity Storage', mimeo, Yale University.