DEPARTMENT OF AGRICULTURAL AND RESOURCE ECONOMICS
DIVISION OF AGRICULTURE AND NATURAL RESOURCES
UNIVERSITY OF CALIFORNIA

THE BEHAVIOR OF MARKETS FOR STORABLE COMMODITIES

by

Brian D. Wright

and

Jeffrey C. Williams

This paper is to be presented at the closing plenary session of the 34th. Annual Conference of the Australian Agricultural Economics Society, University of Queensland, Brisbane, Australia, February 15, 1990.

California Agricultural Experiment Station
Giannini Foundation of Agricultural Economics
February, 1990

THE BEHAVIOR OF MARKETS FOR STORABLE COMMODITIES

Brian D. Wright and Jeffrey C. Williams

One consistent policy prescription in the history of economic advice on commodity markets has been that prices should be stabilized in a symmetric band around the mean to reduce the "boom and bust" gyrations typical of commodity prices (Keynes, Nurkse, Newbery and Stiglitz, Knudsen and Nash). Attached to this policy advice is a little proviso for completeness, to tie up a little technical loose end: Adjust the mean along the long-run trend. As Streeten (1986) puts it, "A good guideline is: keep domestic cereal prices in line with an estimated trend of future world prices (estimated by a reputable authority)..."

This piece of economic policy advice is of particular interest because so many stabilization schemes have tried to follow it (Gardner, 1985, and Gilbert, forthcoming, include recent surveys). It is worth further attention because they almost never succeed for very long—and I do not mean long in the sense of the Keynesian long run. The founders easily survive the life span of the typical scheme, physically if not financially. Why this lack of success?

Economists have contributed few insights of value. A precise argument for pricing at near long-run marginal cost, rather than at short-run marginal cost as usually prescribed for other problems, is rarely offered. A frequent rationale given for failure is "insufficient financing," which usually means not too much more than offering "lack of funds" as an explanation for bankruptcy.

In this paper I shall focus on some elementary positive problems with following a price band buffer stock rule for a storable commodity. I shall deal in feasibility, not optimality and concentrate on two issues:

1. How easy is it to identify the long-run trend for a price band rule to follow?"
2. What can a price band rule achieve in the best of circumstances?

1. Finding the Long-Run Trend

Measured quarterly or annually, actual time series of spot prices for major commodities have two common features. First, they display considerable positive autocorrelation Years with high prices tend to follow years with high prices and low prices to follow low prices. Second, time series of spot prices have spikes, that is, years when the price jumps abruptly to a very high level relative to its long-run average.

Spot sugar prices, shown in Figure 1 annually for 38 years, provide a good example of both features. The spikes combine visually, but the tendency for the (deflated) price to remain in the doldrums for years on end, such as the stretch over the mid-1980s, is unmistakable. The first-order serial correlation of the series is 0.53.

Perusal of the World Bank's (1989) report on commodity prices reveals commodity after commodity with these two features. For example, a plot for copper looks much the same as sugar. Spikes in copper prices transpired over 1973-74, 1979-80, and 1987-89, while the mid 1970s and early and mid 1980s had prolonged low prices. The two spikes in the 1970s also occurred in other base metals, in the principal grains, and in energy products—each time heightening concern about general inflation (Cooper and Lawrence, 1975, and Bosworth and Lawrence , 1982).

These two stylized facts about spot commodity prices are the natural result of storage. Whereas a series of spot prices for a nonstorable commodity subject to independent disturbances in harvests (and no systematic change in long-run demand or supply) would behave as pure white noise, a time series for a storable commodity has an autoregressive structure, because storage spreads unusually high or low excess demand over several periods. Moreover, the skewness of the distribution of price under storage produces a time series in which increases in price, while rare, are larger than typical price decreases. This gives a price series for a storable commodity something of a sequence of cusps. Finally, because of storage, the variance of price changes depends on the particular current spot price. The contrast in variance is most obvious between a very low current spot price, which finds an overflowing stockpile able to buffer price changes, and a very high current spot price, which finds the stockpile empty and unable to buffer price change. Storability causes heteroskedasticity in the time series.

Positive autocorrelation, sometimes of more than first-order, is true of all the important endogenous variables (price, consumption revenues (and price), except the realized harvest.

Because a large carryout makes it less likely that the stockpile will be exhausted, storage itself is highly serially correlated, as is expected price. Expected price here is analogous to a one-year-ahead futures price for a commodity such as corn. The sequence of such futures prices (measured annually) should be more highly autocorrelated the lower are storage costs.

Figure 2 shows a series of 120 spot prices drawn from our rational expectations, competitive, profit-maximizing storage model with linear consumer demand with elasticity -0.2, supply elasticity 0.0, coefficient of variation of i.i.d. harvest disturbance of 10 percent, storage cost of 0, and interest rate r= 5% per period, with an infinite horizon. (See for example Wright and Williams 1988 or forthcoming.) Because the eye wants to impose order where there is chaos, it often sees patterns when the stochastic process is pure white noise. In Figure 2 the patterns are real, however. The model's spot prices have the two stylized facts: occasional peaks and positive serial correlation. Based on a yet longer series than used in Figure 2, 10,000 periods to be exact, these numbers are effectively the population parameters. The first-order autocorrelation is 0.67, not much different from that seen for sugar. Part and parcel, in the storage model the spot price in period t is also correlated with price in period t-2, t-3, and t-4, albeit the connection gets ever fainter.

The degree of first-order serial correlation in the spot price is endogenous, ultimately dependent on the potential for storage. The higher storage costs are, the weaker the connection between periods and the lower the proportion of the variance in the spot price that can be explained by the previous spot prices.

A positively autocorrelated series by its nature has high prices tending to follow high prices, a "boom," and low prices following low prices, a "bust." Thus, several years running of good times or of bad times should be commonplace for a storable commodity.

Yet the serial correlation has another, perhaps pernicious, effect. It makes it much more difficult to determine whether an apparent boom or bust results from a temporary shock or from a change in the long-run average price. Inferences about the long-run price from small samples are inescapable in many investment and policy settings. For example, MacAvoy (1988, p.xi) relates that for him and other members of the board of Amax Corporation "while those of other materials recovered over the 1983-1986 period, the question became that of which metals prices had remained at startlingly low levels for over five years." The question could be rephrased to ask whether Amax should invest in new mines, which would not come into production for several years. Another example is the U. S. farm bill, which is revised about every five years. The U. S. Congress finds itself in the role of judging whether prices in the intervening five years have indicated a shift in the long-run average price. But five years is too brief an interval from which to deduce long-term changes.

A reasonable approach to whether the long-run price has changed over the last five years is from the perspective of confidence intervals. One might imagine statements of the form: "With 99% confidence we can say the long-run mean price has shifted down from the previous average; indeed, with 90% confidence we can say the long-run price is $10 below the previous long-run mean." These statements would presumably be made after calculating the average price and standard deviation over the five-year period, and constructing the confidence interval from the t-statistics in the universal tables. Every elementary statistics book recommends these techniques.

Behind Figure 3 are confidence intervals constructed in such a fashion with 50,000 samples of each length drawn from the same long series behind Figure 2. The standard of $88 or $112 is one standard deviation from the long-run mean price of $100. Thus, Figure 3 tells the chance of concluding with 95% confidence that the mean price is above $112 or below $88. Because the long-run mean price is, in fact, $100, this probability should be very, very small if the procedure is performing as expected. The probability, however, is not very, very small. There above a 5% chance of falsely concluding that the long-run price is $12 from $100, especially that the price is too low. The problem with the procedure is not in the notion of confidence intervals, but the implicit assumption of the conventional t-test that each observation is independent of the others. With a sample of spot prices from a storable commodity, that assumption is not tenable. The positive serial correlation causes the sample standard deviation to be a considerable underestimate of the population variance of the sample mean (Flavin, 1984). As a result, the confidence intervals are much too small. With positive serial correlation, a sample of size five has the information of a sample with independent observations of perhaps size three.

The effect of serial correlation can be seen in Figure 4 in a form more directly relevant to Amax's board. The implicit belief behind the question of how metals prices could be depressed for five years in a row is that, in a normal market, runs of such length should be very rare, as indeed they are for a nonstorable commodity. Figure 4 tests that belief. It asks, from a starting point in period 0 with a price of $100, the long-run average: What is the chance of a run of prices below $100 of various lengths over the next 10 periods. Not surprisingly, the chance of one poor period in the next 10 is high. Yet, the chance for a storable commodity having a depression of 5 periods' duration is surprisingly high, above 15%. Thus, one explanation of the sustained depression in metals prices in the mid 1980s is that such depressions are only to be expected for easily stored commodities such as metals.

Even with a much longer time series and the most modern methods, identifying the trend in a time series is no easy task. The point is well illustrated by a case which has attracted much attention and generated many studies over its long history. This is the statistical debate on the trend in the net barter terms of trade between primary commodities and manufactures (see Figure 5), initiated by Prebisch(1950) and Singer(19500), which acquired renewed urgency in the mid-1980's.

Despite a recent flurry of empirical studies on long term movements in commodity prices, the debate remains unresolved.. Spraos (1980) fitted a simple log-linear time trend variable to the data in a regression estimated via OLS, and found no trend in the postwar period. Sapsford (1985), interpreted the results of Spraos in the light of a possible "omitted" structural break in 1950. By introducing a dummy variable and correcting for serial correlation through the Cochrane-Orcutt technique, Sapsford was able to recover a negative trend in the net barter terms of trade on post-war data, too. Thirlwall and Bergevin (1985), using quarterly data for disaggregated commodity price indices on the postwar period, also fitted exponential time trend models, finding evidence of either constant or deteriorating terms of trade. Grilli and Yang (1988), constructed new price indices and estimated a simple time trend model (correcting for serial correlation). They found significant downward trends in the net barter terms of trade.

As Cuddington and Urzúa (1987, 1989) noted, all these studies (except Grilli and Yang) appear to have overlooked the importance of the serial correlation reflected in the price series. In the absence of any inspection of the statistical properties of the univariate representations of the series, all inferences that have been drawn are potentially subject to spurious regression problems. In a regression of a variable against a time trend and a constant, the distribution of the OLS estimator does not have finite moments and is not consistent if the error process is nonstationary (Plosser and Schwert (1978)), and tests of a time trend are biased towards finding one when none is present, if the disturbance is nonstationary (Nelson and Kang (1984)).

The problem appears thus to be the appropriate description of the error process and, therefore, of the series at hand. Cuddington and Urzúa (1987, 1989), following the identification approach suggested by Box and Jenkins (1976), find that the series they analyze appear to be nonstationary in the mean. In their study of the Grilli and Yang indices (deflating by the United Nations Manufacturing Unit Value) they reject the deterministic trend model in favor of a stochastic trend by testing the null hypothesis of non-stationarity in the price series using the tests proposed by Dickey and Fuller (1979) and Perron (1988). Excluding a one-time drop that they assume occurred in 1920, they conclude that no deterioration has occurred in the net barter terms of trade from 1900 to 1983. On the other hand Ardeni and Wright (1990) use the structural time series approach of Harvey (see for example Harvey and Todd (1983)) that requires no prior testing for stationarity. They find that the Grilli and Yang series appears to be trend stationary, with the trend declining at about -0.6 percent per year. If the 1920 dummy is arbitrarily added as in Cuddington and Urzua, the series remains trend stationary, but the magnitude of the annual trend falls to -0.14 percent. Therefore after more than fifty years of statistical measurement and research effort, economists can confidently say that the net barter terms of trade has stationary mean and negative trend. Or it has a unique and unexplained structural break in 1920 and is either nonstationary and trendless or it is stationary with weak trend. Then of course there are the cycles to consider. With this kind of advice, "adjusting for trend" in running a buffer stock for a sustained period would clearly be a piece of cake.

2. What can we expect from price band schemes?

Analysts of public programs and actual managers of those programs are in rare agreement when it comes to the design of a scheme to stabilize commodity prices. All put forward the standard of a price band scheme. For example, the analytical framework of the recent book by Ghosh, Gilbert, and Hughes Hallet (1987) is centered around price band schemes, from the econometric exercises to the policy simulations. A few of the more recent analytical examples of an analytical focus on price bands following in the footsteps of earlier writers such as Keynes (1974) are Gardner (1979), Hallwood (1979), Gardner (1982), Miranda and Helmberger (1988) and many others. In many actual international commodity agreements, the manager of the public stockpile is charged with keeping price within some band, with rules mandating accumulation of stocks at the bottom of the band and release at the top, often with some management discretion within an intermediate price range (Gilbert, forthcoming and Gardner, 1985). For example, the various International Cocoa Agreements have had a ceiling and a floor price symmetric about an "indicator" price, with this price band decomposed into a trigger range in which the buffer stock's manager can intervene at his discretion., and a nonintervention range. Many U. S. farm programs have had what amounts to a floor price and a much higher release price.

Price-band schemes have a superficial attractiveness and logic. Seemingly, the disruptions of price changes are reduced by efforts to keep prices in a narrow band. Seemingly, the symmetry of the band around the long-term mean price favors neither consumers nor producers and guarantees no great stock buildup. By the same implicit reasoning, the distribution of observed prices is close to symmetric, and accumulated net profits should hover around zero. On the other hand, intuition might suggest that supply response to the program may cause problems of excess stocks, so supply is best made unresponsive, if possible. Most obvious, the restriction on the release of public stocks to a price at lease equal to the top of the band seems a judicious and feasible storage policy.

None of these beliefs is valid in general. Most important, symmetric price-band schemes have an inherent tendency for a rapid and enormous accumulation of stocks, *unless* supply is responsive to price. Their effect on the distribution of price is not symmetric. Moreover, the requirement that public stocks be released only at the top of the band frequently leaves them in store when they would have higher social value if consumed immed ely. Thus, price-band schemes have substantial deadweight losses compared to other market-stabilizing schemes such as deficiency payments or price floors. Nor does it seem that a coalition of producers alone would prefer them.

Some of the analytical support for public storage under a price-band scheme stems from a failure to specify the alternatives to such a program. Many authors use price bands synonymously with buffer stocks, as if they suppose the only way to operate a buffer stock is with different floor and release prices. A buffer stock is more general and can be taken as synonymous with public storage, whatever the rule for public intervention.

Within the broader category of public buffer stocks, a price-band scheme involves two prices—P^F, the floor price at which the government is willing to buy any amount offered to it, and P^B, the minimum price at which the government will release anything from its buffer stock. Naturally, P^B is greater than or equal to P^F.

Conventional price-band schemes, with P^B and P^F symmetric around a plausible long-run price, have an intrinsic tendency to accumulate very large stocks. Indeed, a stochastic steady state may not exist; in expectation, accumulation of stocks may continue indefinitely. These properties are not the result of the interaction of private storage or production with the public policy. Nor are they the result of misidentifying trends in production or consumption. Rather, they result from the prescribed inflexibility of the buffer stock, which can only release its stocks at P^B or higher.

These general observations are best illustrated with a relatively simple example. Consider a specification of our model as described above where the consumption demand curve is linear, new production is perfectly inelastic with a mean of 100 units, the harvest is normally distributed with coefficient of variation= 10 units, and there are no trends to average yields or to demand. The long-run average price without storage is \$100. Also suppose only the government can store in this closed economy and that it uses a price-band scheme. Without elastic supply and private storage, any strange behavior must be attributed to the government's storage policy.

As the top of the price band, P^B, is raised for a given P^F, the average amount stored increases explosively. This feature of price-band schemes can be seen starkly in Figure 5, which which each P^B was simulated for 100,000 periods. With P^F set at \$80, a symmetric P^B is \$120. Yet, if P^B is set at even \$117,, average storage is enormous compared to the average under a simple floor scheme (see the observation for P^B = \$80). At the symmetric P^B of \$120, average storage in that particular run of 100,000 periods is close to 15 times average production.

Although the average storage in a simulation of 100,000 periods is lower for wide price bands, it is very large compared to the average storage under a floor-price scheme that stands ready to buy at the same P^F and sell at any price no less than P^F,as in Wright and Williams (1988). Periods without any storage still occur but become extremely rare, merely 0.14% in the case of the band scheme of 90-110% of P^N, the mean price.

It should also be emphasized that these values in Figure 6 for average storage under symmetric bands with no supply response are not steady-state values because a steady state does not exist. Simulations 10,000 periods long would show lower averages, while simulations 1,000,000 periods long substantially higher averages. Rather, the tendency is for continuous accumulation of stock and reducing of consumption below mean output., in contrast to simulations of a price floor below P^N, or of purely private storage, both of which have a stochastic steady state.

When planned production is elastic, its response to the negative effect of current accumulation on returns to output next period can put a bound on the expected accumulation as in the path for supply elasticity= 1.0, illustrated in Figure 7. Storage is expected to approach, after many periods, its steady-state mean of 31.4 units. Because a stochastic steady state exists, the long-run effects of a price-band scheme on price distributions, mean consumption, and producers' welfare can be studied. (Of course, the more elastic supply is, the more stable is the free market price and the weaker is the case for government stabilization.) Accordingly, in the remainder of this chapter only cases with elastic supply will be examined.

When we turn attention to the welfare significance of price-band schemes, the assumption of no private storage, convenient for demonstrating the tendency for large accumulations of stocks, must be dropped. No welfare analysis without private storage can purport to be accurate, for the welfare effects of purely private storage will be misattributed to the price-band scheme. Modeling of private storage, depending as it does on expectations of future behavior, requires stochastic dynamic programming.

More generally, when the public carryin is positive, private storage is distorted, sometimes upwards and sometimes downward, relative to the socially optimal level given the public storage behavior. This happens because a price-band scheme by definition imposes inflexible management on the public stockpile. Room is left in our example for flexible private storage. Consider public behavior in period t when the public carry in, S^g_{t-1} is 10 units and the new harvest is such that price is \$110. Because \$110 is just below the top of the band, \$112.50, none of the 10 units is released. Nevertheless, the price expected for the next period, t+2, is below \$110, \$102.29 to be exact. Any private stocks are therefore sold. Such conditions should also be a message to release some of the public stockpile immediately, because its current marginal value is higher than its expected marginal value the next period.[1] Because in this and similar instances the carryin of 10 units is not released (from the perspective of the preceding period, t-1), price in the current period, t, is higher than it would otherwise be. This higher expected price induces private storage in period t-1, despite the existence of a public carryout.[2]

Thus, this price-band scheme constrains the public stockpile to store even when the current marginal social value of its holdings is higher than the *undiscounted* expected future marginal value. A floor-price scheme with the same P^F never does this. Hence, a price-band scheme has a higher social deadweight loss. Figure 8 makes this clear. In that figure is plotted the deadweight loss as a function of P^B. The excess burden of a symmetric price band—$87.50 and $112.50—is some 17 times that of a price floor scheme at $87.50.

Moreover, the present value of public expenditures on the scheme increases considerably as P^B is raised, as can also bee seen in Figure 8. With the competition of private storage, the expected profits from public storage is at best zero under all circumstances. The rule that the public agency can release stocks only at P^B, above P^F, exacerbates the cost of interest payments and warehousing.

Figure 8 also shows the capitalized value of the change in the stream of net revenues to producers compared to solely private storage. This capitalized value, equivalent to "producer wealth" if producers are taken to own their land and the initial private stocks, is shown net of the present value of public expenditures on storage.[3] This information can answer whether producers would be willing to tax themselves (lump sum) to run a price-band scheme. The answer is yes, but a price-band scheme is only slightly preferable to a straight price-floor scheme. More surprising, the P^B that most favors producers is not at all close to the level symmetric with P^F. Such a symmetric scheme is usually what people have in mind when they recommend price-band schemes. Producers in this instance of linear demand and supply elasticity almost surely would prefer some scheme other than a price band—destruction of stocks, deficiency payments, a price floor—given that the government is prepared to spend some set amount (in present value). In as much as a linear demand curve makes stabilization especially attractive to producers, the conclusion appears inescapable that for less favorable demand curves price-band schemes are far from producers' first choice.

Of course, price-band schemes are rarely put forward with the explicit objective of increasing producers' wealth. Generally, the immediate objective is presented as a reduction in the variance of price. The way in which this benefits producers and/or consumers is not directly discussed.

Price-band schemes do reduce the variance of price. Nevertheless, that simple characterization misses the complex alteration in the probability distribution of price. The distribution of price for three cases are plotted in Figure 9. One is the distribution with no public intervention, to serve as a frame of reference. The important comparison is between the price floor scheme with $P^F = \$87.50$ and the price-band scheme with $P^F = \$87.50$ and the symmetric $P^B = \$112.50$. Although the price-band scheme reduces the percentage of prices above $112.50 from 14.9% to 6%, it primarily rearranges the distribution within the range $87.50 through $112.50. By far the most common price becomes P^B, where there is a mass point shown on the diagram. The frequency of P^F, in contrast, falls from that under the price floor scheme. Because of private storage, which coexists with a price below about $97, the distribution in the range between $ 87.50 and $112.50 is highly skewed.

3. Conclusion

The time series behavior of commodity prices has characteristics that can be largely explained by storage. The identification of trends in these series is a very difficult task, and the results are so uncertian as to make them useless for pricing policy.

Even if trends are somehow perfectly identified, the common prescription of a symmetric price band for stabilizing prices has effects on prices, producers, and the public budget not generally recognized. At least not until it is too late.

Footnotes

[1]The signal is similar with the true marginal social value rather than the expected price.

[2]The private storage industry's attention to these opportunities persists until expected profits are zero. From that fact it follows that, in expectation, the price-band scheme must run a deficit.

[3]The present value of the stream loss of consumer surplus can be inferred as the curve for the social deadweight loss plus the curve for producer wealth.

References

Ardeni, Pier Georgio and Brian D. Wright (1990), "The Prebisch-Singer Hypothesis: A Structural Time Series Analysis", mimeo, University of California, Berkeley.

Bosworth, Barry P. and Robert Z. Lawrence (1982), Commodity Prices and the New Inflation, Washington, D.C.: Brookings.

Box, G. E. P. and G. M. Jenkins (1976), Time Series Analysis, Forecasting and Control, San Francisco: Holden Day.

Cooper, Richard N. and Robert Z. Lawrence (1975), "The 1972-75 Commodity Boom", Brookings Papers on Economic Activity, 3, 671-715.

Cuddington, J. T. and C. M. Urzua (1989), "Trends and Cycles in the Net Barter Terms of Trade: A New Approach", Economic Journal, 99, 426-442.

Cuddington, J. T. and C. M. Urzua (1987), "Trends and Cycles in Primary Commodity Prices", mimeo, The World Bank, Washington.

Dickey, D. A. and W. A. Fuller (1979), "Distribution of the Estimators for Autoregressive Time Series with a Unit Root", Journal of the American Statistical Association, 76, 427-431.

Flavin, Marjorie A. (1983), "Excess Volatility in the Financial Markets: A Reassessment of the Empirical Evidence", Journal of Political Economy, 91, 929-956.

Gardner, Bruce L. (1985), "International Commodity Agreements", mimeo, The World Bank.

Gardner, Bruce L. (1979), Optimal Stockpiling of Grain, Lexington, Mass: Lexington Books.

Ghosh, S., Gilbert, C. L., and A. J. Hughes Hallett (1987), Stabilizing Speculative Commodity Markets, Oxford: Clarendon Press.

Gilbert, C. L. (forthcoming), "International Commodity Agreements: Design and Performance", World Development.

Grilli, E. R. and M. C. Yang (1988), "Primary Commodity Prices, Manufactured Goods Prices, and the Terms of Trade of Developing Countries: What the Long Run Shows", The World Bank Economic Review, 2, 1-47.

Hallwood, Paul (1977), "Interactions between Private Speculation and Buffer Stock Agencies in Commodity Stabilization", World Development, 5, 349-353.

Harvey, A. C. and P. H. J. Todd (1983), "Forecasting Economic Time Series with Structural and Box-Jenkins Models: A Case Study", Journal of Business and Economic Statistics, 4, 299-315.

Knudsen, Odin and John Nash (1987), "Lessons from Price Stabilization Schemes in Developing Countries", mimeo.

MacAvoy, Paul W. (1988), Explaining Metals Prices, Boston: Kluwer.

Miranda, Mario J. and Peter G. Helmberger (1988), "The Effects of Commodity Price Stabilization Programs", American Economic Review, 78, 46-58.

Nelson, C. R. and H. Kang (1984), "Pitfalls in the Use of Time as an Explanatory Variable in Regression", Journal of Business and Economic Statistics, 2, 73-82.

Newbery, David M. G. and Joseph E. Stiglitz (1981), The Theory of Commodity Price Stabilization: A Study in the Economics of Risk, Oxford: Clarendon Press.

Nurkse, Ragnar, "Trade Fluctuations and Buffer Policies of Low-Income Countries".

Perron, P. (1988), "Testing for a unit root in a time series with a changing mean", Manuscript, Princeton University.

Plosser, C. I. and G. W. Schwert (1978), "Money, Income and Sunspots: Measuring Economic Relationships and the Effects of Differencing", Journal of Monetary Economics, 4, 637-660.

Prebisch, R. (1950), The Economic Development of Latin America and its Principal Problems, United Nations, Lake Success.

Sapsford, D. (1985), "The Statistical Debate on the Net Barter Terms of Trade between Primary Commodities and Manufactures: A Comment and Some Additional Evidence", Economic Journal, 95, 781-788.

Singer, H. (1950), "The Distributions of Gains Between Investing and Borrowing Countries", American Economic Review, 40, 473-485.

Spraos, J. (1980), "The Statistical Debate on the Net Barter Terms of Trade Between Primary Commodities and Manufactures", Economic Journal, 90, 107-128.

Streeten, Paul (1986), "What Price Food? Agricultural Policies in Developing Countries", Economic Development Institute, The World Bank.

Thirlwall, A. P. and J. Bergevin (1985), "Trends, Cycles, and Asymmetries in the Terms of Trade of Primary Commodities from Developed and Less Developed Countries", World Development, 13, 805-817.

Williams, Jeffrey C. and Brian D. Wright (forthcoming), Storage and Commodity Markets, Cambridge: Cambridge University Press.

Wright, Brian D.and Jeffrey C.Williams (1988), "The Incidence of Market-Stabilizing Price Support Schemes", Economic Journal, 98, 1183-1198.

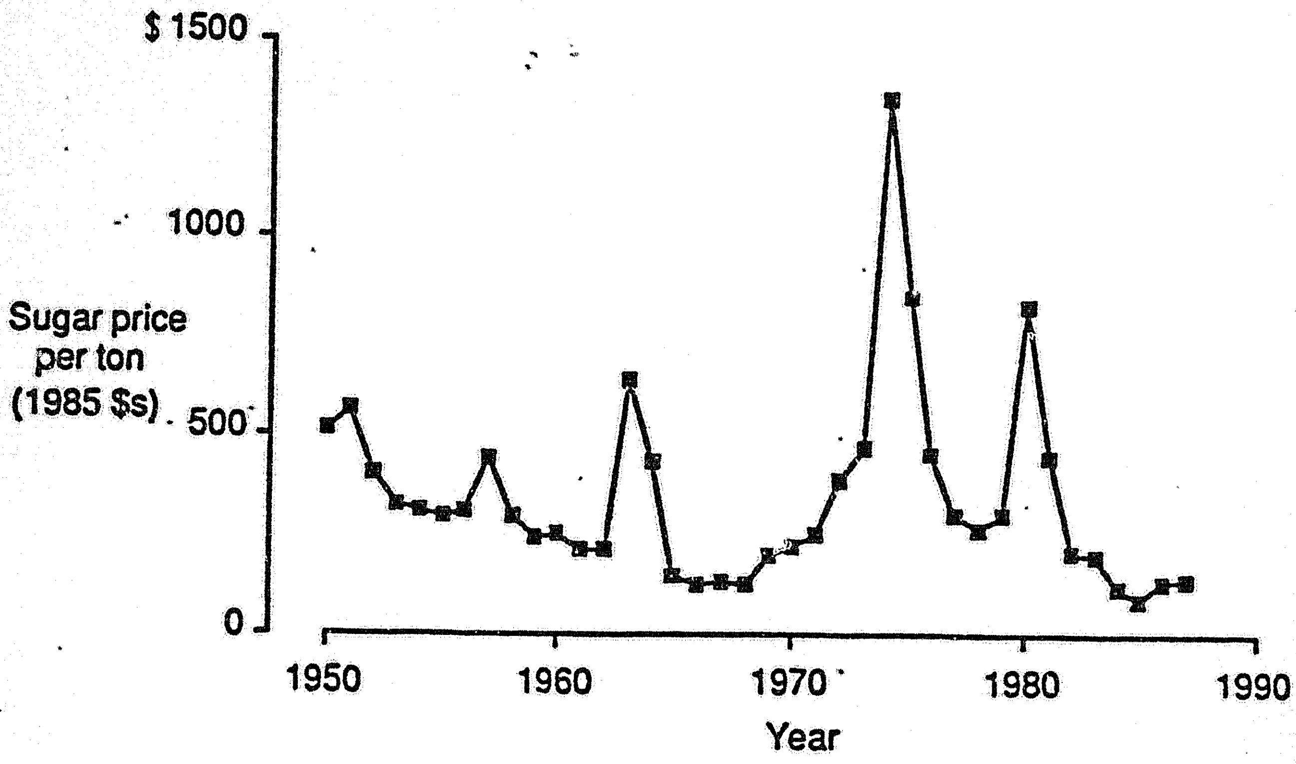

"World" price, f.o.b.s. Caribbean ports, deflated by U.S. GNP deflat
Source: World Bank (1989, p. 74, vol. 2)

Figure 1 Spot price of sugar, annually 1950-1987.

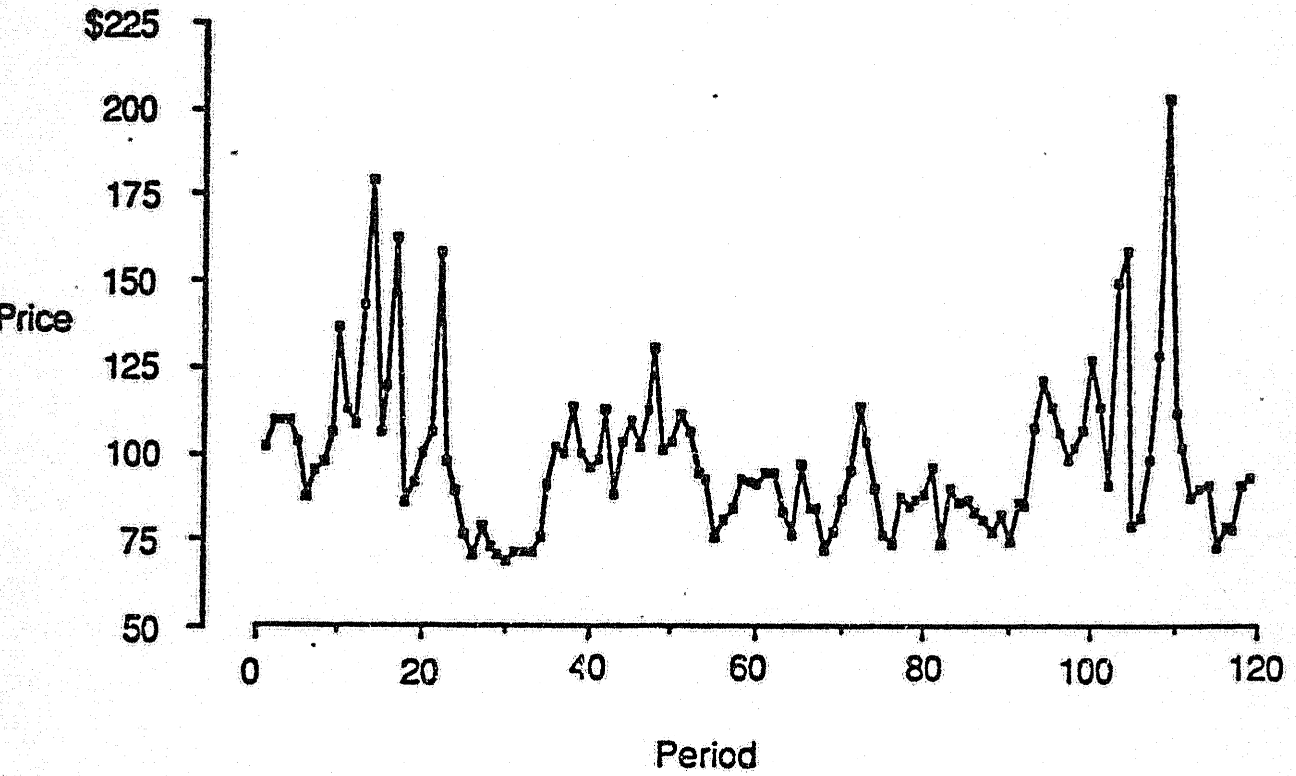

Figure 2 A time series of price

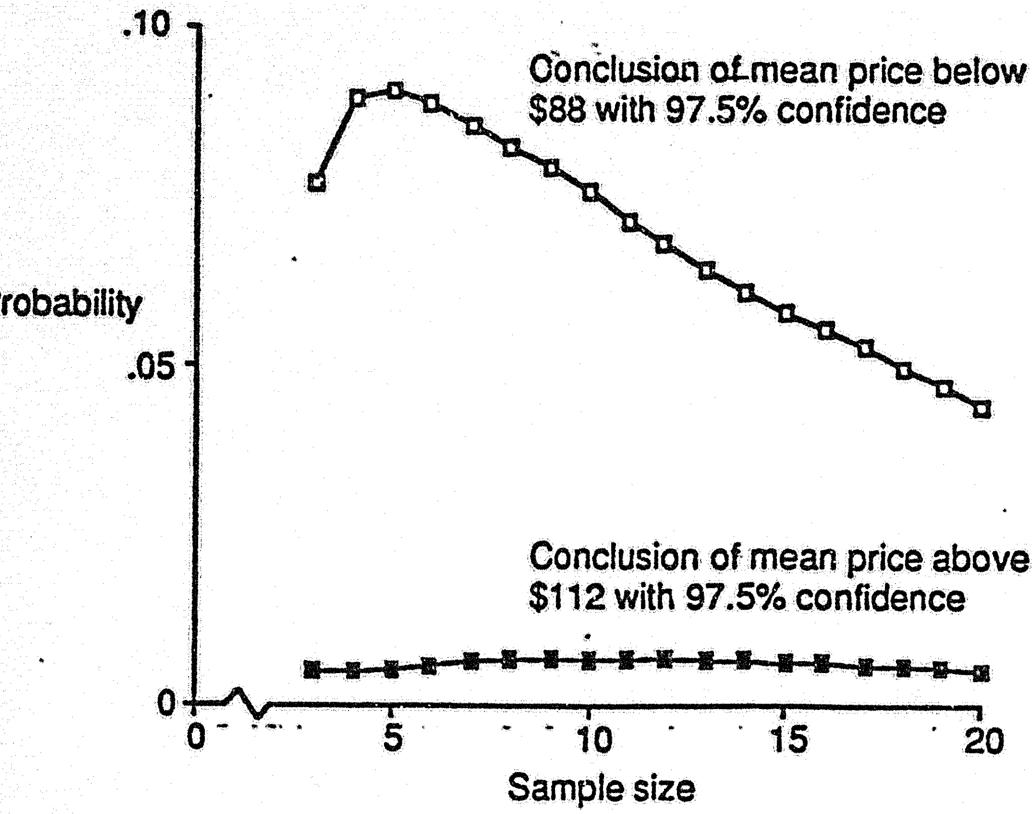

Figure 3 Effect of serial correlation on estimates of long-run mean price (true mean is $100)

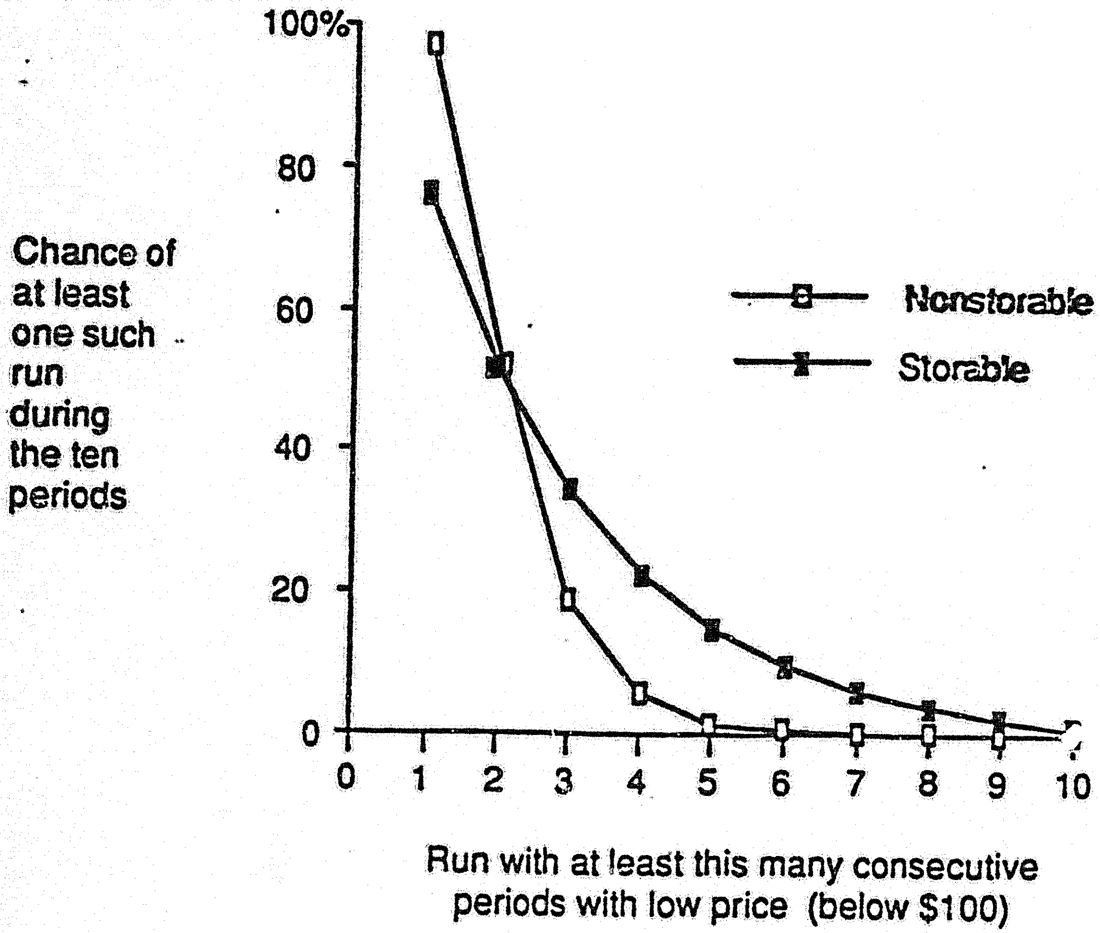

Figure 4 Tendency for sustained low prices

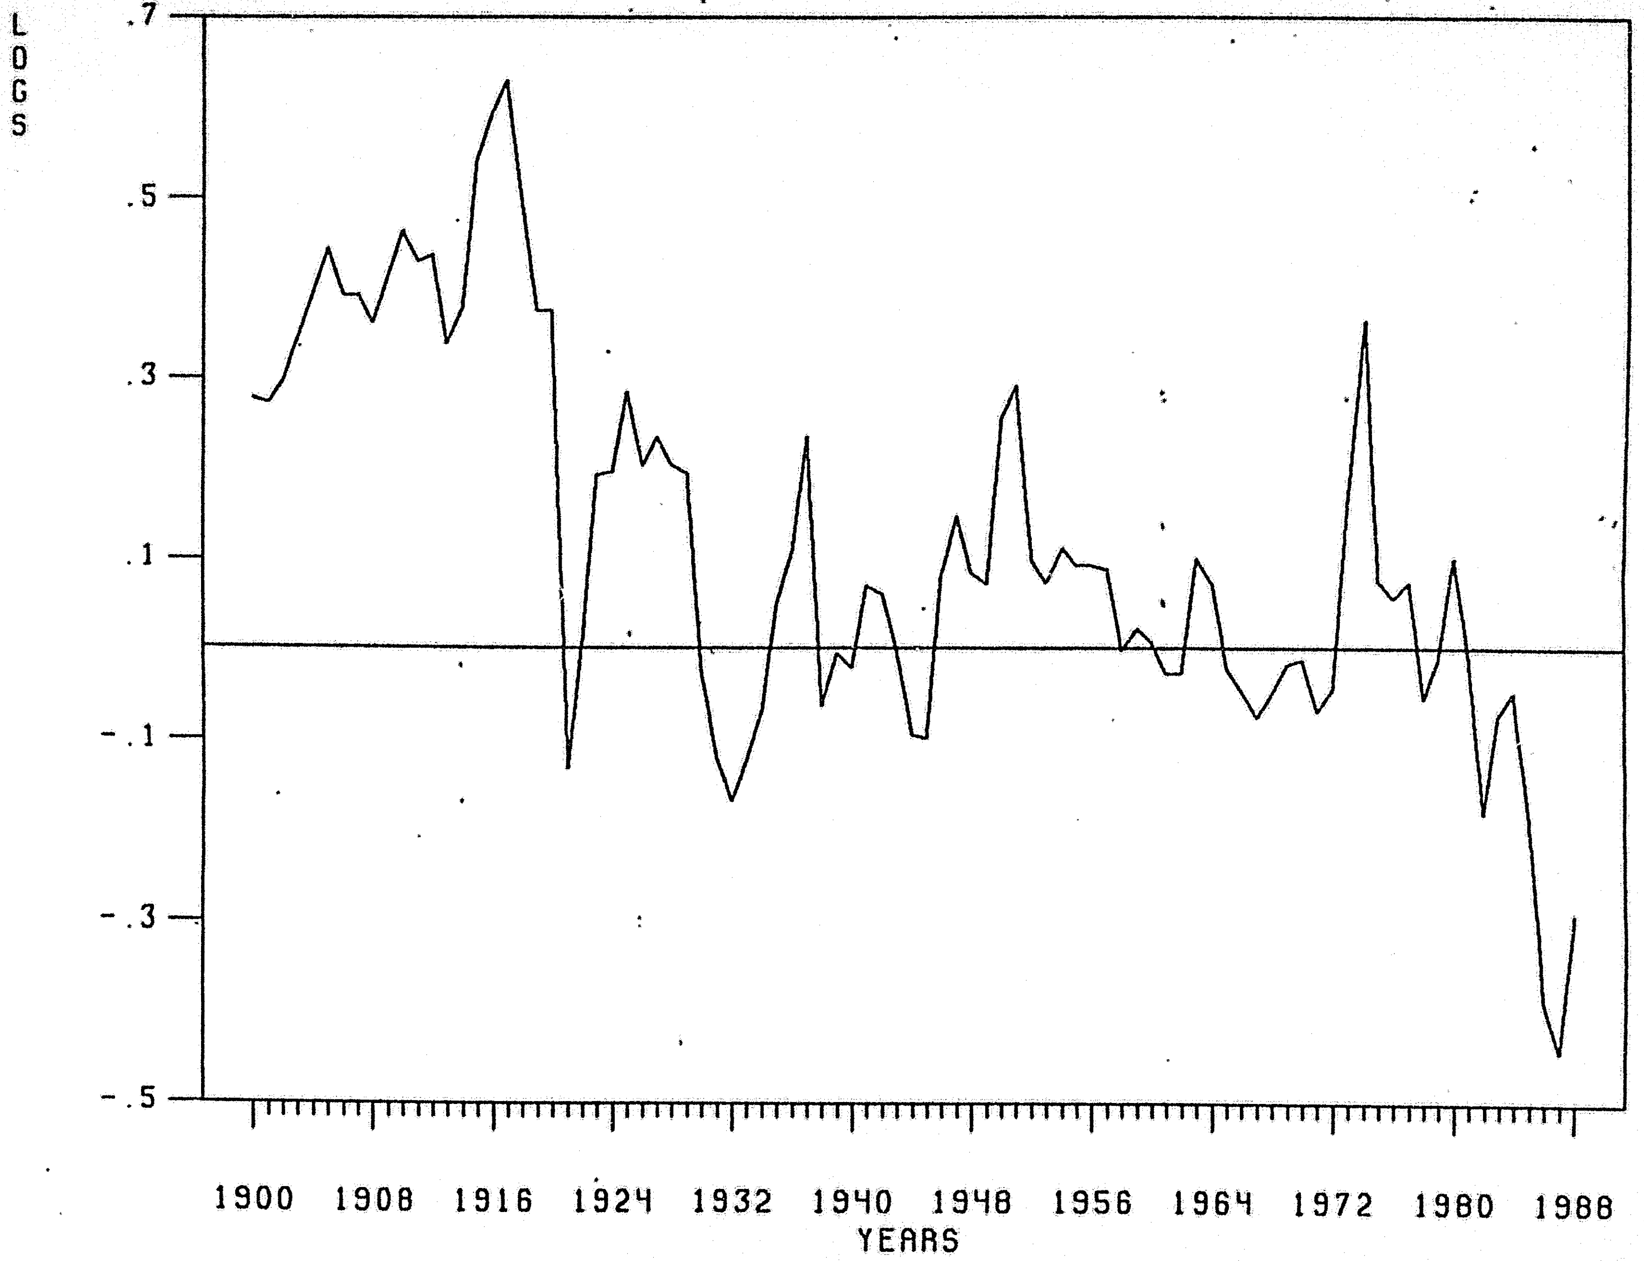
LOGS
.7
.5
.3
.1
-.1
-.3
-.5
1900
1908
1916
1924
1932
1940
1948
1956
1964
1972
1980
1988
YEARS

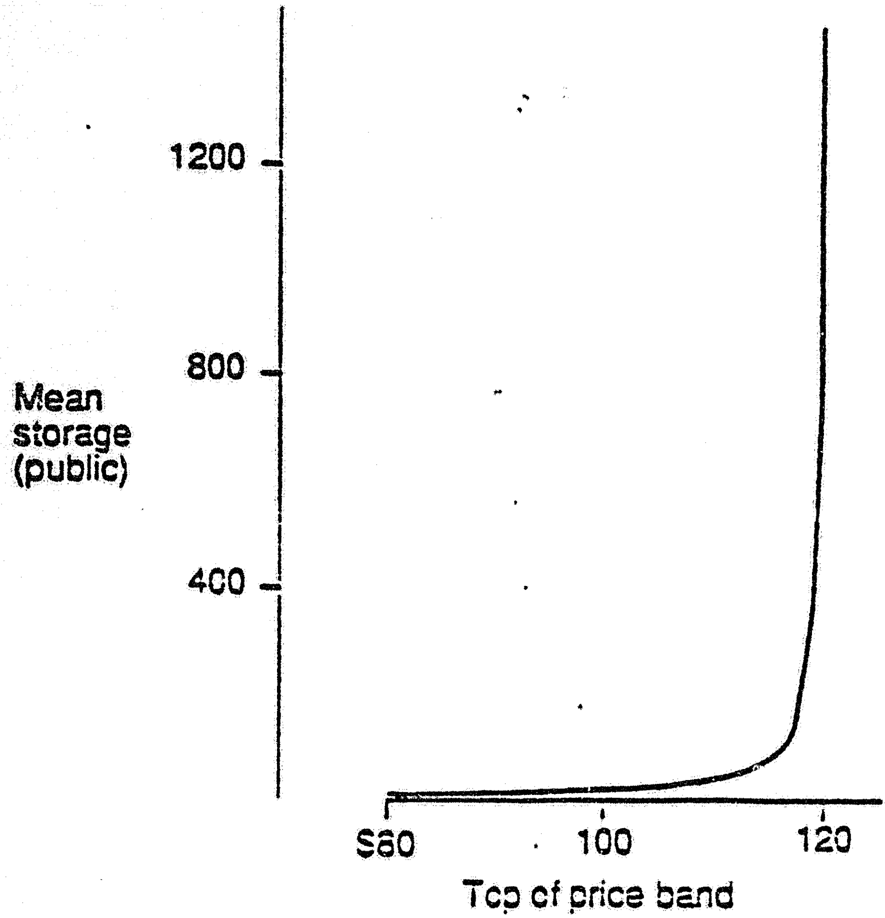

Figure 6 Explosion in quantity stored with higher top to price band

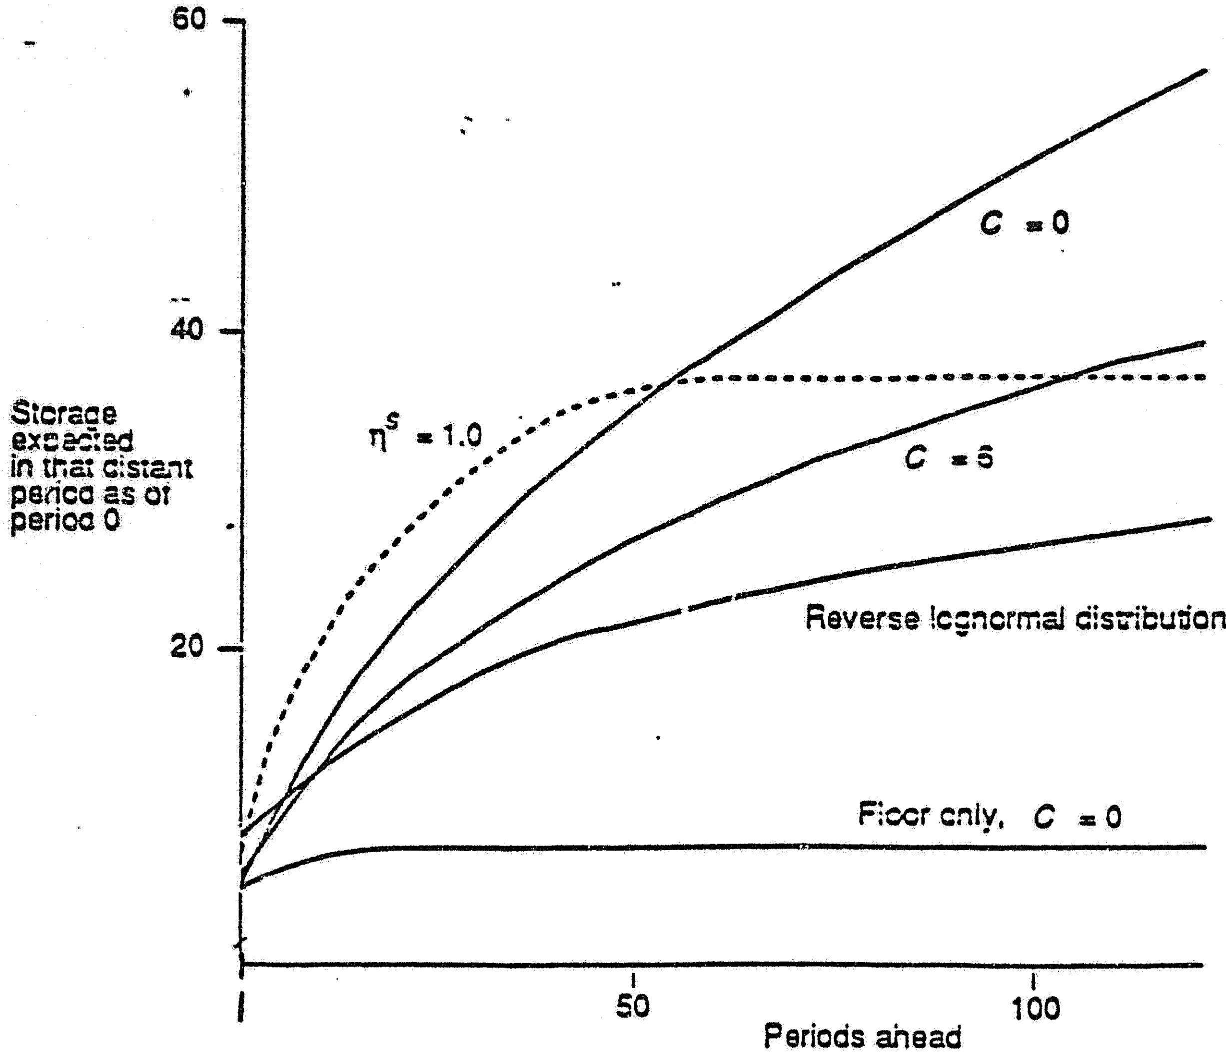

Figure 7 Expected time paths of storage

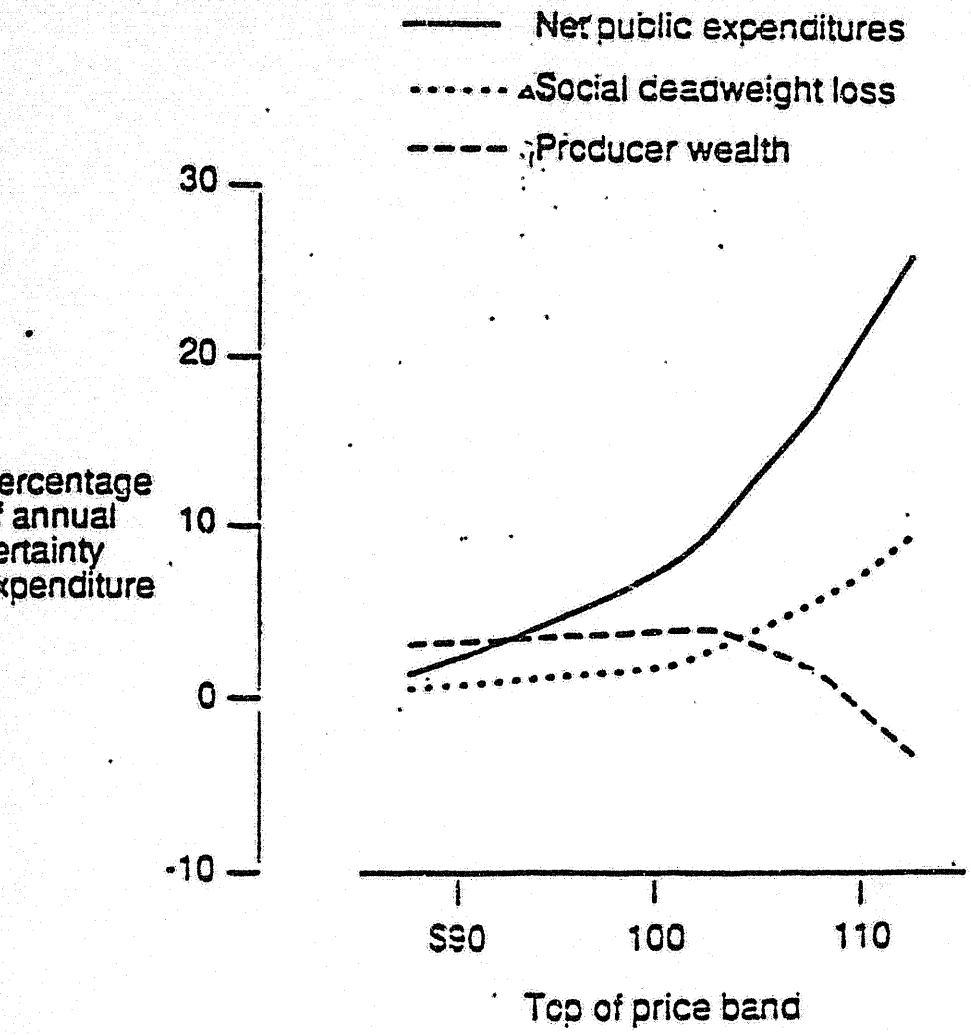

Figure 8 Welfare effects of a price band scheme

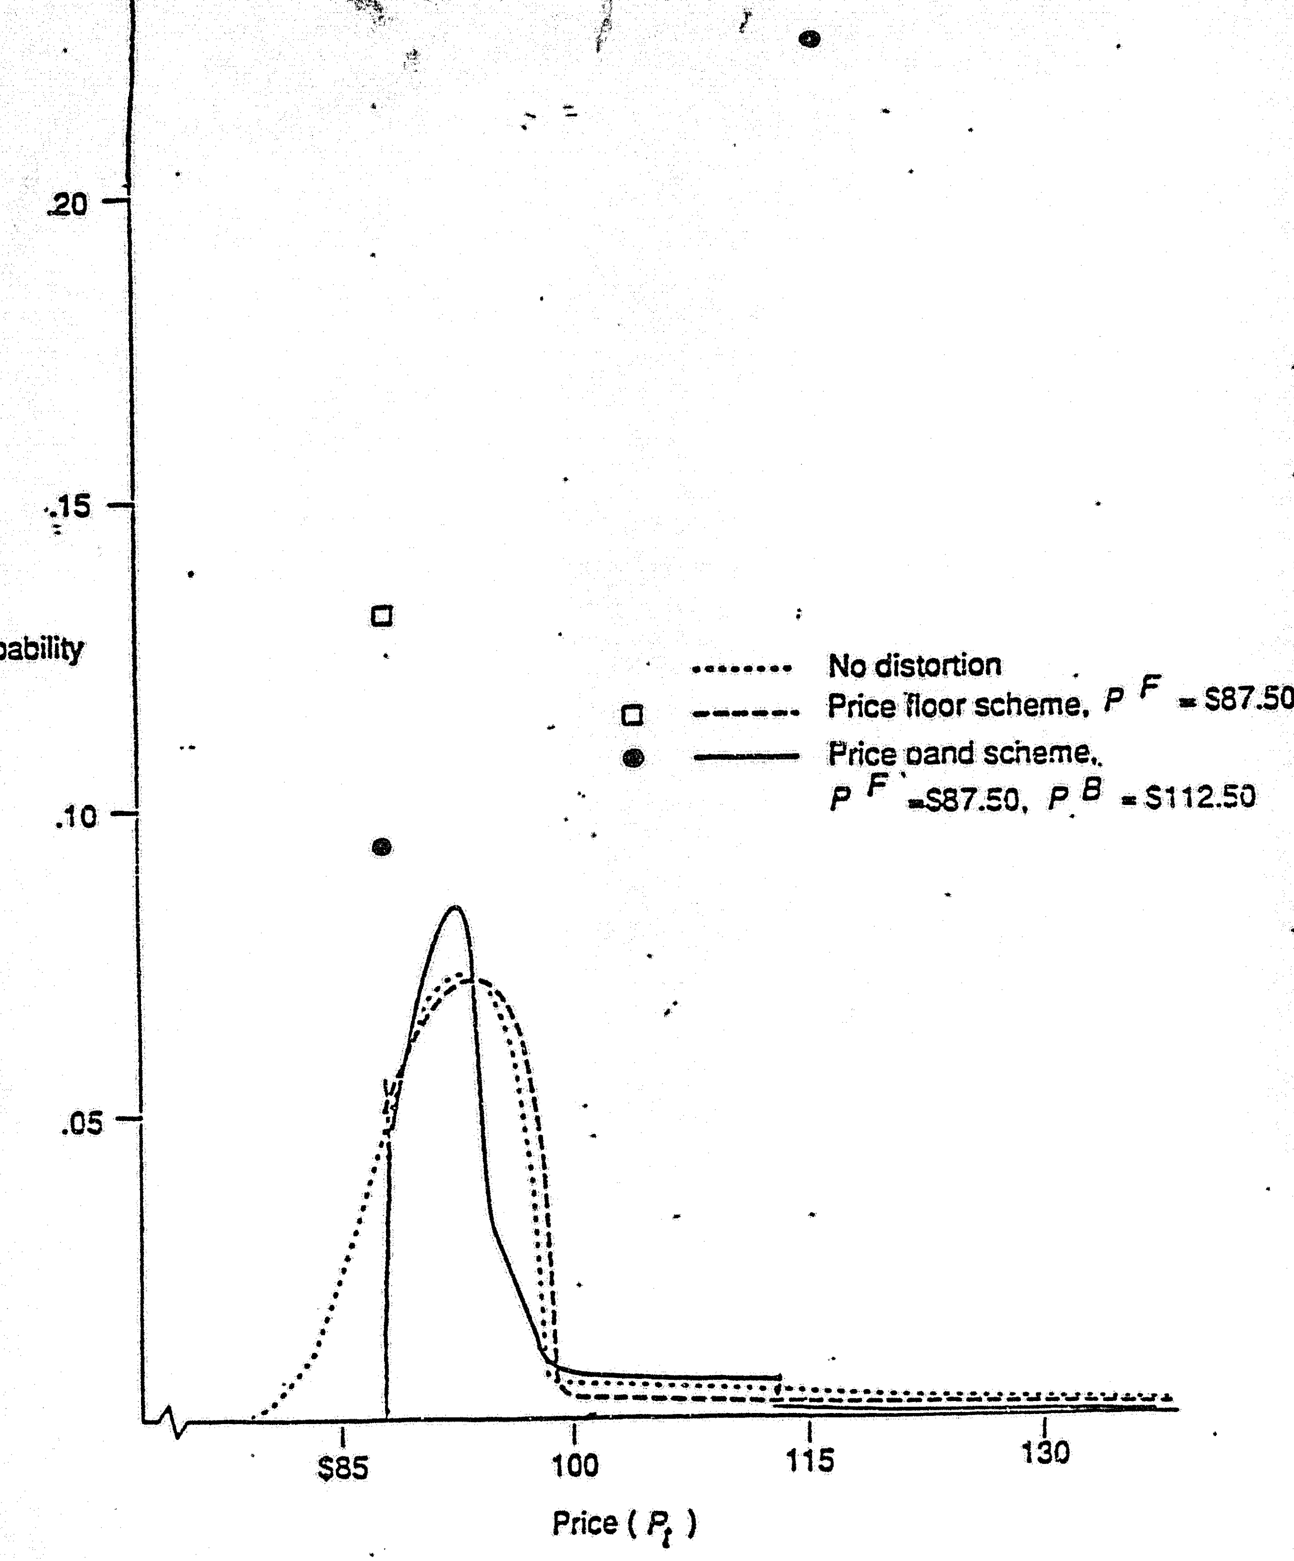

Figure 9

Price distributions