OPTIMAL USE
OF
CROP DISEASE-RESISTANCE GENES

by

DUNCAN KENNETH MACPHERSON

This thesis is submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Australian National University.

April 1977

Except where specific reference is made to the work of other people, the research reported in this thesis is original and was done without collaboration.

Duncan Macpherson

Duncan Macpherson

April 1977

"The use of composite wheat varieties comprising strains similar in appearance but different in resistance to stem rust, is one of the ways of counteracting the effects of different races of stem rust. This picture shows the process of preparing the composite seed." (Rockefeller Foundation (1966), p.41)

ABSTRACT

This study is a theoretical analysis of the problem posed to the breeders and users of disease resistant plants by the planting of large areas of monocultures and the depletion of natural gene pools. In particular, it is an analysis of the use of what are known as multiline crops (Jensen, 1952), crops which are agronomically homogeneous, but which are heteregeneous with respect to the disease reaction of individual plants. In the process of analysing the argument that multiline crops may have a use in combatting the evolution of plant pathogens, some light is shed on the roles of other modes of the use of genes for disease resistance.

A review of the history of breeding for disease resistance in the 20th century is given, which shows that although some disadvantages in the growing of large areas of monocultures had been observed by the mid 1930's, the first suggestions of means for counteracting these disadvantages were not made until the late 1940's. Among the earlier suggestions was the use of multilines. However it was soon pointed out that to plant crops which did not possess every available resistance gene in each plant was to invite avoidable losses and to hasten the appearance of "superraces" of pathogens able to attack all the available resistance genes. No conclusive answer to these objections has yet been made and, in consequence, there is an open debate on the use of disease resistance genes which centres on finding the answers to four questions: which genes should be used in which plants, and when, and where?

It is the contention of this study that the debate can be furthered by phrasing these questions as part of a problem in optimal control. A detailed description is given of the analogy between an

optimal control problem and the task of deciding on the best use of resistance genes. In the process, a number of biological factors which have been suggested as relevant to this task are discussed and it is shown that they can be integrated into the optimal control formulation of the debate. However, only a fraction of the information needed for a resolution of the debate is currently available.

The main analysis of this study is directed towards the question of whether there is any reason to suppose, on the basis of available information, that planting multiline crops may be the optimal way of using crop disease resistance genes. The technique adopted is to simulate the spread of disease in large-scale agriculture and to use an optimal control algorithm to discover what crop composition gives an optimal result. The simulation is based on current concepts of disease spread, and two alternative simple economies are modelled in order to allow a comparison of different judgements of optimality.

The main result of the study is that it is not self-contradictory to assert that multilines can be optimal, despite the objections to their use given above. However it appears that the beneficial effects to be gained from their use are dependent on the nature of the criterion of optimality used, and that more substantial benefits of the same kind can be achieved by crop rotations. Moreover, the uses for multilines discovered in this study are transitional: they are used for a few crop cycles only before a reversion to conventional practice occurs. On this evidence, multilines are not of practical use, but there are some indications in the results of effects that might make them so.

The main recommendations of this study are:

(1) that the next stage in theoretical research into multilines is to investigate the possibility that their sustained use can be optimal,

(2) that future numerical investigation of optimal gene use must be directed towards simplification of techniques, since the method used here is so time-consuming as not to be applicable in practical situations,

(3) that experiments which can be used to suggest whether multiline crops are an optimal mode of disease-resistance gene use are now possible.

Table of contents

Table of contents

Table of contents

Table of contents

Table of Contents

Table of Contents

Table of contents

ACKNOWLEDGEMENTS

I wish to thank the members of my guidance committee, Drs. I.R. Cowan, I. Cruickshank, and D.R. Marshall for discussions which helped me to avoid many pitfalls in the early stages of the project. I also had helpful discussions with Dr. C.E. Harwood, Dr. H. Rawson, and Prof. J. Warren-Wilson.

The manager of the A.N.U. Computer Centre, Mr. G. Blackeby, made the completion of the project possible by giving priority to the very time-consuming computation involved. Mr. I.R. Simpson of the Centre helped me untangle many baffling bugs in the programs.

The Head of the Department of Environmental Biology, Prof. R.O. Slatyer, served on my guidance committee and provided unfailing support in many difficult situations. My senior supervisor, Dr. B.R. Trenbath, provided close and thorough criticism at all stages, and was very patient with divergent thinking when convergent thinking was required, and vice versa.

I am grateful to the staff of the Department, especially Ms D. Lee and Mr. P. Firth for assisting with equipment and giving practical advice. Mrs. A. Clugston and Mrs. I. Brewer typed the thesis, despite the very inconvenient schedule involved in writing it.

Finally I would like to express my appreciation to my wife, who earned the greater part of our income and maintained a belief that the subject matter of the thesis was worthwhile.

Financial support was provided by a Commonwealth Postgraduate Scholarship, supplemented by the Australian National University.

1. Introduction

1.1. Aims and Layout of this Study

This study deals with the choice of adaptive uses of crop disease resistance genes in the face of the mutability and the powers of dispersal of major crop pathogens. In particular, it aims to analyse the case for what are called multiline crops (Jensen, 1952, 1965; Borlaug, 1953, 1958, 1965; Browning and Frey, 1969) as being such an adaptive use. In the process, some light is shed on other modes of use of resistance genes.

That a choice of modes is seen by plant breeders to exist today is a sign of a growing uncertainty about the best use of resistance genes and of an increasing knowledge of the genetics and epidemiology of crop pathogens. The growth of this uncertainty derives from increasing dissatisfaction with some of the consequences of conventional plant-breeding methods, and from the existence of an increasing number of suggested alternatives to these methods, none of which number seems clearly superior. In this chapter the causes of this dissatisfaction are reviewed in the perspective of the development of conventional plant breeding for disease resistance, and then the principal alternatives are set out, with special reference to the history of the multiline concept.

This historical treatment of breeding for disease resistance will involve some mention of genetic and epidemiological factors, but the main discussion of these will be reserved to Chapter 2, where their relevance to different modes of resistance gene use and especially multiline crops will be discussed critically. In Chapter 3 a unifying language will be proposed for the discussion and comparison of different modes, and the general method for making such comparisons described.

In Chapter 4 a theoretical approach to the application of this method will be set out and an implementation of this approach via simulation modelling of epidemics in multiline crops will be described. Chapter 5 discusses computational aspects of this implementation. Chapter 6 introduces a number of computer experiments that test hypotheses about the interactions of some of the factors discussed in Chapter 2 with different modes of gene use. Finally, Chapter 7 discusses the conclusions of the study.

1.2 There is Uncertainty about the Best Use of Resistance Genes

1.2.1. Conventional Plant Breeding for Disease Resistance

The purpose of this section is to review briefly the salient points in the development of breeding for disease resistance in this century. The limitation to this century is more than a convenience for although selection for resistant plants has always proceeded through natural evolution, and has doubtless been deliberately practised since the beginning of agriculture, the beginning of a scientific approach to plant breeding is often (e.g. Stakman and Christensen, 1960) set at 1900.

The landmark event seems to have been the application by Biffen (1905) of the rediscovered genetics of Mendel. Biffen found that some resistance to yellow rust in spring wheat was inherited via a single recessive gene. It now became possible not only to select resistant plants when they were revealed by epidemics but to plan breeding programs in which the parents of crosses and the necessary number of crosses to be sure of recovering the resistant progeny could be intelligently chosen.

When the further discovery was made by Stakman *et al.* (1918) (quoted in Stakman and Christensen, 1960) that natural populations of *Puccinia graminis tritici* were a mixture of physiologic races with virulences effective against different wheat varieties, two further steps became possible. First, many apparently anomalous results in the inheritance of resistance, in which the resistance of varieties changed from time to time and place to place, could be explained as being due to the use of impure isolates of the pathogen. This gave added confidence to breeders guided by Mendelian ideas (Stakman and Christensen, 1960). Second, the breakdown of resistance in cultivated varieties could be understood as the result of changes in the composition of the pathogen population. The task for the plant breeder with respect to disease resistance could now be formulated as being to make these changes harmless by incorporating more resistance genes in the genome of the crop.

This point of view on the plant breeder's task was well suited by the development of the backcross method in 1922 (Harlan and Pope, 1922, quoted in Briggs and Knowles, 1967). This technique is well adapted to the successive concentration of resistance genes in the one variety and so encourages the continued and widespread use of particular varieties that have proven agronomically successful. Thus the whole thrust of breeding for resistance came to be directed at producing single plants which were as far as possible agronomically ideal, while incorporating in them the best known sources of resistance. In pursuit of this ideal great efforts were made to seek out sources of resistance in older crop varieties throughout the world, and large programs to screen these collections were undertaken (Harlan and Martini, 1936).

The successes of this method were so marked, as in the quick recoveries that were made from the disastrous epidemics of wheat stem

rust in 1916, when Ceres wheat replaced Marquis, and 1935, when Thatcher replaced Ceres (Stakman and Harrar, 1957), that it became possible to think of breeding for resistance as approaching a final success, and this view was put forward by, among others, Biffen (1931) and Stanton *et al.* (1934). However, rather than achieving final success, resistance breeding has tended to repeat a cycle of resistance breakdown followed by introduction of a resistant variety, whose resistance in its turn is found to be only temporary, as when stem rust again became a major problem in North America in 1953 (Stakman and Christensen, 1960) as in the succession of new races of oat crown rust (Browning and Frey, 1969) and as in the way that wheat stem rust races effective against the resistance gene *Sr6* reappeared from apparent extinction in Australia with the reintroduction in the 1960's of wheat that used this gene (Watson and Luig, 1968). Although the cycle has moved at ever higher levels of scientific understanding and technique, it has not been halted, nor can it be safely said that its amplitude has decreased. The great epidemic of Southern leaf blight of corn in 1970 in the U.S.A. (Hooker, 1972), was made possible by the application of a technique (cytoplasmic male sterility) that could not have been explained in terms of the simple Mendelian ideas that Biffen used seventy years before.

1.2.2. The First Warning - Genetic Erosion

The cycle of resistance breeding described in the previous section might be regarded as being as acceptable as anything short of the final success once hoped for except for one consideration. The stock of resistance genes, though large, is surely finite. For a policy that depends on the periodic introduction of new genes, three ends are in view.

(1) The mutability of the pathogen is unable to cope with all possible resistance genes simultaneously, and so final success in breeding for resistance is possible.

(2) The stock of resistance genes is exhausted without achieving any permanent security for the crop.

(3) The rate of change of the pathogen population's composition is so great that plant breeders cannot keep up, producing a state more like (2) than (1).

Since the stock of potential breeding material includes not only old varieties predating scientific plant breeding but primitive crops and even wild relatives of crops, these three points may appear rather academic. What first gave them force was the observation, in the 1930's, that the new varieties produced by plant breeders were not only replacing old crops but also the primitive crops in the developing world. Thus the onset of condition (2) was being hastened, while no guarantee of achieving condition (1) could be given. According to Harlan (1975), the first to bring this erosion of genetic reserves to the public notice were Harlan and Martini (1936), who were already aware that, for instance, grain from California was being used to replace traditional varieties of barley in Africa in times of famine.

This homogenising of crops, although it gave cause for concern about the future of all areas of plant breeding, implies the need for genetic conservation, rather than the cessation of conventional plant breeding. Yet the conservation of wild plant populations, which is the only practical way of maintaining their gene pools (Frankel, 1974) is not only difficult, but, by an unpleasant irony, is bound also to conserve the pathogen populations from which races may emerge to challenge the cultivars of the developed world (Watson, 1970a). Thus conventional methods seem to procede irreversibly to an undesirable end.

However it is not clear either that this undesirable end is close enough to require immediate action, or that there exists a better means of breeding for resistance. With sensible conservation the available genetic resources may last until human food is being produced in some quite different manner. Even if genetic resources do not last so long, conventional breeding methods may be the best of a bad lot. One thing is clear. There will be much opposition if "sensible" conservation is allowed to interfere with the access of the developing world to new, high-yielding cultivars (Frankel, 1974).

1.2.3. The Second Warning - Large Epidemics

Not long after Biffen (1931) had expressed the hope of solving the major problems of breeding for disease resistance, and at roughly the same time as the concept of genetic erosion originated, a second pessimistic interpretation of the course of plant breeding began to be put forward. Stakman in 1938 (N.E. Stevens, 1942), writing in German, seems to have been the first to observe, and take as a danger sign, that the success of new varieties meant that the host environment of crop pathogens was more uniform than ever before. Thus the breakdown of resistance in a single variety meant a very rapid growth of epidemics. It also meant a very large loss. Although N.E. Stevens (1939, 1942) publicised these facts in English, he seems to have elicited little response at first. This may have been because the plant breeding community thought the dangers less pressing. Biffen (1931) was perfectly aware of the effect of large homogeneous crop areas on epidemics, but thought that the success of plant breeding would render this effect irrelevant.

The story of the response to the perceived threat of uniformity of crops is to a large extent the history of multiline crops, and so

will be dealt with in the next section. However the dangers of excessive uniformity in crops are now so well recognized as to amount almost to a cliche (Adams *et al.* 1971; Harlan, 1975; National Academy of Sciences, 1972; Day, 1973). Thus the first element of the uncertainty spoken of in section (1.1) is well established in the minds of the plant breeding community. Something is widely held to be wrong with the conventional (and still dominant) methods of breeding for disease resistance. The alternatives, and the debate over whether they are potentially superior, are the subject of the following sections.

1.2.4. The Responses to the Warnings

1.2.4.1. Types of Response: the Alternatives that have been Suggested

A recurring feature of the history of breeding for disease resistance is the slowness with which theoretical debates are resolved. In part this may be due to the complexity of the subject, to the time-lag involved in experiments and also, perhaps, to a suspicion of "theory" among many of those involved. Although, as has been shown, the elements of a major critique of conventional methods were complete by 1940, no alternatives seem to have been suggested for a decade, and no firm conclusions about the value of these alternatives have yet been reached.

The materials for a complete history of the alternatives do not seem to exist; it may be that the amount of thought devoted to many of them has been small. The approach adopted here will be to give, in this section, a list of the alternatives as they are seen now, and then, in section (1.2.4.2) to give a relatively detailed history of the multiline concept. This is the most complex idea to be discussed and it will be seen, in section (1.2.4.3), that it can be used as an

organizing concept for discussion of the other alternatives.

The presently known alternatives to conventional disease resistance breeding may be summarised as follows (International Atomic Energy Agency, 1971):

(1) The use of homogeneous crops in which each plant incorporates as many sources of resistance as possible (sometimes known as pyramidding of resistance genes).

(2) The use of susceptible varieties, though only in the absence of pathogens that can damage them severely.

(3) The use of so-called horizontally or generally resistant, or tolerant varieties. These terms have not been clearly defined, and will be discussed in Chapter 2.

(4) The use of different varieties in rotation.

(5) The use of different varieties in spatial patterns. Depending on the scale of the pattern this might be done at the level of:

(a) major geographical regions,

(b) farms or fields within farms,

(c) neighbouring plants - the multiline option.

1.2.4.2. The Multiline Option

The originator of the term "multiline" seems to have been Jensen (1952). His concept of a multiline crop was of a mixture of agronomically uniform plants, differing only in their response to the pathogen. Jensen believed, on "theoretical grounds" to be examined later, that

> A multiline crop would be expected...to possess the characteristics of longer varietal life, greater stability of production, broader adaptation to environment, and greater protection against disease.

Two elements seem to have combined to suggest the multiline idea to Jensen. The first was the idea of non-homogeneous crop mixtures as a means of buffering the yield of a crop against environmental fluctuations. In fact much of his paper is taken up with reviewing the performance of mixtures under conditions that were not explicitly stated to include disease. The validity of this idea of buffering, and the extent to which it applies to the behaviour of crop/pathogen systems will be taken up in section (2.2).

The second, not entirely distinct, element was the idea of "spreading of risk". In Jensen's view, the inhomogeneity of a multiline lowered the chance of a devastating epidemic. Although "small losses" would be expected to occur more often than in conventional crops, the susceptible components of the multiline could be removed from the next year's crop and perhaps reintroduced later. As will be seen in Chapter 2, the concept of the spreading of risk remains at the heart of the question of the best use of resistance genotypes, while the beneficial effects of inhomogeneity are by no means as clear as they seemed to Jensen.

Two sources seem to have been important in suggesting the application of risk spreading to the crop/pathogen situation. The first of these was the work of Rosen (1949). Rosen summarised in an abstract some possible ways of countering the problem that, at that time, oats resistant to race 45 of oat crown rust were susceptible to Helminthosporium blight and vice versa. As a short term solution he suggested the use of unselected progeny of a cross of parents from each resistance type, which would give a mixture, some of which would resist one disease and some the other. Jensen, while taking the point that disease losses could be lowered in this way, believed that it was important to produce an agronomically uniform crop, though he

later modified his views on this point (Jensen, 1966).

The second source of the risk-spreading idea seems to have been the work of R.B. Stevens (1949), who suggested that varieties which had become susceptible be reused after a long enough interval for the race that had attacked them to have disappeared. This suggestion of rotation as a means of disease control may deserve the credit for being the first constructive suggestion of an alternative to conventional methods of breeding for disease resistance.

The year after Jensen, Borlaug (1953) published an abstract in which he advocated a very similar program to Jensen's. He advocated the use of what he called "multilineal" cultivars, which would have the property of "remaining rust resistant indefinitely". Since this property was to be achieved by the successive replacement of susceptible varieties, Borlaug's multilineal varieties would be rather like Achilles' ship, which was still Achilles' ship after its hull planks, mast and rigging had all been replaced.

After this suggestion the pace of innovation slowed, and in the next mention of multilines (Borlaug, 1958) their production was still being spoken of in the future tense. However, in this paper two new and important ideas were introduced. The first idea was the criticism of multilines that remains the most serious reason for doubting their efficacy: that the epidemics that cause the "small losses" of Jensen (1952) are the best possible breeding ground for a "superrace" of the pathogen that can attack all components of the multiline, since they would greatly increase the opportunity for a mutation for increased virulence to occur, and for the mutant to disperse successfully to a previously immune host. The second idea was an attempted counter to this criticism: that the absence of supperace epidemics in nature implies that a "balanced" multiline system can be set up in which the

superrace is not fatally destructive. As will be seen in several parts of Chapter 2, the notion of 'balance' remains an important one in this context.

Although no theory existed to suggest how such a balanced system could be set up, during the 1960's several programs were started to produce multilines for practical use. According to Browning and Frey (1969) these were:

(1) The New York program, associated with N.F. Jensen. Little has been published about the results of this program, and specific mention of a multiline program was dropped from the group's annual report in 1970 (Cornell University, 1970).

(2) The Rockefeller Foundation program, associated with N.E. Borlaug, and concentrating on wheat varieties for Central and South America. One multiline, the 10-component Miramar 63 was introduced commercially in Colombia. Two years after its introduction, two of its components became susceptible and were replaced (Rockefeller Foundation, 1965). Although wide use of this variety was planned it was not achieved, presumably because of the advent of the short-strawed wheats. The heir to the Rockefeller Foundation program, the Centro Internacional de Mejoramento de Maiz y Trigo, began in 1968 to develop a modern wheat multiline based on its own cross number 8168. By 1973, 285 possible components of an anticipated 15-20 component, not entirely agronomically uniform, multiline were being tested in 31 countries (CIMMYT, 1974). The reports of these tests are general and emphasize the good agronomic performance of the components rather

than their behaviour as a mixture or their effects on the pathogen population.

(3) The Iowa program, associated with J.A. Browning and K.J. Frey. This program has produced some commercial oat cultivars for use in the southern U.S.A.

In each case the areas planted have been too small to constitute a direct test of the superiority of multilines in retarding pathogen evolution. The chief result has been that even when races virulent on a fraction of the multiline are present, the loss of yield is less than that fraction of the disease free yield. The need to tolerate even these losses, when a pure line of a fully resistant plant might be planted, has not yet been demonstrated, and, by the same token, the principles of constructing a multiline are not clear. In what seems to be the only published account of the method of making up a multiline, Frey *et al.* (1971) give their rule as being to ensure that not more than 40% of the multiline is fully susceptible to any of the pathogen races present at the time, but the derivation of this number is not given. (See also the frontispiece).

Thus it was only after the decision was made by several people to produce multilines that experimental work was done to investigate the basis of the assumption of multiline superiority. The most notable epidemiological work has been that of Clifford (1968), Leonard (1969a, b, c), and Cournoyer (1970), who have observed the dynamics of epidemics in multiline plots of oats, and Sumner and Littrell (1974) who have followed epidemics in experimental plots of corn multilines. Work at the epidemiological level is mentioned here simply to show how wide the gap still is between promise and proof in regard to the effect that multilines would have if used on a large scale. The direction in which the results of epidemiological

and other work point will be discussed in Chapter 2.

Browning and Frey (1969) nevertheless sum up the promise of multilines in terms of extending "indefinitely" the "useful life" of resistance genes, "removing the rust hazard" and avoiding the risk of "homogenizing the pathogen population on a global scale". Given the current state of knowledge, this is an optimistic assessment. The present position of multilines is an unresolved standoff between optimism and caution.

1.2.4.3. The Other Alternatives

The purpose of this section is to review briefly the other alternatives listed in section (1.2.4.1) in order to show that some of them need not be considered any further separately from multilines and that the others may share certain attributes with multilines.

From the definition of multilines given in the previous section, together with the reasonable assumption that the same multiline need not suit all situations, it is plain that the various uses of different crop varieties in spatial and temporal patterns are special cases of the use of multilines. For example, if, over a wide area such as the wheat belt of N.S.W., adjacent farmers used different varieties, this would be the special case of multiline use in which one component of the multiline dominated in each farm, but different multilines of this type were used in each farm. Conventional methods, of course, are a still more special case. Thus all of these patterned methods will share in varying degrees any advantages of "spreading the risk" and any disadvantages of acting as a breeding ground for the superrace.

The situation with regard to the use of "generalised resistance" or "tolerant" varieties is more complicated because there exist conflicting definitions of these concepts and because some of the debate turns

on technical genetic questions, which will be discussed in Chapter 2. Despite these difficulties a preliminary account can be given which indicates the issues involved and brings out another important idea which has influenced thought about plant breeding.

The earliest use of "tolerance" in this context seems to be that of Caldwell *et al.* (1958), for whom a variety was tolerant if it could "endure severe attack ... without sustaining severe loss". The issue was somewhat confused by van der Plank (1960), who contrasted tolerant varieties to susceptible varieties as those which, while not fully resistant, delayed the spread of an epidemic. He suggested that the widespread use of such varieties might effectively delay epidemics out of existence. To make matters worse, Browning and Frey (1969) quote van der Plank (1960), but use tolerance in a sense close to that of Caldwell *et al.* (1958). They suggest that multilines should show "synthetic" tolerance, presumably because many of each generation of pathogen propagules will meet nonsusceptible tissue.

Tolerance is also said to be significant by Browning and Frey because overcoming tolerance should "provide no mechanism for new races to prevail over established races". This appears to be true for tolerance in the sense of Caldwell *et al.* (1958), because decreasing the yield of a plant is of no direct significance to the pathogen's survival and reproduction if it does not simultaneously increase the infection level. However it is not true for tolerance in the sense of van der Plank, because a decrease in the delaying effect of a plant on a pathogen race constitutes a selective advantage for that race. Indeed the advantage of tolerance (*sensu* Caldwell *et al.*) may be only verbal, because mutations for increased vigour in the pathogen can confer selective advantage and decrease yield without affecting "tolerance" at all.

A third reason for considering tolerance to be significant is its allegedly polygenic character (Simons, 1972). In this it is joined by "generalised resistance", with which it is sometimes identified (Watson, 1970b), though generalized resistance is also defined as resistance which is effective to some extent against all races (Caldwell, 1968) or as resistance which does not involve the hypersensitive reaction in which the invading pathogen is killed (Watson, 1970b). Whatever form it takes, the significance of polygenic resistance is either that it is by definition less vulnerable to single mutations in the pathogen, or that the combination of genes will "place before the fungus barriers which are more difficult to negotiate." (Watson, 1970b). Generalised resistance is also sometimes known as "horizontal resistance" (van der Plank, 1963). Browning and Frey (1969) also claim that multilines should have "synthetic" horizontal resistance, which though not complete (since some components are vulnerable) is effective against all races (except, presumably, a superrace).

While there is general agreement that tolerance and generalised resistance are desirable properties for crop plants, no specific proposals for their use have yet been suggested. The important idea that appears to be contained in the proposals featuring tolerance and generalised resistance is that immunity is probably unachievable in the long run and that the task facing plant breeders is going to be to minimise the effects of unavoidable disease. The way this minimisation is to be carried out is not yet clear.

1.3 Summary

All the alternative ways of breeding for disease resistance reduce to giving answers to three questions:

(1) Which sources of resistance should be used in which plants?

(2) When should these plants be used?

(3) Where should these plants be used?

The conventional answer to this question has been, in essence: one new major gene for resistance bred into one plant to be used in all times and places. As a result of mixed success with this approach, uncertainty about its correctness began to be expressed in the 1930's. This was the first part of the uncertainty mentioned in section (1.1).

A number of different approaches have since been suggested. The earlier ones concentrated on varying the answers to questions (2) and (3). With increasing understanding of the genetics of crop/pathogen interactions, it has become possible to imagine varying the answer to question (1) as well, by using polygenic as well as monogenic resistance. However, no consensus exists on the best approach or combination of approaches and so the second part of the uncertainty mentioned in section (1.1) still persists.

2. Factors Affecting the Mode of Use of Resistance Genes

2.1. Introduction

The purpose of this chapter is to examine the nature and significance of factors that may affect the mode of use of resistance genes, and that, in particular, may justify the use of multilines. The factors that have been proposed as important in this context form only a small and apparently miscellaneous subset of the factors currently studied under the headings of disease resistance and crop epidemiology and the behaviour of plants in mixtures. The two following sections discuss only as much of the general background of these areas as is necessary to evaluate the factors in question. In section (2.2) the behaviour of mixtures in the absence of pathogens is discussed. Section (2.3) deals with pathogen-dependent factors, proceeding from a definition of resistance to a consideration of factors at the genetic, physiological and epidemiological levels. Section (2.4) summarises the conclusions about the significance of the various factors.

2.2. Non-Pathogen Factors

2.2.1. Nature of the Factors

As was mentioned in section (1.2.4.2), the original impulse towards the use of multilines was partly that advantages had been claimed for the use of mixtures of varieties in the absence of pathogens. These postulated advantages are of two kinds: a greater average yield for the mixture than for its highest-yielding component, or a greater stability of yield for the mixture than for its most stable component. Jensen (1952) recognised that these advantages, if they exist would apply to a diseased mixture, but he also drew an analogy between the

greater stability of yield in a mixture and the avoidance of crop failure, or "spreading of risk", that would occur if the multiline were so composed that all of its components were never vulnerable to the pathogen races present at one time.

Simmonds (1962) reviewed the experimental evidence for the existence of these advantages of mixtures, and found that while mixture yield was occasionally higher than the pure stand yield of the highest yielding component, the stability advantage was more common. A possible explanation for increased stability is that if the components of the mixture react differently to the random variation of the environment, then only part of the mixture will be adversely affected at any time, and the chance of the whole mixture yielding below its mean is correspondingly reduced. An increase in yield, on the other hand, seems to require some kind of mutually beneficial interaction between the components. Such cooperation is apparently less often found than simple differences between mixture components (Trenbath, 1974), and when it does occur it is usually unexplained (Simmonds, 1962). One example which is interesting because the mechanism is partly understood is the case of heterosis in sugar beet, where Curtis and Hornsey (1972) have been able to suggest that a mixture may be so chosen as to maximise the amount of outcrossing in the crop and consequently maximise the relief of inbreeding depression.

Marshall and Brown (1973) have derived a criterion for the stability of mixture yield to exceed the stability of the most stable line in the case where stability is measured by variance of yield and the mean yield of the crop is given by

$$Y_m = \sum_i \lambda_i \, (Y_i + I_i) \qquad \ldots (2.2.1)$$

where the λ_i are the proportions of the components of the mixture

the Y_i are the component yields in pure culture

the I_i are the weighted interaction effects

The condition for a mixture to be more stable than the most stable component in an equal mixture of n components is

$$C_Y + C_I < (nV_{min} - V_I - V_Y)/(n-1) \qquad \ldots (2.2.2)$$

where n is the number of components in the mixture

V_{min} is the variance of the yield of the most stable component in pure culture

C_Y is the mean covariance of the yields of the components

C_I is the mean covariance of the interaction effects

V_Y is the mean variance of the yields

V_I is the mean variance of the interaction effects

Clearly, this relation is more likely to hold the more negative the covariances are, that is, the more unlike the responses of the components to the environment are, or, in Jensen's terms, the more the risk is spread.

2.2.2. Significance of the Factors

Although mixtures which yield better than their best component are unusual, they may have some significance for the use of multilines as a means of disease control. The example quoted above from Curtis and Hornsey serves as a reminder that the components in an outbreeding mixture may differ in their characteristics as pathogen hosts from the same components planted in pure culture. Experimental evidence is needed to show whether the effect is positive or negative: heterosis might result in a plant whose greater vigour enabled it to resist the debilitating effect of the pathogen, or in a plant whose rapid metabolism made it an ideal source of metabolites for the pathogen.

Of more general possible significance are the effects of mixtures on the stability of yield, and it is these that will be concentrated on here. In particular, the identification of 'stability' with low variance of yield facilitates a precise discussion. The model of Marshall and Brown (1973) assumes that the source of variation behaves randomly and that the yield has no systematic trend with time. The general notion of 'spreading of risk', though not as precise in its expression, seems to embody the same assumptions.

However the presence of a pathogen can often invalidate both assumptions if the host and pathogen populations are not in some stable state or limit cycle. It is particularly clear from the events described in Chapter 1 that in the case of extensive monocultures that the evolution of the pathogen is, as Johnson (1961) put it, "man-guided" and tends always to produce a decrease in yield. It may sometimes be the case, for example in the kind of traditional mixed agriculture in India described by Aiyer (1949), that host and pathogen interact stably. To prove that they do would be difficult because it would be necessary to model the behaviour of the host and the pathogen and their interaction with the rest of the mixed-agriculture system and then show that (1) the model was stable to the perturbations produced by mutations in the pathogen, and (2) the model was sufficiently like the real system for stability of the model to imply stability of the system. Even if they are not strictly stable such systems may be effectively stable if the probability of the pathogen dispersing successively to a plant of the correct species is so low that even the superrace cannot reproduce fast and selection for virulence has little effect on yield. However in the cases of interest to most of the world today neither stability nor effective stability prevails.

Thus the behaviour of a crop/pathogen system, unlike that of a crop/environment system, consists of two parts: a random variation and a systematic interdependence of crop and pathogen. Where the first dominates, the spreading-of-risk conception of the mixture is appropriate because Marshall and Brown (1973)-type assumptions hold, and the plant breeder can trade off some of the average yield attainable in exchange for a greater stability of yield, without intensifying the disease problem. Where the systematic component dominates, as in the closer approaches to complete genetic uniformity, spreading-of-risk is not an appropriate description, because the crop affects the pathogen population, and the focus of the plant breeder may shift from minimising the variance of yield to maximising the absolute value of yield, which, under the influence of the pathogen, tends downwards.

Therefore the behaviour of multilines in the absence of a pathogen or in the case where the multiline does not noticeably influence the evolution of the general pathogen population is only a partial guide to their usefulness as a means of disease control. Uncertainty remains, largely because of the interaction of the crop with the pathogen population. While it is unclear whether crop/pathogen systems can (or should) be stabilised (more will be said about this later in the chapter), it is clear that the analogy between the response of mixtures to random environmental factors and their response to pathogens should be used with great caution.

2.3. Pathogen-Dependent Factors

2.3.1. Introduction

The long term results of any mode of resistance gene use will depend on the way resistance is inherited in the crop, the way virulence is inherited in the pathogen, the way the individual plants grow when diseased, and the higher-order interactions which may be summed up under the heading of the dynamics of epidemics. There are several ways of classifying the various factors which have been said to be relevant to the choice of a mode of gene use. None of these classifications gives complete and mutually exclusive subsets of factors which can be discussed separately without overlap in answering the questions posed in Chapter 1: which genes in which plants, when and where. The following sections are organised in a progression from the genetic to the physiological to the epidemiological level. This corresponds roughly to a progression from host/pathogen interactions of a low order to those of a high order, from hostXpathogen, as in the infection of a single plant, to hostXhostXpathogenXpathogen, as in the competition of plants infected by different pathogens. The increase of complexity involved is particularly evident in the higher orders of interaction.

As a starting point for the discussion it is useful to have a definition of resistance. Wood (1967) has defined resistance as the "properties of a plant which reduce damage caused by a pathogen". This is a wide definition if damage is defined widely, and could conceivably extend to "damage" done to the genetic diversity of a crop. It also emphasises that resistance is best quantified by a vector of values rather than by a scalar. These values represent components of resistance, which are often measurable separately. For example the germination rate of pathogen on host and the rate of production of propagules per unit infected biomass are measures of components of resistance, though

if any one component is large enough it may be difficult to measure the others in practice. From an epidemiological point of view all these properties can sometimes be summed up in a single parameter, the relative rate of increase of the pathogen (van der Plank, 1963). This is a short-hand description that inevitably involves a loss of information about the host pathogen interactions. Though the relative rate of increase may be the most convenient measure of the interaction between single pairs of hosts and pathogens, it can be expected to vary with the composition of a multiline, and so cannot be thought of as a constant property of a multiline/pathogen interaction.

2.3.2. The Genetic Level

2.3.2.1. The Nature of the Factors

The definition of resistance given above already indicates the possibility of separate genetic control of different components of resistance. The initial observation of Biffen (1905) was of a single recessive gene which conferred such a high level of some component of resistance that the overall growth of the pathogen was almost completely stopped, and the question of separate effects on separate components became academic. A majority of the sources of resistance discovered since have been, like the first, oligogenic (Wheeler, 1975), meaning that one or a few so-called "major" genes confer a noticeable increment in some component of resistance. Person and Sidhu (1971) reviewed 912 reports of resistant reactions whose genetic basis had been investigated and found that in over 95% of cases the genes reported were oligogenic. An unknown proportion of the remainder that reported polygenic resistance were held by Person and Sidhu to be suspect on technical grounds. A common pattern among "major" resistance genes is that they condition a reaction called the hypersensitive reaction (Wood, 1967), in which the invading propagule is killed in its attempt to germinate and

consequently the spread of disease is entirely stopped. Because of the prevalence of this type of resistance it has generally not seemed important until recently to enquire what component of resistance is directly affected by a gene.

This preponderance of reports of oligogenic resistance may reflect the fact that it is easier to search for the causes of discrete differences in resistance than to attempt to assay resistance quantitatively, as a variable affected by many genes with a cumulative effect. However it is becoming clear that at least some resistance is inherited quantitatively, under the control of a number of genes (Simons, 1975; Luke *et al.*, 1975), though in some cases the number of genes involved is small (Luke *et al.*, 1975). At the same time, physiological studies have found individual genes that have discrete, but small, effects on components of resistance (Slesinski and Ellingboe, 1969; Stuckey *et al.*, 1974) and that could form part of a polygenically inherited pattern of resistance. When the effects of individual genes must be combined by the breeder in order to produce a high level of resistance in the plant, it becomes more important to enquire what component of resistance is being affected, and this importance is reflected in current attitudes in breeding.

The resistance of the plant to the pathogen depends also on the genetic constitution of the pathogen. The most significant attempt to systematize the genetics of interactions for resistance and susceptibility between hosts and pathogens has been the gene-for-gene hypothesis of Flor (1955). As stated by Flor (1971) it is that "for each gene that conditions resistance in the host there is a corresponding gene in the parasite that conditions pathogenicity". If this is true, a susceptible reaction is produced only when a gene for virulence is present in the

pathogen and the corresponding gene for susceptibility is present in the host.

As Wheeler (1975) points out, it is not clear whether the gene-for-gene hypothesis can be disproved, for if a particular interaction does not seem to follow the rule, then the discrepancy can be explained by the presence of other, non-complementary genes in the interaction. However, there are many interactions that do follow the gene-for-gene pattern (Flor, 1971) and this suggests that many pathogens and their hosts have evolved together over a long period of time. Since unchallenged resistance exterminates the pathogen and unchecked aggressiveness in the pathogen tends to eliminate the host, this long period of coevolution suggests that perhaps the host and the pathogen formed a system that was stable before the advent of modern agriculture. Mode (1958) showed that it was possible to stabilise a model of a gene-for-gene system by appropriate choice of selection coefficients, and suggested that stability could be maintained over long periods by a system of balanced polymorphisms with suppression of crossing-over.

2.3.2.2. Significance of the Factors

At the level of genetic manipulation, neglecting for the moment the way in which resistance genes express themselves, these patterns of inheritance raise two questions for the plant breeder:

(1) Should he concentrate on oligogenic or polygenic resistance? In particular, are multilines unnecessary if polygenic resistance is available?

(2) Do the known patterns of resistance justify fears of genetic erosion? In particular, can induced mutations remove the criticism of conventional plant breeding that it may exhaust the stock of resistance genes?

(1) The importance of using polygenically based resistance is sometimes argued (Robinson, 1968) on the grounds that oligogenic resistance is more easily overcome by mutations in the pathogen race. If mutation rates are equal at all loci in the pathogen, and the gene-for-gene hypothesis holds, this advantage of polygenic resistance would hold because more mutations would be required to overcome the larger number of genes. However if the crop genes in question operated by giving additive doses of some component of resistance, successive mutations in the pathogen would result in an accelerating return to susceptibility as the increasing pathogen population had more opportunities to mutate, and the advantage would be temporary. A case of this kind has been reported (Simons, 1972). It is a process that resembles the step-wise mutations for virulence that might be expected in multiline epidemics and which are the main disadvantage for which multilines have been criticised (section 1.2.4.2). This process of successive mutation might be arrested if the polygenic grouping was in some way cooperative so that all genes had to be made vulnerable before susceptibility appeared. The simplest case of this occurs when a number of single genes, each conditioning a hypersensitive reaction and so a high level of resistance are simultaneously incorporated in the host. Other mechanisms for cooperation do not seem to have been suggested. It ought not to be overlooked, however, that there are cases known in which a single major gene has conditioned resistance for decades without becoming susceptible to a new race (Lupton, 1972).

The imperfect correlation between the number of genes involved and the endurance of resistance, together with the blurring of the distinction between polygenically and oligogenically determined resistance tend to suggest that polygenic resistance should not be treated differently in principle from resistance conditioned by major genes. Macer (1972) discusses some of the practical problems involved

in handling plants with polygenic resistance. In the first place, more elaborate selection processes are needed to detect plants with relatively low levels of resistance and whose progeny segregate in other than a simple Mendelian manner. Also, the polygenic character is masked if there is a single major gene in the plant, so that major genes can only be incorporated at the later stages of a breeding program.

However there can be no question of breeding for one type of resistance rather than the other. A balanced position (Simmonds, 1962; Watson, 1970a) cannot afford to ignore either, even though the details of how best to partition effort between different techniques are at present unclear. In particular, the choice of multilines need not neglect polygenic resistance. There is no theoretical reason why multilines should not employ polygenic resistance, either uniformly in the genetic background of all lines, or distributed nonuniformly by the same rules that are used to distribute the major genes.

(2) The rate of erosion of genetic resources is also affected by awareness of patterns of inheritance because this awareness affects the use that is made of genetic material. In the first place, the selective transfer of major genes from gene pools into crops without their accompanying polygenes (the "Vertifolia effect" of van der Plank (1968)) must mean that polygenic resistance is being removed from available genetic resources faster than major gene resistance. But on the positive side there is the possibility that induced mutations may help to make good the losses.

Although there has been much experimentation with mutagenesis (International Atomic Energy Agency, 1971; Roane, 1973), and some genes for resistance not known from natural sources have been found (Jorgensen, 1971), there is as yet little sign that practical

contributions have been made to breeding for disease resistance (Smith, 1971; Williams, 1975). It also appears that such contributions as are made are likely to be made via single major genes whose transformations can be more easily observed in screening the large number of usually deleterious changes that are produced by mutagens. Consequently, it "is most improbable that the balanced gene complexes which are the consequences of long term selection for adaptation can be assembled by mutation breeding techniques" (Frankel and Bennett, 1970).

Thus although it is difficult to say whether we are in danger of exhausting the supply of resistance genes (Day, 1974), it does not seem that increasing genetic knowledge has created the possibility of either a static equilibrium or a dynamic equilibrium in which genes can be replaced as fast as they become vulnerable over an indefinitely long time. The adverse consequences of conventional breeding techniques discussed in sections (1.2.2) and (1.2.3) must still be regarded as real threats. It remains necessary to consider alternative methods.

2.3.3. The Physiological Level

2.3.3.1. The Nature of the Factors

At the physiological level the genetic basis of resistance is expressed in many ways, the details of which are not all of interest in the present context. In most cases (Wheeler, 1975) the links between the genes involved and the physiology of the disease reaction and between the disease reaction and the success or failure of the invading pathogen are not known. Hence some of the ideas to be discussed here are not, strictly speaking, physiological concepts but serve to organise the phenomena of resistance where the physiological correlates of resistance are obscure.

Two important organising ideas are the concepts of general (versus specific) and induced (versus preformed) resistance. Some resistance is believed to be of a passive preformed nature in which the ordinary functioning of the plant reduces damage. This is a wide category that includes such things as a life cycle that avoids the most favourable period for pathogen development (Shaner *et al.*, 1975), the possession of a thick cuticle (Day, 1974), and a generally vigorous metabolism. However, most attention at present is focussed on induced resistance in which the growth of the pathogen is retarded or stopped by a reaction which does not take place until the pathogen attempts infection. Current theories suggest (Day, 1974, Chapter 5) that this type of resistance is usually specific and is the type of resistance governed by 'gene-for-gene' interactions. The role of the gene for resistance in the plant is to enable the recognition of some characteristic byproduct of the avirulent pathogen and initiate the production of relatively nonspecific anti-fungal agents, which may be the substances known as phytoalexins (Kuc, 1972).

Less is known at the physiological level about the forms of resistance in which specific recognition does not appear to play a part. Perhaps as a consequence the terminology of 'general resistance' is the most confused of any area relating to disease resistance. Terms used include general resistance, slow rusting, field resistance, mature plant resistance, durable resistance and tolerance. Schafer (1971), reviewing tolerance, gives eight slightly divergent definitions that have been used. Much of the confusion seems to stem from a failure to specify closely enough what measurements would define a value of the type of resistance in question.

For instance, tolerance is included here as a form of resistance because all the definitions of tolerance center on the idea of a lesser

yield reduction for a given infection level, and because yield reduction is reasonably counted as damage to the plant. However reference has already been made in section (1.2.4.3) to the confusion that has resulted from the failure to specify at what time the infection level is to be measured. For example, if infection level at time t_1 affects infection level at time t_2, which in turn determines the yield reduction, then "slow rusting" (Luke *et al.*, 1972) from t_1 to t_2 produces "tolerance" of the infection level at t_1, without reference to any tolerance of infection level at t_2. On the other hand, if the yield reduction depended on infection level at t_1 as well, then we would probably not wish to speak about 'tolerance' but about 'general resistance'. In a model in which the dependence of yield on infection levels at different times was explicit, questions about 'tolerance' or 'general resistance' would be replaced by questions about particular model parameters, and no ambiguity would arise.

Terms such as 'tolerance' and 'mature plant resistance' do not seem to cause serious ambiguity in practice. What is not clear, despite the wide use of the term, is the meaning of the term 'horizontal resistance'. Van der Plank (1963), the originator of the term gave two separate definitions:

> (1) "horizontal resistance reduces r" (r is the relative infection rate) (1963, p.120)
>
> (2) "when the resistance is evenly spread against all races of the pathogen we shall call it horizontal" (1963, p.174)

and he expanded on these by saying "horizontal resistance has been called field resistance ... generalised resistance ... and other names" (1963, p.120). The contradiction in these definitions can be seen in the attempt to classify the effect of a single gene that gives a degree of protection against one race of a pathogen. Such a gene (some of Flor's (1955) genes come under this heading) is apparently to be

considered 'vertical' because it affects only one race. Yet vertical resistance "does not reduce r, because reproduction is normal in races not governed by the resistance" (1963, p.120). In this system of definitions resistance is not considered to reduce r unless it reduces r for all pathogen races. Thus all resistance is defined as horizontal until it becomes vulnerable to a single race, after which it becomes vertical, and perfect resistance is vertical unless it is effective against all races. The confusion seems to result from a failure to distinguish between the mean level of resistance to various races and the variance of resistance to various races, since horizontal resistance is in one case defined by its absolute level (reduction of r) and in another by its variance (evenness with respect to races). The adoption of this system of definitions has led not only to advocacy of the use of "horizontal" resistance (usually supposed to be polygenic) with varying degress of enthusiasm (Watson, 1970b; Nelson, 1972; Robinson, 1968), but also to such statements as

> By definition, any given resistance mechanism must be either horizontal or vertical; it cannot be both. (Robinson, 1968)

2.3.3.2. Significance of the Factors

The significance of the physiology of specific host/pathogen reactions for someone seeking to answer the three questions of section (1.3) lies in the fact that induced resistance is a common phenomenon, and that it can be effective against pathogens other than the avirulent spore that initiated it (Day, 1974). Thus in a number of cases the reaction to a avirulent spore results in a temporary resistance to a normally virulent spore that chances to fall near the same place though it would normally not be "recognised". This effect is known as the cross-protection effect, and Johnson and Allen (1975) have proposed that it may be significant in multilines, where the more susceptible

plants would be a source of cross-protecting spores. The effect is known in wheat (Cheung and Barber, 1972), flax (Littlefield, 1969) and oats (Kochman and Brown, 1975), and although there is some controversy about its mechanism (Kochman and Brown, 1975), it does not appear to be a laboratory artefact. The size of the effect that might be expected in a multiline under field conditions does not seem to have been estimated. Presumably the number of cross-protecting spores would have to be comparable to that of virulent spores, which suggests that high proportions of the more susceptible plants would have to be grown. There is also the possibility of induced susceptibility occurring (Brown, 1975), but this is relatively rare and should not come as a surprise to the breeders of a multiline, since it too is reproducible experimentally.

The significance of the various posited types of general resistance is less clear, partly because they themselves are not clearly defined. The key is probably the same as that given in section (2.3.2.2) for the significance of polygenic resistance, with which general resistance is often identified (Simons, 1975), correctly or incorrectly. Thus general resistance is thought of as resistance which is less vulnerable to mutation than specific resistance and whose conquest, since it is usually not perfect resistance, gives less of an advantage to the successful race. Luke *et al.* (1972) give an example of a quantitatively inherited form of resistance that has apparently provided general resistance for many years. However examples of general (in the sense of enduring (Caldwell, 1968)) resistance governed by single genes are also known (Simons, 1972).

If general resistance is defined only by the fact that no race has yet been found that decreases its effect, then it would be foolish to base a breeding program on it. Moreover the correlation between

polygenes and durability of resistance which might provide a basis for predicting 'general' resistance is imperfect. In the absence of other explanations of why some resistance is durable and other resistance is not, it would seem that general resistance is not a very useful idea. At the same time, a history of durable resistance should not be ignored. Within a model of gene use and its effects, general resistance can perhaps be consistently modelled as resistance whose complementary races (in the sense of races that can attack its genetic base in some way) have low probabilities of arising. It is an interesting question how such probabilities could be estimated in practice.

If the "general" property of general resistance can be modelled in this fashion, and its specific manifestations can be modelled as indicated in the discussion of tolerance, the choice of different genes and their use can be resolved as a quantitative rather than a qualitative question. The beginning of such a quantitative treatment will be found in the following chapters. What will not be found is an attempt to model "horizontal resistance", since the inconsistencies in its definitions make it impossible to model.

2.3.4. The Epidemiological Level

2.3.4.1. Nature of the Factors

The factors relevant at the epidemiological level can be classified under seven headings, corresponding roughly to an increasing order of host pathogen interaction:

(1) the probability that new, virulent races of the pathogen will appear

(2) the probability that a pathogen propagule will find a susceptible host

(3) the relative survival rates of different pathogen races

(4) the interaction of diverse plants subject to one pathogen

(5) the interaction of diverse pathogens on the same host

(6) the interaction of diverse hosts and diverse pathogens

(7) other high order interactions.

(1) The probability that a new, more virulent race of a pathogen appears at a given time has two components. The first is the probability that the new race may arrive from another area, where it already exists. It is believed, for example, that some new races of *Puccinia graminis tritici* appearing in Australia may have entered from Africa (Luig and Watson, 1970). The second component is the probability that the new race arises by a change in one of the races already in the area under consideration. There exist a number of mechanisms by which this may take place, both sexual and asexual. Reviews are given in Watson (1970b) and Webster (1974). The basis for changes in virulence is mutation, but other processes such as heterocaryosis and hybridization can lead to the effective expression of the new gene. From the point of view of the plant breeder seeking to define his expectation of the pathogen's evolution, all these processes appear in the form of probabilities of occurrence dependent on the size of the pathogen population.

(2) The probability that a pathogen propagule will find a susceptible host is a function of the proportion of susceptible plants within the dispersal range of the pathogen. It will evidently decline as the proportion of the susceptible host in the crop declines, both because the hosts are, on average, father apart and because other plants intervene. Even if the whole crop is susceptible, the probability will not be unity, because many propagules will not land on any kind of plant. Even apart from this mortality, the probability will not in

originally supported by a number of field reports of the survival of simple races in mixed populations, notably of *Phytophthora infestans*. Since the original introduction of the concept of stabilising selection, some more systematic collections of field reports have been made (Watson, 1970b; Luig and Watson, 1970) and a number of experiments have been carried out to measure the progress of selection in mixed pathogen populations under more controlled conditions (Leonard, 1969b; Brown and Sharp, 1970; Martens, 1973; Volin and Sharp, 1973). Work in this field has been reviewed recently by Brown (1975).

While some of the work that has been done has produced evidence that stabilising selection occurs in particular cases (Leonard, 1969b), the consensus from both field and laboratory seems to be that stabilising selection is not a universal or even a general phenomenon (Brown, 1975). Thus Watson (1970b) reported that one of the two most prevalent strains of *Puccinia graminis tritici* in Australia carried genes controlling virulence on two resistance genes currently unused, and Martens (1973) found that a simple race of oat stem rust (*Puccinia graminis avenae*) that showed superior competitive ability in the growth cabinet was inferior in the field.

Van der Plank (1968) attempted to adjust the theory of stabilising selection to account for these exceptions by introducing the concept of the strength of resistance genes, defined "in terms of the strength with which stabilising selection acts against the gene's complementary race". This modification has been strongly criticised (Nelson, 1972) as depriving the hypothesis that simplicity correlates with vigour of any predictive power. In fact some genes which van der Plank designated as strong have since been shown to be weak (Brown, 1975). More than this, however, the terminology of 'weak' and 'strong' genes is self contradictory, because there is no unique "complementary

race". The same gene in different genetic backgrounds has different complementary races and so may be rated strong and weak at the same time.

Whatever the resolution of the debate about stabilising selection the underlying problem of accounting for the scarcity of superraces still remains. One possible solution might lie in a reassessment of the possible rate of evolution of wild pathogen populations. It may be that only in the last century or so has there been any significant probability of the superraces that arose on a particular local variety of a crop making a successful transfer to a crop in another region. Obviously, evidence on this point will be very hard to obtain.

(4) A common phenomenon in crop mixtures of plants of uneven competitive ability is that the inferior competitor grows less well in the mixture than in a pure culture, while the superior competitor grows correspondingly better (de Wit, 1960). The presence of different levels of resistance in a heterogeneous crop during an epidemic tends to ensure that the competitive position of heavily attacked plants will deteriorate so that compensatory growth of the less attacked plants may occur. Examples where compensatory growth has been observed experimentally as a result of the presence of disease are relatively uncommon, but in some cases it is strong enough to allow an almost normal crop yield in spite of heavy damage to part of the mixture (Trenbath, 1977). Thus Sibma *et al.* (1964) found in comparing the yields of mixtures of susceptible and resistant tomatoes with and without the nematodes that affected the susceptible plants, that the resistant plant, which was the inferior competitor when the nematodes were absent, was raised to parity with the susceptible plant in their presence. This is an example of a host X pathogen X host interaction.

(5) The discussion of stabilising selection above proceeded in terms of the intrinsic rates of increase of different pathogens, without reference to any interactions between races on common hosts. Differential rates of increase imply that some races are more efficient users of the plant resource than others and so obtain a larger share than their initial proportion in the population might indicate. However more direct interactions between races are possible. Leonard (1969a) found that the total infection level reached by a joint inoculation of two races of *Puccinia graminis avenae* on oats was less than the sum of the separate infection levels resulting from separate inoculations. Since the infection levels were low, so that plant resources should not have been limiting, Leonard inferred the existence of some mutual inhibition of races.

This area of pathogen competition has not yet been thoroughly explored, and the existence of such direct effects is hard to prove unambiguously. Another possible explanation of Leonard's result is that the doubly inoculated plants were inferior hosts as a result of the greater inoculum density and consequent loss of competitive ability relative to unattacked plants (Trenbath, 1977). The only result to set beside that of Leonard as an example of such a direct pathogen X host X pathogen interaction seems to be the counter-example of induced susceptibility reported by Brown and Sharp (1970). The primary type of p X h X p interaction appears to be competition for the plant resource but this has not been verified at the physiological level.

(6) Higher order effects are, understandably, still less well understood. One that has been observed is the adaptation of pathogens to grow better on particular hosts from the range on which they are virulent. Watson (1970b), Leonard (1969b), and Caten (1974) report effects of this kind in which after several generations on one host a pathogen loses some of its vigour on its other hosts. This is treated

here as a host X pathogen X host X pathogen interaction because such divergent evolution in a multiline would tend to reduce the mean vigour of complex races by denying them the opportunity to adapt. No mechanisms for the effect have yet been suggested, though one obvious presumption is that pathogen races are heterogeneous with respect to other host specific characters than those determining immunity, and that prolonged selection on one host reduces this heterogeneity.

Other high-order effects on which no work has yet been done can be envisaged. For instance, it is conceivable that compensatory growth of the less affected components of a multiline may expose them to increased pathogen attack. Alternatively, the multiple infections of simpler components may decrease their capacity to provide cross-protecting spores to more complex components. At present, such effects, if they occur, have not been separated experimentally from lower order effects.

(7) The effects described under headings (1) - (6) do not take into account the special adaptations of individual pathogens or the many-faceted responses of hosts. A susceptible plant in a mixture is not just a source of propagules with a given leaf area. It is a shorter plant that consequently lives in a different microenvironment and responds to environmental stimuli in a different fashion. These areas of plant response are too complex to be approached with any confidence at the moment, and since they have not assumed any role in the theoretical debate on gene use they will not be discussed any further in this study.

2.3.4.2. Significance of the Factors

The factors which have been listed in the previous section can be organised with respect to their significance for gene use under the following three headings:

(1) whether the mode of gene use encourages the appearance of new races

(2) whether the mode of gene use minimises crop damage

(3) whether a host/pathogen system with the superrace present can be stabilised in a beneficial way.

(1) That the use of multilines, by permitting avoidable epidemics, encourages the appearance of new races by giving increased opportunities for the pathogen to mutate has been a major criticism of this mode of gene use for twenty years (Borlaug, 1958). At the same time, the opposite point of view has sometimes been put in the suggestion that widespread disease-free crops exert a "selection pressure" on the pathogen (Mode, 1958; Wolfe, 1975). This usage is misleading. What is created by modern monocultures is not "selection pressure" but "selection opportunity". If a new race is not in existence, the largest and most homogeneous crop will not make it any more likely to appear. It is the consequences of the appearance of a new race that are affected.

Thus the criticism of multilines as encouraging the appearance of new races is valid. Anything other than the concentration of all available sources of resistance into a single plant is subject to this criticism. If this were the only consideration involved in gene use, the only reason for using multilines would be that they allow the simultaneous use of different alleles of resistance genes (Browning and Frey, 1969). Apart from this special case, if multilines or geographical mosaics are to be used it must be for reasons connected with headings (2) or (3) above.

(2) Many of the factors described in the preceding sections can be expected to cause an interaction between the mode of gene use and the extent of yield loss due to disease. Cross-protection, compensatory growth, adaptation of races, mutual inhibition of races and decreased probability of finding a host may all, if they apply, serve to ameliorate the effects of disease spread in multilines and the special cases of multilines. It remains unclear whether these effects make a case either for tolerating losses when the absence of a superrace makes them avoidable, or for tolerating the presence of additional pathogens when the superrace is present.

(3) One factor that might tip the balance towards the use of multilines is the prospect of stabilising the pathogen population. While stability has not usually been precisely defined in this context, it has been regarded as an important attribute of multiline systems (Browning and Frey, 1969). The usage of stability has varied from the conventional mathematical usage of Mode (1958), through the concept of a managed common growth rate of all pathogen races (Leonard, 1969c), to the wish for the avoidance of sudden irruptions of new pathogen races (Knott, 1971).

The most careful investigation of the stability of man-managed crop/pathogen systems (as distinct from natural gene-for-gene systems) has been that of Leonard (1969a,b,c). He observed that in small plots of oats the rate of increase of races of *Puccinia graminis avenae* was proportional to the fraction of susceptible plants in the plot. He also found that among his races a wider range of virulence implied a lower rate of growth on susceptible hosts. He suggested that, given these two relationships, it should be possible to regulate the composition of a multiline so that all races present increase at the same rate. If the relative vigours of the races do not depend on the

composition of the pathogen population, this rate is the rate at which the superrace alone grows. This is a kind of stability but, as Leonard (1969c) pointed out it offers few advantages. However if there is also mutual inhibition of races (as there may have been in Leonard's case), then the natural rate of increase of the superrace can be slowed by the presence of simpler races to an extent greater than would be indicated by simple saturation of the host. In this case Leonard suggested that a sequence of multilines in which the simpler components predominated could be used to make the simpler races predominate, with the result that the superrace would be inhibited. The composition of the multiline could then be adjusted so that this inhibited rate of growth was the norm for all races.

The principal theoretical disadvantage of such a scheme is the apparent rarity of both stabilising selection and the additional effects needed to turn a neutral Leonard stability into an advantageous Leonard stability. No mechanism has yet been suggested for achieving stability without stabilising selection. However, given stabilising selection, an obvious practical question concerns the level of simple races that would have to be tolerated to achieve the suppression of the superrace. Unless mutual inhibition between races were very strong, stability might be achieved at the cost of crop failure. Thus the probability of stabilising multilines is not great, and the probability of achieving a beneficial stability is somewhat lower.

2.4. Summary

The factors discussed in this chapter can be reviewed under three broad headings:

(1) the controversy over the utility of different types of resistance

(2) the range of host/pathogen interactions of various orders to be found in epidemics, particularly multiline epidemics

(3) the reason for seeking a stable pathogen population.

(1) The controversy over different types of resistance has caused some confusion among those seeking to decide the best use of resistance genes. It is clearly helpful to have some verbal organising concepts for the variety of types of resistance and inheritance of resistance. In particular it may be of heuristic value to have expressions for the mean value and the variability of resistance to different races conditioned by a given gene or set of genes. However the evidence quoted in sections (2.3.2) and (2.3.3) shows that the distinctions between general and specific and between polygenic and oligogenic resistance are not sharp. From the point of view of making quantitative decisions about the use of genes these unclear distinctions cannot be of great significance.

(2) A range of factors affecting the progress of epidemics has been considered. Some of them appear to mitigate the effects of epidemics in multilines; some of them suggest disadvantages of multilines. Many of them involve second-, third- and higher-order effects. Thus in the first place the choice of a mode of gene use must be a quantitative one arrived at as the result of a quantitative balancing of factors. And in the second place, simple linear expressions are unlikely to be useful guides to finding that quantitative balance.

(3) The review of stabilising selection and the stabilising of pathogen populations has brought out two points. First, it is likely that stabilising selection is the exception rather than the rule (Martens, 1973). Thus a search for stability does not offer a guide to the use of resistance genes in most cases. Some other guide is

still needed in these cases to resolve the controversy over methods of breeding for disease resistance.

Second, beyond the question of whether stability is attainable is the question of whether it is desirable. The response of plant breeders to the experience of the first half of the century seems to have been, in part to identify the problem as instability (variously defined) of the pathogen population, and then to identify the solution as stability. Yet if the decision is made to manipulate the pathogen population, there is a good case that stability should be avoided. It might be said instead that the appropriate tactic is to keep the pathogen "off balance" by planting lines whose corresponding races are rare, if that is possible. This in fact is the essence of a rotation strategy. Rotation does not give the answer to the problem of the superrace when there is no stabilising selection, but then it is not known what does. Neither stability nor instability is the key to the problem of gene use. In the following chapters an alternative will be suggested.

3. The Choice of Breeding and Planting Policies as a Problem in Optimal Control

3.1. The Choice of Language for Discussing the Multiline Question

The problem of choosing the best mode of use of resistance genes has resisted solution partly because neither the advocates of conventional plant breeding methods nor the advocates of the various alternatives have been precise enough about the aims of programs for breeding disease resistant plants. It is not clear for example, why the possibility of achieving an equal growth rate of all pathogen races (Leonard, 1969c) should be regarded as a point in favour of multilines: such a growth rate, if it were high enough, might be consistent with complete crop failure. On the other hand, there has been no account given of why and, more important, in what way, conventional practices are superior, given the points raised by multiline advocates, such as the devastating effects of epidemics when large areas are planted to monocultures. The parties to the debate have largely talked past each other, rather than discussing the topic in a common language.

It is the purpose of this chapter to suggest that this poor communication can be improved by using at least the language, and possibly the methods, of optimal control. The first step will be to make the suggestion plausible by setting out, in this section, the elements of an optimal control problem and stating the analogy between these elements and the elements of the gene use problem. In subsequent sections the suggestion will be made persuasive by discussing each element in more detail with the objects of:-

(1) introducing the notation of optimal control theory,

(2) making some important simplifying assumptions that expose the basic structure of the gene use problem,

(3) preparing the way for the modelling approach described in Chapter 4,

(4) deducing some theoretical necessary conditions for the use of multilines to be an optimal use of crop disease resistance genes.

The following list is one possible way of setting out the elements of an optimal control problem:-

(1) A system to be controlled.

(2) A criterion of optimality.

(3) A set of manipulable control variables which can influence the optimality of system behaviour.

(4) A sequence of times at each of which an optimal set of current values for the control variables must be decided on. A sequence of such decisions will be referred to as a policy.

(5) The task of achieving the best possible value of the criterion of optimality. For example the task might be to find the sequence of investment decisions that yields the maximum return, given the amount of money to be invested and the response of the market to the operations of buying and selling.

With regard to the first element, it is easy to think intuitively of the crop plants, the disease organisms, other host plants (if any), and the environment as forming a "system" of some kind. This notion is neither very confining or very fruitful, but it serves to establish this part of the analogy at the level of natural language: a more precise characterisation of systems in general and this system in particular will be begun in section (3.2).

With regard to the second element, it can be seen that while the various points of view of multilines do not ignore the choice of a criterion of optimality, there is a tacit divergence in their choices. Opponents of multilines give high importance to the immediate losses to be expected from the use of susceptible plants, while advocates seem more concerned with stability of yield and the husbanding of stocks of resistance genes in the long term. If there is disagreement over the sense in which the use of resistance genes is to be judged optimal, then there is not likely to be agreement on the optimal manner of use. Accordingly, there is need for a commonly accepted criterion of optimality in the argument over multilines, and the recognition of this need would amount to an advance in the understanding of the problem. The choice of such a criterion, however, is not automatic, and will be discussed in section (3.3).

The nature of the problem suggests an obvious set of control variables. The proportions of the different genotypes planted in the crop are manipulable, and can obviously affect any reasonable criterion of system optimality. In fact the genotypic proportions are probably not a complete set of control variables, as will be discussed in section (3.4), but they are certainly classifiable as control variables.

Finally it can be seen that the sequence of crop plantings forms a sequence of decision points at which values of the control variables are chosen. Thus what is done on the sowing date can be thought of as a control action. The plant breeding program is a constraint on the control action at these points. If a genotype has not been produced it cannot be used, and if it costs more to produce then less of it will be used than would otherwise be the case.

The problem of gene use can now be stated as being to choose a sequence of crop compositions which are optimal according to the criterion of optimality accepted by the plant breeder. This is obviously a difficult problem, since it requires that the outcomes of various policies be calculated in advance so that they can be compared. However, unless this can be done to the satisfaction of all concerned, with all the mooted significant factors being considered at the same time in the calculation, the gene use problem cannot be said to have been solved. Because the problem is at present debated in terms of an opposition between diverse crops and monocultures, with the balance of opinion leaning towards the monoculture, the question that will be addressed in this study is whether it is reasonable to believe that it is sometimes optimal to use multilines. This question will be referred to as the "multiline problem", and the null hypothesis of the study is that the answer to the question is in the negative.

3.2. Defining the System to be Controlled

The first element of an optimal control problem listed in section (3.1) is a system to be controlled. So far, the concept of a system has been undefined. It is possible to discuss this concept in a highly formalised manner (e.g. Windeknecht, 1971) or in a highly generalised manner emphasising its general applicability (e.g. von Bertalanffy, 1967). What will be used here is the simple and precise definition due to Bellman (1971): a system is a combination of a state vector $x(t)$, and a rule for determining the behaviour of this vector over time. This rule may well depend on an input vector of some kind. The state vector and the rule governing it are now considered in order.

3.2.1. Defining the State of the System

The state vector itself needs defining, and yet the notion of a state is a singularly elusive one. According to Kalman (1963) it is "intuitively speaking ... the minimal amount of the past history of the system which suffices to predict the effect of the past on the future". Kalman's more mathematical definition is beyond the scope of this work. Another way of considering the notion of state is to treat the choosing of a state vector as corresponding to drawing the boundaries of the system, since it is this choice that separates those variables considered to characterise the system from those regarded as input to it. There is thus no unique choice of state vector.

Environmental factors that affect the multiline problem appear as causes rather than effects and so it is natural to regard them as inputs rather than states, and to draw the boundary of the system around the crop plants and the disease organisms. Since the genotypic proportions of the crop are control variables, and the rules governing the response of the plants to the pathogen, the environment and the passage of time are essentially fixed, by elimination the state vector must then consist of measures of the abundances of the various pathogen genotypes. Special significance will, of course, be attached to the abundance of pathogen propagules at the date at which the sowing decision must be taken, since it is to these that the control action must respond, and it is these which form the feedback information in the system. Therefore it is these presowing abundances that will be treated as the state vector, though it will be assumed that they also give the intensity of the first infection of the crop. Abundances at times between the first infection and the next control decision will be treated as intermediate transformations of the state vector.

3.2.2. Describing the System Evolution Rule

Even without assuming a particular form for the system evolution rule some statements about its nature can be made with confidence.

(1) The rule is nonlinear.

(2) The rule is stochastic.

(3) The rule is nonserial.

These points will be taken in order with the object of clarifying the nature of system behaviour and placing some initial limitations on the way in which the search for optimal policies must be carried out.

3.2.2.1. Nonlinearity

Crop/pathogen systems exhibit many nonlinearities. There are environmental thresholds for disease development (Waggoner and Parlange, 1974), saturation levels of disease severity, and nonlinear yield-loss curves (James, 1974), to name only some of the most obvious effects. Also, not only are the relations between many of the important variables nonlinear, they are mappings from functions to points: for instance, the number of new disease propagules produced on a given day depends not only on the conditions on that day, but, in varying degree, on the whole previous time course of the epidemic.

These considerations, taken together, almost certainly ensure that no simple analytical control law will be discovered. They also rule out the use of most of the simpler traditional techniques of control, so that a numerical treatment of the problem becomes mandatory. In addition there is the possibility that the nonlinearities may be necessary to any optimality of multilines: the optimal policies for linear systems are drawn from the class of what are known as bang-bang controls (Porter, 1969), which in this case would sometimes correspond to rotations of pure line crops. Thus to linearise the system for the

sake of simplicity may result in effectively begging the question that is being asked. Whether a linear model of the rule is adequate for some particular practical purpose is a somewhat different question, though one that is unlikely to have an affirmative answer, if only because of the strong perception by plant pathologists of the non-linearity of the systems with which they deal.

3.2.2.2. Stochasticity

Crop/pathogen systems are also stochastic in nature. This statement does not refer primarily to unpredictable weather changes, although such unpredictability will in practice result in much of the difficulty involved in deciding on optimal gene use. Rather, the basic unpredictability of system behaviour results from the occurrence of mutational changes in virulence and from the irruption of new pathogen races from outside the system boundaries. While a first simplification of the problem might ignore the variation in the environment, either to gain an initial understanding of system behaviour, or because a particular environment under study behaved similarly in each crop cycle, it could not ignore the appearance of new races, since without them a clean crop would stay clean for ever. Thus the multiline problem is inescapably stochastic.

3.2.2.3. Non-Seriality

This stochasticity leads on to the nonserial (Nemhauser, 1966) nature of the system, because at any moment in the crop cycle the system may change direction as the result of the appearance of a new race and begin to evolve along a path in state space quite different from the one it was following before. Another way of saying this is that the trajectory starting at any point in state space has many

branches. In fact this property turns out not to have a large effect on the method of calculating the optimal policies, but it may be of some assistance to think of the whole system as undergoing a kind of Markov process, in which each state is defined by which genotypes are present and which are absent, and the transitions between states have probabilities depending on the state and the frequency of mutations in the pathogen population. The final absorbing state of the process is the one in which all known resistance genes are vulnerable. The control action must take some account of the various pathways by which this final state may be reached.

3.3. The Choice of an Optimality Criterion

3.3.1. The Necessity of an Economic Component in the Criterion

It is easy to show that some simple and obvious criteria of optimality for the behaviour of crop/pathogen systems have unacceptable properties. Maximising total yield over one year may be shortsighted if the resulting composition of the disease population causes great losses in subsequent years. Maximising average yield over a period of years is compatible with achieving alternating shortages and gluts. An algorithm designed to find the control policy that minimised the variance of yield might achieve it by a uniformly zero yield unless other considerations were included. Drawing upon these cautionary examples, the chooser of a criterion might decide to maximise average yield subject to a constraint on the coefficient of variation of yield, but the choice of the constraint would be arbitrary unless guided by information on the range of variation that is tolerable. In short, if the defects of simple a priori criteria are to be remedied by trading one criterion off against another, some trade-off scale based on human judgements of utility is indispensable. Also, if some information is

available about utility preferences it makes sense to deal directly with this rather than to work in terms of arbitrary criteria whose familiarity may not imply economic trustworthiness. It remains true that the criterion chosen must be numerically definable. Extending "indefinitely" the "useful" life of a resistance gene (Browning and Frey, 1969) is only acceptable if the definition of "useful" is acceptable, and although this is an economic judgement in embryo, it is hard to see how it can directly be made more precise.

Thus it is necessary in devising a criterion to give an account of who is to benefit from the optimal policy and in what way. The optimal policies prescribed as a result of an optimal control exercise are bound to differ depending on the answers given to these questions. Moreover, there is unlikely to be a single pair of answers which will coincide with the judgements of all those who have an interest in a solution to the multiline problem. A consensus may develop in the future but the procedure adopted in this study will be to distinguish two apparently very different cases, in the hope of illuminating the range of criteria that might be chosen, and of exposing some principles that might apply in a range of situations. In each case what is offered is not a definitive treatment, but a simple model based on some of the salient economic concepts as a starting point for further explorations. In Chapter 4 two models are introduced that are intended to describe respectively an agribusiness in a rich country, and a society of poor subsistence farmers, so that the utility of different policies can be compared in each context.

3.4. Necessary Conditions for Multiline Use to be an Optimal Control Action

The basic facts about the choice of control variables and the timing of control actions have been set out in section (3.1). However it will be shown below that there are some fairly plausible necessary conditions which must be satisfied if multiline use is to be an optimal control action, and that these conditions necessitate the extension of the set of control variables in an unexpected and counter-intuitive way. In order to demonstrate these necessary conditions two basic assumptions about the biology of multilines will be made.

(1) the multilines to be considered are of the Iowa type (Browning and Frey, 1969), composed of nearly-isogenic plants, differing only in their complement of major genes for resistance.

(2) the superrace of the pathogen can grow equally well on all lines in the mixture.

These assumptions will be carried forward in the rest of this study except where specific exceptions are made. Frequently, slight deviations from the assumptions will not be significant.

It is now possible to prove the following simple theorems:

Theorem 1.

If for a given crop/pathogen system, a given criterion and some state, the optimal control action is a multiline, then the optimal control action will result in avoidable disease.

Proof:-

By Assumption 1, the composition of the multiline cannot affect the criterion in the absence of disease. By Assumption 2, the composition of the multiline cannot affect the criterion by affecting

the dynamics of the superrace in the absence of the simpler races. Therefore if the multiline is optimal it must be because of the presence of simpler races. But the disease caused by simpler races is avoidable by planting the most complex cultivar (the "superline") exclusively.

A second and more surprising result can also be proved.

Theorem 2

If for a given crop/pathogen system, a given criterion and some state A, the optimal control action is a multiline, then there is at least one other state B for which the optimal control action is a multiline together with the deliberate addition of quantities of the simpler races to the system.

Proof:-

By Theorem 1, the simpler races are present in state A. Let V(A) be the cost of the control action of planting the optimal multiline as a response to the state A. Let V'(A) be the cost of planting a pure crop of the superline instead. Now

$$V'(A) > V(A) \qquad \ldots (3.4.1)$$

because V(A) is optimal. Let B be the state equal to state A except that the simpler races are absent. Then

$$V'(A) = V'(B) \qquad \ldots (3.4.2)$$

because the simpler races cannot grow on the superline.

But $$V(B) = V'(B) \qquad \ldots (3.4.3)$$

by Assumption 2, because the superrace can grow equally well on all lines and the simpler races are not present, in state B.

Thus $$V(B) > V(A) \qquad \ldots (3.4.4)$$

But the difference V(B) - V(A) occurs only because of the presence of the simpler races. Thus if it costs less than some set amount per propagule to produce the simpler races artificially, the optimal control

response to state B involves the addition of the simpler races to the system.

In the stochastic case statements about the presence or absence of the simpler races must be replaced by statements about the expected amounts of the simpler races. Only slight restatements are needed, as follows:

Theorem 1a.

If for a given crop/pathogen system, a given criterion and some state, the optimal control action is a multiline, then for this control action the expected amount of disease will include some avoidable disease.

Proof:-

This is only a restatement of theorem 1.

Theorem 2a.

If for a given crop/pathogen system, a given criterion and some state A, the optimal control action is a multiline, and the optimal control action when the expected amount of the simpler races is less than some set amount is not a multiline, then there exists a state B for which the optimal control action is to plant a multiline and deliberately add some of the simpler races.

Proof:-

This is the same as the proof of Theorem 2, except that the additional assumption about the lower bound on the effective amount of the simpler races is intended as a realistic indication of the situation faced by the designer of an optimal policy, for whom the expected amount of the simpler races is unlikely to be zero.

These theorems are necessary conditions for multilines to be optimal. Clearly, slight deviations from the assumptions will not affect the conclusions, though it is impossible to say how slight 'slight' is. If for example, there was a strong adaptation effect for the superrace growing on different lines, Assumption 2 would no longer hold (more strictly, there would no longer be a single superrace) and multilines without avoidable pathogens would be more likely to be optimal. It is less easy to construct sufficient conditions in terms of the factors discussed in Chapter 2 because their effects are complex and depend on the multiline composition. I do not know of a sufficiency conditions of the same significance as the results given above. In section (3.5.3), after some useful notation has been introduced, some fairly weak sufficiency conditions will be given.

Theorems 2 and 2a are in fact a generalisation of Leonard's proposal, discussed in section (2.3.4.2), for the manipulation of multilines so that the amount of the simpler races present increases. The generalisation arises because the need to increase the amount of the simpler races can now be seen as a necessary feature of the optimality of multilines in general and not just of those multilines where (as in the Leonard case) there is mutual inhibition of races.

The need to increase the amount of simpler races present can be met in two ways; by the release of propagules from greenhouse cultures or by the growing of multilines. So far the tacit assumption has been that glasshouse cultures are used and that their cost is negligible. In what follows the release of pathogen races (other than the superrace) will be included as a normal measure in the optimal control of a crop/pathogen system. It will be assumed that any control pathogens are released at the same time as the crop is normally infected, which is computationally convenient and realistic (if such an untried measure

can be described as realistic), because the earlier the pathogen is added, the more effective it is expected to be.

The differing roles of naturally produced and cultivated pathogens can be understood in terms of their relative cost. If cultivated pathogens are cheap and the conditions of Theorem 2 hold, then it will be sensible to add pathogens directly to a crop believed to be in need of them. If cultivated pathogens are expensive, there will be a case for anticipating their use by growing a multiline in a previous crop cycle. This anticipatory use will have a cost of its own, so that as the cost of cultivated pathogens rises the advantage of using a multiline in state B of Theorem 2 will gradually disappear.

3.5. The Task of Finding an Optimal Control

3.5.1. Optimal Control: Further Concepts and some Notation

The method commonly used for the solution of optimal control problems of the kind posed here is what is known as approximation in policy space (Bellman, 1957: Kushner, 1971), which gives an optimal feedback control for the system. Before the notation to be used is introduced, these concepts require further explanation.

Optimal control problems admit two kinds of answer: an open or a closed loop (feedback) control. In open loop control, a policy is specified in advance for given initial and final conditions and a given number of decision stages. Control then consists of implementing the precalculated decisions at each stage as the process unfolds. This type of control is particularly associated with the name of Pontryagin (Pontryagin *et al.*, 1962) and the Maximum Principle he derived. Here the idea is that a function of the criterion of optimality and the system equations can be constructed with the property that if it is

maximised at each stage of the process with respect to the control variable values at that stage, then the resulting policy will be optimal. The use of the Maximum Principle is most common in deterministic, continuous problems, though the extension to discrete-time problems is often possible (Jordan and Polak, 1964: Halkin, 1964), and though maximum principles have been proposed for stochastic systems (e.g. McReynolds, 1972).

By contrast, in closed loop, or feedback control an optimal control action is derived for each current state of the system, so that a response is made to system behaviour, rather than a course of action being laid out in advance. This approach is associated with Bellman (1957) and his development of Dynamic Programming, and has the advantage that under some circumstances each decision has to be made optimal with respect to only one set of control variables, and not all the sets of control values that must be chosen in subsequent stages.

These two methods of optimal control are not fundamentally different, but represent two alternative views of the same mathematical subject matter. The equations of Pontryagin's maximum principle can be shown to be logically equivalent to Bellman's equations for optimal feedback control (McAusland, 1969). The tendency to handle deterministic problems more frequently as open loop problems and stochastic problems as closed loop problems seems to be partly a result of the historical accidents of the development of the subject. However, in at least one kind of stochastic problem it is more reasonable to use a closed loop approach. This is the case in which information about the parameters of the source of random variation can be gained from observing the results of the previous state of the random process. Where random variation appears, as in many engineering problems, as noise with well defined statistical behaviour, the accumulation of information from stage to stage may not be great. However, in the crop/pathogen system

with its non-serial, irreversible character, knowing that a race has arrived makes a marked change to estimates of the future behaviour of the system, and because of this, the following notation, from Kushner (1971), applying to the feedback control of stochastic systems, is appropriate.

$\pi(i)$: The control policy used starting at state i. A policy is a sequence of control decisions.

$V(\pi(i),i)$: the cost (in terms of the criterion of optimality) of the control policy (i), starting from the current state i.

X : the state vector of the system.

$u(X)$: a control decision taken in response to X.

$k(X,u(X))$: the cost of single stage of a policy where u(X) was applied to X.

ρ : discounting factor, representing the ratio of the value of an asset available in a year's time to its value if it were available now.

$E_i^{\pi}(y)$: expected value of some quantity y, given state i and policy $\pi(i)$.

The general problem of control can now be stated as being to discover, for each admissible state i, some $\pi^*(i)$ such that

$$V(\pi^*(i),i) <= V(\pi(i),i) \qquad \ldots (3.5.1)$$

for all allowable π, that is, to minimise V.

The general statement of equation (3.5.1) simply establishes $V(\pi(i),i)$ as a measure of the utility of the outcome of a state and a policy. There are many possible ways of comparing different costs and

setting up utility scales, and the axiomatic treatment of rational behaviour under conditions of uncertainty and risk is a controversial subject (Balch and Fishburn, 1974). The range of ways of comparing costs extends from setting each value of V equal to the expected total cost of the sequence of costs k to setting it equal to the greatest possible total cost of any sequence of costs k. The use of the expected cost can be criticised as being insensitive to the occurrence of catastrophic costs occurring at low probability (Luce and Raiffa, 1957). The use of the maximum cost as the figure to be minimised (a minimax approach) can be criticised as being unduly pessimistic in cases, such as this one, where there is no intelligent adversary seeking to maximise the cost (Shubik, 1975).

In this study the procedure that will be adopted will be to set V equal to the discounted sum of the expected values of k. This procedure is a standard one in works on optimal control (Kushner, 1971), and the meaning of its results is easy to understand. When it is adopted, the optimal control problem takes the form:-

$$V(\pi^*(i),i) = \min_{\pi} E_i^{\pi} \rho(k(i,u(i)) + V(\pi(X_1),X_1)) \quad \ldots (3.5.2)$$

That is, the optimal control $\pi^*(i)$ given the initial state i is that control which minimises the expected value of the sum of the cost of the first stage of the policy and the cost of the infinite policy beginning with the output, X_1, of the first stage. The equation is called functional because V is a function of i, and recursive because V appears on both sides of the equation with different arguments. It is the process of solving this equation that is known as approximation in policy space (Bellman, 1957). The equation itself is an expression of the principle of optimality (Kushner, 1971). The significance of the discounting factor ρ, will be explained in section (3.5.2.1). Some possible ways of making the cost k more pessimistic

and some uses of a minimax approach will be discussed in Chapter 4.

3.5.2. Three Important Simplifications and their Significance

In the three sections that follow the structure of the problem is simplified by the use of three assumptions. These undoubtedly restrict the general applicability of any results of this study, but in each case there are compelling reasons for accepting the restriction.

3.5.2.1. The Use of Policies of Infinite Length

The first simplification to be made is that only infinite time horizons for planning policies will be considered. This is computationally convenient, but it is also realistic from two points of view. The convenience arises because, as it is easy to show (Kaufman and Cruon, 1969), for an infinite horizon the choice of an optimal control depends only on the current state. This can be seen intuitively by observing that each decision in an infinite process is the beginning of another infinite process, so that the optimal decision in any state is the optimal decision for an infinite process starting from that state. Thus the position of the given state in a sequence of decisions is no longer significant and only the current state need be taken into account.

The realism of the simplification can be supported by two arguments. In the first place, in a situation in which resources can be invested at a profit outside the agricultural enterprise the costs of epidemics in the future should be discounted. When this is done, the contributions to V of long term consequences of policies stretching in the limit to infinity can be included without disturbing the result. In the second place, if the stock of resistance genes is finite, they must be husbanded from now to infinity as a nonrenewable resource.

Thus both the pragmatist making sound economic decisions and the conservator may wish to use an infinite planning horizon. The latter, however, must be careful about his choice of optimality criterion, because if he places an equal value on the costs of crop failures in all epochs, that is, he does not discount, the calculated cost of his policy may well increase beyond any bound, making the existence of the minimum of V doubtful.

3.5.2.2. The Omission of Environmental Variation

The second major simplication that will be made is to ignore the effects of environmental variation on the behaviour of the system. The chief reason for doing this is that to include the environmental variation would increase the amount of computation required by some power of the number of different environmental states considered. However the form of the criterion has been chosen in such a way that if multilines are suboptimal without environmental variation they will also be suboptimal with it.

Suppose that the variability of the environment is expressed as a set of environments, $s_j \varepsilon S$, each with an associated probability of occurrence p_j. Then the expected cost of the optimal policy is

$$V(\pi(i),i) = \sum_j p_j \, V(\pi(i), i, s_j) \qquad \ldots (3.5.3)$$

If for all $s_j \varepsilon S$, $\pi(i)$ is a pure line of the most complex cultivar, none of the component $V(\pi(i),i,s_j)$ of the summation can be reduced by the use of a multiline. Consequently, the expected cost cannot be reduced by the use of a multiline, and consequently the introduction of a variable environment cannot make multilines optimal. This reduces the significance of the omission of environmental variation from this study.

Assumption 3

The minimum cost associated with a state increases monotonically with the amount of the superrace present in that state, all other factors being constant. (This assumption can be viewed alternatively as a necessary quality of a 'sensible' cost function).

Theorem 3

If for a given crop/pathogen system and a given criterion there exists a state A in which a multiline can be grown which decreases the output of the superrace from that state below the pure line level without increasing k(u(A),A) over the pure line level, then the optimal control action for state A is a multiline.

Proof:-

The only way of decreasing the superrace output is by the presence of the simpler races, and hence the use of a multiline. By Assumption 3, decreasing the output of the superrace decreases the cost incurred over all subsequent crop cycles. Thus if it does not increase k(u(A),A), and does decrease the superrace output, a multiline can be optimal.

A similar, but weaker, result is obvious:

Theorem 3a

If for a given crop/pathogen system and a given criterion there exists a state A in which a multiline can be planted which decreases k(u(A),A) below the pure line level without increasing superrace output, then a multiline is optimal in state A.

Comment

This result is called weaker because of the likelihood that, in the absence of mutual inhibition, the effect of the simpler races will be to increase k(u(A),A).

In fact, any real multilines that are optimal will presumably balance out positive and negative effects on $k(i,u(i))$ and $V(\pi(X_1),X_1)$. It can be expected that at sufficiently high infection levels the simpler races can decrease the superrace output by denying it some of the crop on which to grow. That these high levels can be achieved without increasing $k(i,u(i))$ appears improbable. The weakness of theorems 3 and 3a results from the difficulty of dealing qualitatively with the quantitative balance of the two kinds of effect on the cost. The quantitative aspect must be handled if multilines are to be used with confidence. This is the task of an optimal control algorithm, which passes judgement on the sufficiency of numerical conditions proposed to it. Chapter 4 considers the process of transforming the qualitative arguments about the theory of resistance gene use into the quantitative form that a numerical algorithm can accept.

3.6. Summary

The idea of using optimal control methods for biological and particularly agricultural purposes although relatively new is now fairly widespread. There is a gradual evolution from the simulation of agricultural systems as an end in itself, to single-stage optimisation of inputs to systems, to sequential optimal control. Arnold and Bennett (1975) and Van Dyne and Abramsky (1975) give annotated bibliographies of the application of linear and dynamic programming to agricultural problems.

The present study is bracketed fairly closely in terms of subject matter by studies of the use of optimal control in the control of pests (Shoemaker, 1974a,b,c) and by studies of optimal prophylaxis in human epidemics (quoted in Banks, 1975). However, the aim of the present study is of a somewhat different kind in that it seeks to

answer a theoretical question rather than to find a specific optimal control for a particular system. This chapter has been mainly concerned to construct a basis for a theoretical approach, though in the process some analytical results have been obtained.

The general conclusion of this chapter is that the language of optimal control is a convenient language for the qualitative discussion of the problem of resistance gene use. It allows a previously unsuspected extension of ideas about the way in which epidemics can be controlled (section 3.4) and gives insight into the way in which successful multilines might operate (section 3.5.3). Most important, it dispels the idea that simple arguments can show what mode of gene use is best, and shows that a quantitative approach is necessary if the criticisms of conventional plant breeding reviewed in Chapter 1 are to be endorsed or set aside.

4. A Theoretical Approach to the Multiline Question via Modelling

4.1. Introduction

The purpose of this chapter is threefold. The first aim is to introduce and discuss the philosophy of an approach to the optimal control of crop composition. The second is to describe a particular implementation of this philosophy. The third is to consider the robustness of this implementation.

Section (4.2) describes, and attempts to justify, an approach via simulation modelling to solving the multiline problem in the optimal control formulation that has been given in Chapter 3. Section (4.3) describes the necessary structure of such a simulation model, and in sections (4.3.1 - 4.5) a particular model based on current concepts in the relevant fields of knowledge is described. The description of the crop/pathogen model itself is followed in section (4.6) by a description of the cost function models used to derive two alternative criteria of optimality, and this description is followed in section (4.6.3) by a discussion of the robustness of these cost functions and the extent to which they represent rational responses to the level of crop yield. In section (4.7), the difficult topic of testing the robustness of conclusions drawn from simulation modelling is discussed, and a type of test, to be used later in the study, described. The discussion of parameter values is reserved to Chapter 6, except for a discussion in section (4.8) of the number and type of resistance genes to be modelled. Section (4.9) summarises the chapter.

4.2. Nature and Justification of the Approach

Chapter 3 has indicated in general terms what information is needed in order to solve the multiline problem for any given crop/pathogen system. What has not been discussed is how this information is to be acquired. The hardest element to acquire of those listed in Chapter 3 as making up an optimal control problem is evidently the system evolution rule, since it is unaffected by human decisions and is a property of large scale phenomena. The most direct empirical way to acquire it would be via a large experiment in which the input/output relations of a range of multiline compositions were measured for a range of input conditions. These measurements, with interpolations, could be used as a direct numerical representation of the rule which could then form the basis for computing an optimal control. Such an experiment would be large. If a three-genotype system were to be taken and quadratic nonlinearities detected, at least 7x3x3 large, well isolated plots would be required in each replicate. Yet this experiment would not solve the problem of extrapolating from a plot to a region, or that of linking the state at the end of one crop cycle to that at the start of the next.

The task of a theoretical approach to the multiline question where the empirical approach is so difficult is twofold. First, to suggest whether there are grounds for supposing that the empirical enquiry is necessary, that is, whether the arguments proposed in favour of multilines are valid, and, second, to simplify the empirical enquiry if possible by suggesting a more discriminating experiment than the one described above. Most of this study is concerned with the first half of this task. A discussion of the second half will be deferred to Chapter 7.

The available techniques for studying whether the arguments proposed in favour of multilines are valid are formal logic, analytical mathematics and numerical mathematics. Formal logic is ill adapted for argument about dynamic systems. The laborious manner in which simple biological conclusions were reached through formal logic by Woodger (1939) is evidence of this. By comparison, analytical methods are well adapted to considering dynamic systems, but their quantitative application is limited to linear systems and special cases of non-linearity (Bellman, 1971). The multiline question involves non-linear dynamics which do not seem to be among these special cases, if only because of their complexity. By elimination the only remaining technique is to synthesize a numerical optimal control of a numerical simulation of the system, and see whether it prescribes the use of multilines. However it is open to question whether numerical simulations should not also be eliminated from the list as well, on the grounds that they are impractical or unconvincing.

Numerical simulation has the disadvantage that there can be no confident generalisations made from the results of one simulation to those of another simulation with different parameter values or different functional forms. In particular, it is usually arguable that the departures from total realism of simulations of complex biological systems deprive their conclusions of credibility. There is thus a dilemma: on the one hand numerical simulation is the only technique that can represent any of the detial of complex biological systems. On the other hand it cannot represent all of the known complexity, nor show what effect its simplified structure (compared to the natural systems) will have on its output. Thus simulations lack credibility as a means of convincing specialists of its application to their areas of knowledge about crops and pathogens.

The utility of numerical simulations can be partially justified, despite the dilemma, by two considerations.

(1) If the results produced by a simulation are not credible, then opinions about modes of gene use arrived at without benefit of simulation are less credible still. Thus even a decision to persist with conventional plant breeding must lack credibility if modelling lacks credibility, because it must be made on the basis of verbal arguments resembling greatly simplified simulations.

(2) In order to test the validity of the arguments for and against multilines it is not necessary to integrate all that is known about crops and their pathogens. The task at hand is not to call upon a body of knowledge in order to explain why multilines do or do not "work", but to assess a body of arguments, that are largely unsupported by evidence and are poorly integrated with each other. The first test of these arguments is to see what conclusion they support when integrated in a model. To do this is a matter of ascertaining and modelling the current level at which knowledge about crops and their pathogens is represented in the debate about multilines. The second test is to see if the conclusions are altered by including the most prominent factors not mentioned in the current level of debate. If both these tests give no support to multilines, then it may be said that there is no good reason known to believe what multilines are optimal.

In such a situation it would be up to the advocates of multilines to suggest why the inclusion of more realism in the model than was present in the pre-model debate should,

in theory, lead to the optimality of multilines. If however the two tests are positive the onus is thrown onto the opponents of multilines. A conflicting pair of results should point to a better understanding of what biological factors are most important in determining the nature of an optimal control.

Whatever the results of these tests, however, our present ignorance about which of the many possible biological factors may be important makes it necessary to emphasize the improvement of the debate about the most obvious factors rather than the inclusion of more realism, since the existing level of debate cannot handle even the more obvious factors quantitatively. Thus when, in the rest of this chapter, models of multiline phenomena are introduced, it is with the intention of maintaining contact with and extending the current level of debate on the multiline question, rather than of organising the full complexity of epidemic phenomena. The conclusions of this study should be viewed in this light.

4.3. Elements of the Epidemic Simulation

The elements that are required of a simulation of an epidemic in a multiline for the purpose of giving a system evolution rule are:

(1) a description of pathogen increase in the crop cycle as a function of crop composition.

(2) a description of the effects of pathogen increase on the crop, especially the effect on yield.

(3) a description of the period between harvest and the next sowing.

At present there is only one published model (Trenbath, 1977) which addresses itself to all these requirements, and there is no consensus on an adequate level of realism for a generally accepted simulation of them. Most published work has been done on the simulation of disease spread. There are a number of expressions that relate infection level to yield loss, though these have not been developed in conjunction with models of disease spread. Very little is known about the interface between crop cycles. In the following three sections these requirements will be dealt with in this order. In each section the current state of simulation will be reviewed, and then a model of the simulation element concerned will be proposed in the spirit of section (4.2).

4.3.1. A Simple Simulation of Disease Spread

A useful concept for making systematic comparisons of different models is Levins' (1968) distinction between general, realistic and precise models. According to Levins there is a tendency for these attributes of models to be mutually exclusive, and this seems to be the case for simulations of disease spread. The most precise models are linear multiple regression models used for predicting disease levels from previous disease levels and from environmental variables (Eversmayer *et al.*, 1973). The most realistic models are those of Waggoner and Horsfall (1969) and Waggoner *et al.* (1972) in which each stage of the life cycles of (respectively) *Alternaria solani* and *Helminthosporium maydis* is modelled in some detail. The most general models are those which seek to fit some non-linear curve to the rising level of disease incidence or spore production (van der Plank, 1975; Jowett *et al.*, 1974).

Linear regression models are quite inappropriate for the purpose of this study because of the need to deal with the second and higher order effects discussed in section (2.3.4.1.), and because they can only be valid over a restricted range of state inputs. By contrast, the highly realistic models offer scope for incorporation of complex effects, but are not usable as general theoretical tools. One reason for this unsuitability is their specificity, which would make the general applicability of any results doubtful. A more immediate and practical reason is their complexity. The EPIDEM program (Waggoner and Horsfall, 1969) consists of about 500 FORTRAN statements and proceeds in time increments of three hours. As will be seen in Chapter 5, an optimal control algorithm must run thousands of simulated epidemics in the process of choosing an optimal control, and a model of the complexity of EPIDEM (even though it does not consider yield effects) is already too complex for use in this fashion.

By elimination, an extension of the current level of debate about disease spread must begin, if at all, with the use of the simpler non-linear curves for disease increase. The basis for most of these curves is the two most striking aspects of the growth of air-borne fungal pathogens: their capacity for exponential increase and their necessary self-limitation by exhaustion of host tissue. The simplest model that can express these two aspects is:

$$\frac{dy}{dt} = ry(1-y) \qquad \ldots (4.3.1)$$

where y is the fraction of plant tissue infected

r is a constant

This has as a solution the logistic equation

$$y = \frac{1}{1+e^{-rt}} \qquad \ldots (4.3.2)$$

The obviousness of this model should not be confused with accuracy. Although disease progress curves often appear logistic, there are other S-shaped curves that sometimes are a better fit to epidemic data (Cournoyer, 1970; Jowett *et al.*, 1974), although it is usually not easy to see the biological basis for choosing them. Van der Plank (1975) has also criticised the logistic equation on the grounds that equation (4.3.1) does not provide for either a latent period before infected tissue becomes infectious or a finite infectious period. He proposed modifying the equation (4.3.1) to include these effects as follows:

$$\frac{dy(t)}{dt} = r(y(t-T_p) - y(t-T_p-T_l))(1-y(t)) \quad ...(4.3.3)$$

where T_l is the latent period of the infection

T_p is the infectious period of the infection.

Since van der Plank's work is undoubtedly at the current 'level of debate', and is expressed in terms of biological mechanisms, and since the equation is apprently reasonably consistent with the behaviour of real epidemics (van der Plank, 1975), equation (4.3.3) will be used as the basis for this study.

The generalisation of equation (4.3.3) to fit the multiline case is not difficult, but it is helpful to have a biological model in mind. The r in equation (4.3.3) (the relative rate of increase) combines two coefficients: the rate of production of propagules per unit area of infected tissue, and the rate at which those propagules convert healthy tissue into infected tissue. The rate of conversion itself depends on the probability of a propagule landing on uninfected tissue and on the expected area that it converts if it does. These separate effects must be kept track of when different propagules have different probabilities of finding a host. Note that van der Plank's equation

assumes a constant plant size of 1 so that the probability that a propagule lands on a plant remains constant. The generalised version of equation (4.3.3) can be written as follows:

$$\frac{d(y_{ij}(t))}{dt} = \sum_m R_{ijm} (y_{im}(t-T_l) - y_{im}(t-T_l-T_p)). (Y_j(t) - \sum_k y_{kj}(t)) - \delta_{ij} \frac{d(y_{ij}(t-T_c))}{dt} \quad ...(4.3.4)$$

where $y_{ij}(t)$ is the area of plant tissue on cultivar j that is infected with race i at time t.

$Y_j(t)$ is the total area of cultivar j.

and

$$R_{ijm} = s_{im} f_{ij} \frac{1}{Y(t) + d} \quad ...(4.3.5)$$

where Y(t) is the total crop area

f_{ij} is the expected area taken up by an i-propagule on a j-cultivar

s_{im} is the production rate of i-propagules on cultivar m in numbers of spores per unit area per unit time

d is a constant which establishes the law governing the probability of a propagule landing on a plant as the total crop area increases.

δ_{ij} is 1 if race i crossprotects cultivar j, and is 0 otherwise.

T_c is the time for which the crossprotection effect lasts.

R_{ijm} corresponds to r in equation 4.3.3 and $(Y_j(t) - y_{kj}(t))$ corresponds to (1 - Y(t)), but with the difference that since 1 is replaced by $Y_j(t)$, it is not assumed that the plants remain at a constant size. The other innovation in equation (4.3.4) concerns the implementation of the cross protection effect. Since the effect is local, not systemic, and induced over about the same time span as

the infection period (Littlefield, 1969), it is reasonable to treat it in the same way as infection is treated: as an area being taken over by a propagule, though in this case the propagule induces immunity. The size of f_{ij} when i is not virulent on j and the length of time for which the effect lasts can be used to control the strength of the effect. The differences between cross protection and ordinary infection are that the relevant s_{im} is zero and that the last term of equation (4.3.4) is used to return cross protected area to the unoccupied category after the effect has worn off (typically six days in Littlefield's (1969) example). The unoccupied area is thus continuously increased at the rate at which it was decreased by cross-protecting spores T_c days before.

The treatment of dispersal and deposition requires further comment. Much work has been done on the effects of environmental factors, chiefly wind velocity, on the removal of propagules from infected tissue (Meredith, 1963), their dispersal (Schrodter, 1960; Roelfs, 1972; Gregory, 1973), and on their deposition on other plants (Chamberlain, 1974). In this study, where a lumped system (section 3.5.2.3) is being treated, these are only of significance in so far as they affect the probability of a propagule finding a host: the spatial pattern of deposition is not represented. Since environmental variation is not being considered either, the only variable that can affect the probability of a propagule landing on another plant is the stage of growth of the crop. In the van der Plank equation, where the crop does not grow, this effect does not occur. However in this model it is assumed that the growth of the crop affects the probability of deposition on healthy tissue of line j according to a hyperbolic law

$$p_j(t) = \frac{\Delta y_j(t)}{(Y(t) + d)} \qquad \ldots(4.3.6)$$

where $$\Delta y_j(t) = Y_j(t) - \sum_k y_{kj}(t) \qquad \ldots(4.3.7)$$

In the computation of the model, equation (4.3.4) is implemented as a difference equation with a time step of a day, because this is more realistic and computationally faster than solving the differential equation. There is therefore a need to arrange the deposition of a finite number of spores which may take up overlapping areas. The procedure which is adopted (suggested by a method in Justesen and Tammes (1960)) is to solve the linear equations

$$\frac{d(\Delta y_j(s))}{ds} = -\Delta y_j(s) \sum_i f_{ij}\, s_i / \sum_i s_i / (Y(t) + d) \qquad \ldots(4.3.8)$$

where s is the number of spores deposited

s_i is the number of spores of race i produced that day

The solution of this equation is

$$\Delta y_j(s) = \Delta y_j(0) \exp(-\Sigma f_{ij}\, s_i .\, s / \Sigma s_i / (Y(t) + d)) \qquad \ldots(4.3.9)$$

and when all the spores have been deposited, i.e. when $s = \sum_i s_i$, the unoccupied area has been reduced in the ratio

$$\Delta y_j\ (\Sigma s_i)/\Delta y_j(0) = \exp(-\Sigma f_{ij} s_i / (Y(t) + d)) \qquad \ldots(4.3.10)$$

with the various races taking up the new infected area in proportion to

$$f_{ij} s_i / \Sigma f_{ij} s_i$$

4.3.2. The Effect of the Disease of Plant Yield

The work quoted in the previous section takes no account of the effect of the disease on the growth of the plant. Similarly, the work that has been done on yield loss resulting from disease takes no explicit account of the spread of disease. The most recent review of the prediction of yield loss from disease intensity is that of James

(1974) which distinguishes three types of approach. All three are examples of the use of a convolution function approach (Sokolnikoff and Redheffer, 1966) in which percentage loss at harvest time is given by

$$x(T_h) = \int_0^{Th} f(T_h-t)g(y(t))dt \qquad ...(4.3.11)$$

where $f(T_h-t)$ is a weighting function for the effect on yield of a function g of infected plant area y at time T_h-t. The three approaches are distinguished by their choice of f

(1) the critical point model, where f is a single 'spike' or Dirac pulse at $t = T_h-t'$, the 'critical point' in crop development at which the disease exerts its effect.

(2) the multiple point model where f is a series of spikes centred at different times.

(3) the area-under-the-curve model due to van der Plank (1963) where f is a rectangle over $0 <= t <= T_h$.

Current opinion (James, 1974) tends to prefer either (1) or (2), with (3) being regarded as less flexible. According to James, the choice between (1) and (2) favours (2) for highly variable or long epidemics or for long periods of yield accumulation. At the level of generality of this study these considerations do not provide a basis for choice. The function that will be used was chosen on the basis that it had been used to provide a very accurate prediction of yield loss from disease intensity by Romig and Calpouzos (1970) who used a critical-point model with g as a logarithmic function so that

$$x(T_h) = a + b\ \ln(\dot{y}(T_h-t')) \qquad ...(4.3.12)$$

for $0<= a + b\ \ln(y(T_h-t')) <=100$

Using this relationship for yield loss of spring wheat to stem rust, they found a correlation between disease intensity at the critical point and yield loss over the 5-95% range with r = .993. It is equation (4.3.12) that will be used in calculations of yield loss in this study.

4.3.3. Incorporating Plant Growth Effects

The last two sections described two areas of work that have not overlapped. Because of the complexity of modelling plant growth it has been simpler for those concerned with yield loss not to enquire into the dynamics of the disease. The gap between the two approaches makes it impossible to simulate the hostXpathogenXhost and higher order interactions of section (2.3.4.1) without extending the current level of debate on epidemic simulation. In this section an extension of this kind which can be used as an optional feature of the simulation model is described.

The literature gives little guide on how to make this extension. A number of simulation models of plant growth now exist which simulate at the level of partitioning photosynthate between plant organs. A review is given by Hesketh and Jones, (1976). None of these studies has generated a simple equation, like equation (4.3.3) which can be used as the basis of a generalised discussion of plant growth: instead their models are complex (and probably must be) and have the same disadvantages for the purposes of simulating plant growth as the EPIDEM approach to epidemiology. Also, the studies of plant growth involve detailed representation of exactly those plant functions that the disease affects, and so their models would have to be drastically changed to take account at the same level of realism of the presence of a pathogen.

An exception, and so far the only indication of how to extend the level of debate is the model of Trenbath (1977). This model has a number of similarities to the one to be presented here in that it uses a growth law for the disease resembling the van der Plank equation, however the physical interpretation of the equation as representing life

stages of the disease is not exploited. The method of implementing plant growth is also at variance from the model described here in that, while a very similar basic growth equation is used the manner in which disease affects the plant's growth is different, and the yield loss is based directly on the total size of the plant at harvest time rather than on the generally used convolution function approach described above. The submodels used to describe plant interaction in the Trenbath model are also more complex than the rather simple one to be introduced here. Partly because of this and partly because the model is continuous, requiring the use of a differential equation solving routine, the Trenbath model takes more than ten times as much computation per epidemic as the model used here (B.R. Trenbath, personal communication), which is a severe disadvantage in the computation of optimal controls.

The basic property of a plant disease's effect on the plant is the removal from normal function of part of the plant tissue. The second property of the behaviour of the common biotrophic pathogens is that they compete with their hosts for the products of photosynthesis. It does not seem possible to model the second property (and more complex behaviour) without some reference to the process of photosynthesis, and this reference would make the model too complex to be used for optimal control. The following submodel, therefore, is intended to extend the level of debate in only three ways:

(1) it mimics the growth of a healthy wheat plant

(2) growth is reduced in proportion to the fraction of tissue infected

(3) compensatory growth by less attacked plants is possible.

The basic growth equation used is

$$\frac{d\Delta Y_j(t)}{dt} = k\Delta Y_j(t)(1 - t/T') \qquad \ldots(4.3.13)$$

where k is the growth constant

T is the time of maximum crop size

The curve generated by this equation is a normal curve with its mode at t = T' and standard deviation $\sqrt{(T'/k)}$. Figure (4.1) shows that this curve can be used to mimic the growth of a healthy wheat plant as recorded in the data of Rawson and Hofstra (1969). Since $\Delta Y_j(t)$ is the area of plant tissue not occupied by disease, growth can be slowed but not reversed by the presence of the pathogen.

In order to simulate compensatory growth without assuming any particular model for the interaction of plants a positive feedback is introduced into equation (4.3.13) such that when a cultivar's proportion of the crop exceeds its proportion in the seed sown, its growth rate is increased. Correspondingly, when (as a result of disease) the growth rate of a cultivar falls behind its fellows, the feedback effect intensifies its disadvantage and further suppresses its growth. The feedback signal for cultivar j itself is given by

$$e_j(t) = (y_j(t)/Y(t).u_j))^f \quad (u_j \neq 0) \qquad \ldots(4.13.14)$$

where u_j is the fraction of cultivar J planted

f is the feedback gain

The feedback effect is then applied muliplicatively:-

$$\frac{d\Delta y_j(t)}{dt} = k\, e_j\,(t)\, \Delta y_j(t)(1 - t/T') \qquad \ldots(4.3.15)$$

Once again, this equation is actually implemented numerically as a difference equation with a time step of one day, for greater speed and a slight increase in realism.

It should now be pointed out that where crop growth effects are introduced in a model of a multiline the standard methods of predicting

yield loss do not apply. In particular, they offer no way of predicting increased yield as a result of compensatory growth by the less affected plants. In the multiline different plants may have different absolute sizes as well as different disease covers and a plant with a 20% cover that has been competing with plants with a 30% cover is likely to yield more than a plant with 20% cover that has been competing with similar plants. There seems to be no element of the current level of debate which covers this question. The approach that will be used here is to take the yield as calculated by the conventional method and then multiply it by a factor resembling the one already used to initiate compensatory growth. Thus

$$z_j = Z_0 u_j \, (100 - x(T_h)) \, (\phi_j)^g / 100 \qquad \ldots(4.3.16)$$

where z_j is the yield of line j m^{-2}

Z_0 is the yield of a clean crop

u_j is the proportion of line j sown

$x(T_h)$ is, as before, the percentage yield loss

$$\phi_j = \frac{Y_j(T_h)}{\sum_j Y_j(T_h).u_j} \qquad \ldots(4.3.17)$$

and g is the gain factor for yield compensation. If g = 0, there is no yield compensation.

Figure 4.2. shows the yield according to the model of susceptible and resistant plants in a replacement series ranging from fully resistant to fully susceptible. In each mixture a quantity of spores sufficient to depress yield by 92% in a fully susceptible crop is applied 40 days before the critical date at which disease loss is determined. At low proportions of the susceptible plant the disease is not able to multiply enough to affect yield at all. As the proportion increases the disease loss gradually increases. Compensatory growth by the resistant plant then becomes noticeable in the concavity

of its yield function, which shows that it is yielding more than in proportion to its representation in the seed sown. However compensatory growth is not (and cannot be in this model) the explanation for the insensitivity of multiline yield to disease at low proportions of the susceptible components. The insensitivity is caused by the inability of the pathogen population to increase under these circumstances.

It is not supposed that the submodel described in this section is a fully realistic simulation of a diseased mixed crop stand. Instead it is conceived as a test to see if the conclusions of a basic model using equations (4.3.4) and (4.3.12) are, as it were, locally stable to increases in realism of this particular kind. Empirical studies may suggest that this, or a similar, approach is a robust enough mimic of actual multiline behaviour to have further use, despite its lack of realism, but this cannot be judged at the moment.

4.4. The Time Between Harvest and Sowing

If information which suggests how to simulate the growth of a diseased multiline is sparse, information that suggests how to model the interface between crop cycles is almost nonexistant: Kiyosawa (1972) seems to be the first to have considered 'overwintering' from a mathematical viewpoint, and he used a simple constant survival probability for propagules between cycles.

Ideally, the interface between crop cycles should be considered as almost as important as the crop cycle itself. It would be pointless to achieve an accuracy of say +/- 10 per cent for within cycle dynamics while having only an accuracy of +/- 90 per cent for between-cycle dynamics. Yet it is obviously much harder to measure the between-

season dynamics. No solution is offered here except to use a very simple model of the interface, whose parameters might conceivably be measured in practice. The model is written as though survival between seasons is a matter of the survival and active reproduction of some of the spores present on the day of harvest on some alternate host, perhaps wild grasses. Kiyosawa's (1972) study applies more to a situation in which spores overwinter passively on crop residues, as in the case of *Helminthosporium maydis* (Sumner and Littrell, 1974). The choice is a largely arbitrary one which was influenced here by the situation of wheat stem rust in Australia.

The basic relationship used here is a hyperbolic one, intended to suggest that small numbers of spores may even increase on the wild host but that large numbers of spores saturate the available host at a level far below that at which the crop would be saturated (about 5% of saturation level). Accordingly the relationship between the output of season n and the input of season n+1 is given by

$$s(in)_{j(n+1)} = \frac{s(out)_{j(n)} K_1}{K_2 + s(out)_{j(n)}} \qquad \ldots(4.4.1)$$

where K_1 is the maximum possible input to the next season in spores m^{-2}

K_2 is the output at which half K_1 is achieved.

4.5. The Introduction of New Races

The occurrence of new races of pathogens is the origin of all problems of resistance breeding and consequently deserves careful consideration in modelling. However the information on which to base quantitative estimates of model parameters is almost totally lacking, so that what is proposed in this section, while simple in concept, is subject to wide uncertainties, both about parameter values and about the appropriateness of the approach.

The probability of occurrence of a new pathogen race is the sum of two probabilities:

(1) The probability that the new race will arise by mutation and recombination within the system

(2) The probability that the new race will be introduced from outside the system.

4.5.1. Mutation

Mutation rates for plant pathogens are hard to estimate either in nature or in experimental situations. In nature it is impossible to observe the first appearance of the new race, or to estimate accurately the size of its parent population and in experiments the number of spores observed is usually too low (an exception occurs with the rice blast fungus (Ou, 1971; Kiyosawa, 1976)) for estimates to be made. Since it is now becoming possible to culture what were once regarded as obligate parasitic pathogens, such as the rusts, on chemically defined media (Bose and Shaw, 1974), it may be possible to obtain better estimates in the future.

The probability that a new race will arise by mutation depends not only on the mutation rate but on the number of opportunities to mutate, which in turn depends on the infection level and the total area of the crop. Thus although the system is being modelled as "lumped" (section 3.5.2.3) it must be scaled to represent an actual area, and the multiline problem is essentially a problem of large areas. Now if the area is large, in any but the lightest epidemic the number of spores is likely to be so large that

$$N\mu \gg 1 \qquad \ldots(4.5.1)$$

where N is the total number of spores produced through the whole crop cycle.

μ is the mutation rate

Under these circumstances the expected number of mutations

$$E(s') >> 1 \qquad \ldots(4.5.2)$$

and the standard deviation of this number

$$\sigma(s') \sim \sqrt{N\mu} \qquad \ldots(4.5.3)$$

is so small compared to the expected number that the occurrence of a mutation is virtually certain. Accordingly in this study mutation is treated as the steady introduction of the expected number of propagules of a new race, which is reasonable since the area being modelled is large (say 10^{10} m^2) and the initial conditions considered are either zero disease or (as will be seen in Chapter 6) represent multiples of 10^{12} spores.

4.5.2. Irruptions

Thus the essentially stochastic element of the behaviour of an extensive crop/pathogen system seems to be provided, in the absence of environmental variation, by the irruption of new races from outside the system. Quantitative information on how to model this process is lacking, except for the hint given by the often repeated statement (e.g. van der Plank, 1968) that the life span of a new major gene in small grain crops is of the order of 5-10 years. This hint gives a basis for estimating the probability of irruptions but not for estimating their distributions of size and race composition.

In practice it will usually be impossible to tell whether the suddenly observed presence of a new race in a particular part of a distributed crop is in fact a large scale irruption from outside the system or is simply the rapid spread that will occur under favourable circumstances from a small unnoticed focus. The important feature of either event, from the point of view of the administrator of a control policy is the suddenness with which serious damage is done to the crop.

Accordingly the procedure that is adopted in this study is an attempt to preserve the significance of irruptions as a source of unpredictable crop damage while making the simplest possible assumptions about the size and composition of the irruptions. It will be assumed that races appear as a fixed quantity of spores, sufficient to cause a noticeable yield loss in a fully susceptible and otherwise clean crop. The appearance of different races will be treated as independent events, each with the same probability, P_{ext}. Thus if n genotypes are considered, it will be assumed that there are 2^n possible irruptions defined by the presence or absence of the fixed quantity of each race. The expected cost of a control action will then be found by simulating the effects of each type of irruption, multiplying each cost by its binomial probability and adding.

4.6. Modelling the Cost of an Epidemic

4.6.1. An Agribusiness Criterion

The first criterion to be modelled in this study is intended to represent the decision-making of an agricultural industry in an advanced country. The industry is managed by an organisation, public or private, which is interested in maximising the present value of the expected cash return, which it does by attempting to avoid shortages and gluts. The factors that this organisation must consider are:-

(1) the yield per unit area, as affected by the pathogen

(2) the area to be sown

(3) the fraction of the product to be stored from cycle to cycle and the fraction to be sold

(4) the price for which each unit of the product will sell

(5) the costs of production.

For the sake of simplicity and in order not to divert too much attention from the biological events, changes in the parameters of the rules governing (4) and (5) will be neglected.

The basis of the calculation of cash return is a simple price/demand curve of a type found in many economics textbooks (e.g. Leftwich, 1970).

$$p = p_0 e^{-bQ} \qquad ...(4.6.1)$$

where p is the price per unit

and Q is the number of units sold (assumed to be all those put on sale)

If there were no cost associated with production then the return for one cycle would be given by

$$R = Qp \qquad ...(4.6.2)$$

$$= Qp_0 e^{-bQ} \qquad ...(4.6.3)$$

which has a minimum at $Q_{opt} = 1/b$, the optimum amount to produce for a one year policy in the cost-free situation.

But

$$Q = q(yA + s) \qquad ...(4.6.4)$$

where q is the fraction put on the market, $0 <= q <= 1$

y is the yield per unit area

A is the area planted to the crop, $0 <= A <= 1$

s is the amount in storage

So

$$R = q(yA + s)p_0 e^{-bq(yA + s)} \qquad ...(4.6.5)$$

If there were no costs associated with R, then all planting policies would be equivalent unless they made it impossible to offer Q_{opt} on the market. So in order to increase realism somewhat and constrain the optimal control to distinguish between possible policies, the following simple cost function is introduced.

Suppose that the cost of producing Q has two parts; a fixed cost c_0 which might refer to maintenance of farms and farmers, and a variable part k_3A, which increases linearly with the amount of land cultivated and might refer to tractor fuel and fertiliser. The return equation is now:

$$R = Qp - c_0 - k_3A \qquad ...(4.6.6)$$

It is now necessary to consider any costs associated with different genotypic ratios that might be used in the crop. Plainly, there would be costs involved in setting up a multiline program. The breeding of multilines has been called "agronomically conservative" (Hooker, 1967) because any improvement (other than the incorporation of a new resistance gene) must be made in a number of lines simultaneously with a corresponding extra cost. There might also need to be expansion of pathogen sampling programs in order to get more accurate information on the composition of the pathogen population. There would also be costs of the continued running of such a program but they would be almost independent of the particular decisions taken in each cycle, provided that multilines were used so frequently that setting up and closing down cost of the multiline facilities need not be taken into account as recurrent costs.

Nevertheless it will be assumed that these costs are negligible compared to the total cost (assumed constant) of any breeding program. To do otherwise would introduce a discontinuity into R at values of y produced by pure lines, and this discontinuity might bring special mathematical difficulties which are best left out of an initial exploration. Thus all costs associated with plant breeding will be absorbed into the fixed cost c_0. Whether this is justifiable must be decided in each practical case.

Three important points must be made about this model. The first is that the introduction of storage of the product between cycles has augmented the state vector. As well as consisting of $x_1 \ldots x_n$ representing pathogen abundances, the vector must now contain an $x_0 = s$ representing the amount stored at the decision point. The component x_0 is likely to affect the choice of policy. For example, if the secular trend of the pathogen is ppsitive, and there is the possibility of storage, it may prove optimal to produce some of next year's crop this year, because even if the price is the same next year, the cost of production will be higher.

The second important point is that the vector of control variables has also been augmented by the addition of q and A. The introduction of q is a straightforward consequence of there sometimes being a maximum of R attainable within the constraints on the system so that a decision must be made on how much to sell. The role of A is more subtle in that it restricts the size of the set of candidates for the optimal policy by requiring a minimisation of $k_1 A$. This adds both realism and interest to the problem, for in real life it is unlikely that the plant breeder would have the luxury of deciding on his optimal policy by selecting among equals with the toss of a coin.

The third important point is that one of the additional control variables chosen here will have a direct effect on system dynamics. The choice of area planted will tend to multiply or diminish the total number of spores produced, though not, in the lumped system, the number of spores produced per m^2. In the extreme case where no crop at all is planted for a year, even the superrace will die out, though it will still return eventually from outside the system.

In a real, distributed system, a reduction in area planted might take the form either of a reduced density of fields of the crop in each locality, or of the absence of the crop from some localities. If this "thinning out" were undertaken on a small enough mosaic, then it would have the effect of altering the local dynamics of the epidemic because spores would be frequently lost to other species of plants. On a larger scale, the local dynamics of epidemics would be unchanged where they occurred, but they would not occur everywhere. In the lumped system it is only this latter assumption that can be modelled without introducing many extra assumptions. Accordingly, when the agribusiness criterion is being used, the state output will be multiplied by the fraction of total area planted.

4.6.2. A Subsistence Farmer's Criterion

The model discussed in the last section applies to a world where Q_{opt} can be attained most or all of the time. If this were not so we might expect to be able to model the situation by setting A always equal to $A_{max} = 1$ and optimise by maximising productivity in all seasons. However in a complex market economy it is difficult to assign a definite limit to the area that can be planted, since it is not usually an absolute physical limit but one determined by the relative profitability of other enterprises, the cost of substitutable resources and other economic factors such as changing consumer preferences (even for basic foodstuffs).

So for a clearer contrast to the rich economy the model in this section is designed to represent a subsistence economy of an extreme type, exaggerating the tendencies described in Clark and Haswell (1967). This new economy produces a single product, which will be called a balanced diet, equivalent perhaps to the "grain equivalent" described

in Clark and Haswell (1967, p53). This product is produced from a fixed amount of land, which might be thought of as a crowded island like Java, with one input - labour, which consumes the balanced diet in order to produce. There is never a surplus sufficient to trade and so the balanced diet is the only commodity in the economy. There is some resemblance between this "economy" and a global society enduring great scarcities because of the loss of its reserves of resistance genes, but the more complex the society, the harder it is to define the "balanced diet" for that society except through a market or a planning system using monetary concepts.

The basic assumption of the model is that the marginal productivity of labour falls as the labour input rises, so that there will be a maximum net production, equal to net consumption, at the point where the consumption needed for each additional amount of labour equals the resulting increase in production of the balanced diet. This maximum might be at a negative or positive value of net consumption and at a zero or positive value of the labour input. It is also assumed that, as in the previous model, the cost of implementing each control decision is part of the fixed costs of production, which in this case correspond to the "resting" metabolism of the population.

Since this is not a money economy it is not obvious how to define the cost of a policy. As various case histories in Clark and Haswell (1967) show, subsistence farmers sometimes have utility preferences that surprise western economists. However, the subjective utility of a policy to one of the hypothetical subsistence farmers will probably depend on two factors; the amount of labour that has to be done and the size of the consumption deficit. It is to be expected that these factors will have to be traded off against each other, this

year's work being traded off against next year's hunger, for example. Yet the two factors are at first sight incommensurable.

One consistent solution is to suggest that the optimum point for the farmer occurs when he is neither hungry nor overweight (fat, the stored balanced diet, has to be worked for), and to make an equal trade-off on either side of this ideal body weight by a quadratic function

$$C = k_2 x_0^2 \qquad \ldots(4.6.7)$$

where C is the cost

x_0 is an extra state variable, the deviation from ideal body weight

k_2 is a constant.

It is now necessary to suggest some specific forms for the rule governing x_0. In the first place

$$x_{0,n+1} = x_{0,n} + f - e \qquad \ldots(4.6.8)$$

where $x_{0,n}$ is the deviation from ideal weight in cycle n

where f is balanced diet produced

e is balanced diet consumed.

A simple law to express the dependence of f on labour might be

$$f = \frac{k_3 w y}{k_4 + w} \qquad \ldots(4.6.9)$$

where w is the labour input

y is, as before, yield per unit area as determined by the pathogen

k_3, k_4 are constants.

While a simple law connecting e and w is

$$e = c_1 + k_5 w \qquad \ldots(4.6.10)$$

where c_1 is a fixed cost equal to resting metabolism

k_7 is a constant.

The labour input to this model is at first sight unusual because it does not (as the area variable in the first model) affect the pathogen output. However some realism can be given even to this simple model if it is supposed that w is work done in such functions as scaring seed eating birds, irrigating the crop at the time of grain maturation or any other task which affects the yield but has little effect on the leaf area.

Two modifications to this basic model have proved necessary in practice. The first is that a small cost has been put on the use of added spores, in order to induce the control algorithm, for the sake of clarity, to remove unnecessary spores. The reason for doing this was that it was found that if the trial policy added spores, and the algorithm made a swift decision to plant a superline, the remaining added spores were undetectable by the system. The second alteration was more important, because the model as described above allows the possibility, if the superrace is vigorous enough, of a regress to infinite negative bodyweight. It is clear that some high cost should be set on death at the minimum tolerable bodyweight, but not clear how this might be implemented. The decision was made to set a lower bound on the bodyweight and assign a cost to any policy that results in a bodyweight less than this bound, equal to the cost at the bound plus a multiple of the transgression. This penalty-function approach to bounding the cost can only be a poor approximation to the real social situation being modelled but the situation itself is not unrealistic (Turnbull, 1973).

Similar points to those made about the agribusiness model can be made about the subsistence model. The state vector has been augmented because another system variable is being carried over from stage to stage. The control vector has also been augmented, by the addition of w, in the interests of realism, but also with the object of making a unique optimum more likely.

4.6.3. Utility and Risk

It may be argued that the cost functions described above do not represent the inconvenience and loss that may be caused by fluctuations in crop yield from cycle to cycle or by variance of yield within cycles. This is equivalent to saying (Luce and Raiffa, 1957) that the discounted sequence of cost functions does not define a linear utility scale. It is of interest to consider methods of avoiding this reproach. There seem to be two different classes of method. One is to treat the choice of an optimal control as a game between two players, the plant breeder and the pathogen. The other is to augment the criterion by adding a measure of the variability of system output, and minimising the value of the sum. We will consider these methods in order.

4.6.3.1. The Multiline Problem as a Game Against Nature

Pesek (1974) has treated a simplified version of the multiline problem as a "game against nature" (Luce and Raiffa, 1957). In his formulation the problem is a zero-sum game played between the plant breeder and the disease with the payoff being the amount of disease and the solution is held to be a minimax solution: the minimum of the maximum amount of disease. This approach is unrealistic in a number of ways. In the first place, the multiline problem is not a zero-sum game because the human objective is only indirectly to reduce disease,

being in fact to minimise a cost. In the second place, and more fundamentally, the question is whether to treat the problem as a game at all.

The reason for questioning the games theory treatment is that the pathogen population is not like one (or even a number) of players each with a range of moves. Instead, each race of the pathogen has just one (mixed) strategy which is determined before the game begins by the dynamics of mutation, recombination and disease spread. The description of the multiline problem as a "game against nature" is thus somewhat misleading since the use of a minimax approach to finding the "optimum" policy is not a "rational" response in the face of a "rational" adversary. Instead if a plant breeder chooses to use a minimax approach he is in effect saying that the cost of an epidemic (as calculated by equations such as those above) does not represent his own utility scale linearly, and he places an absolute priority on avoiding the highest costs that may occur. This approach will be avoided in this study because, as was said in section (3.5.1) it seems to be excessively pessimistic in the absence of further information about agribusiness and subsistence farmers' utility scales.

4.6.3.2. Placing a Cost on Variability

In at least one respect the cost functions described above do take account of variability. Both provide storage for the crop, which is the most ancient and direct way of smoothing out fluctuations in yield. The agribusiness cost function, however, places no cost on this storage, only an upper bound, while the square-law form of the subsistence cost function tends to act as a suppressor of variability in storage.

The use of measures of variability in making a criterion more responsive to fluctuations was prefigured in section (3.3.1) by the suggestion that the variance of yield be minimised. In this section measures of the variability of the cost functions already defined will be considered. One measure that is easily obtained is the variance of the one-stage cost k(u(i),i), since its probability distribution is constructed in the course of calculating its expected value. Similarly, the variance of $E(V(\pi(i),i))$ can be found. However, these variances are measures of within-cycle, or, equivalently, within-state variability, while between-cycle variability, which is of prime interest, is omitted.

To see how between-cycle variance can be handled in the optimal control setting, consider an n-stage decision process. In stage (n-1), the decision that minimises between-stage variance for the whole process will minimise

$$(k(x_{n-1}, u_{n-1}) - K_n)^2 = (k(x_{n-1}, u_{n-1}) - ((n-1)K_{n-1} + k(x_{n-1}, u_{n-1}))/n)^2 \quad \ldots(4.6.11)$$

where k is the within stage mean of the cost

K_n is the grand n-stage mean

In this form the minimisation of the variance is treated as a dynamic programming problem at the expense of augmenting the state vector to include the overall mean cost at the ith stage K_i, and the stage number i. This augmentation markedly increases computation, but even apart from this the method cannot be extended to an infinite policy horizon because i would increase without bound. This is equivalent to saying that no individual decision in an infinite policy can affect the variance of the cost. An alternative method, which does not calculate a true variance but which can be applied to an infinite policy horizon

is to define a new cost function $C(\pi(i),i)$ so that

$$V(\pi(i),i) = (C(\pi(i),i) + V(\pi(X_1),X_1) \qquad ...(4.6.12)$$

where $$C(\pi(i),i) = \alpha k(i,u(i)) + \beta var(k(i,u(i)) + \gamma(k(i,u(i)) - k(X_1,u(X_1)))^2 \qquad ...(4.6.13)$$

and α, β, γ are the weighting coefficients for mean cost, within cycle variance and between cycle variance respectively.

This procedure adds a discounted measure of variability (rather than a true variance), to the cost function V, which would represent both within and between cycle variability. The amount of computation involved in each function evaluation would of course be roughly doubled through the need to compute $k(X_1,u(X_1))$ as well, and it might also be necessary to minimise over two sets of control variables at once (u(i) and $u(X_1)$).

These uncertainties have contributed to the decision not to attempt to include additional expressions for the cost of variability in the optimality criteria of this study. The main reason has been that their inclusion would complicate the interpretation of the results of an investigatory study such as this. The confusion would arise, firstly, because of the addition of arbitrary α, β and γ to the other model parameters, but mainly because in all versions of the model that have been tried reduced variability is most easily achieved by increasing the disease level, so that the stochastic irruptions have little effect.

4.7. Uncertainty and Robust Conclusions

Section (4.6.3) has dealt with finding an optimal control that represents rational behaviour under the conditions of risk represented by the (known) probabilities of new races arising. In this section the impact of events whose probabilities cannot be measured or to which probabilitiers cannot be assigned. In the context of simulation modelling, the first type of uncertainty occurs as a result of it being impossible to measure system parameters (including probabilities) without error. The second type occurs as a result of not knowing what features of the functional forms used in the model are inessential or wrong.

About the second type of uncertainty nothing can be done in general. Sometimes it is possible to base a qualitative argument upon insight or experience with numerical work and derive a result that is not dependent on these particular forms. To a limited extent this was done in the theorems of Chapter 3. The first type of uncertainty may, however, be combatted to some extent. What technique is used will depend on whether the modeller's task is to decide on an optimal control in practice or to undertake a theoretical analysis like this study.

4.7.1. The Practical Case

If the uncertainty to be combatted is impeding the choice of a practical optimal control, what is required is a means of choosing parameter values from within the range of uncertainty such that the calculated expected value of the cost corresponds to its true expected value. If information is available, as a result of experiment, which gives the probability distribution of the different parameter values, this "uncertainty" can be treated as a form of risk, from a new

stochastic input that affects parameter values, and the expected cost can be calculated as before. If all that is available are upper and lower bounds on parameter values, then it is harder to decide on what is "rational" behaviour. A clear exposition of the difficulties of deciding on rational behaviour under uncertainty can be found in Luce and Raiffa (1957): an indication of the current state of debate can be found in Balch and Fishburn (1974). One possible approach is again a minimax in which the control is chosen so as to minimise the maximum value of the expected cost in relation to possible parameter values.

$$V(\pi^*(i),i) = \min_{\pi} \max_{p} E_i^{\pi} (V(\pi(i),i)) \qquad ...(4.7.1)$$

where p stands for the set of possible parameter values.

The use of the minimax procedure to fit models to data is a known alternative to the familiar least squares procedure (Demyanov and Malozemov, 1974). The use of such a method when the criterion of goodness of fit is also a cost function, as in the case of the method suggested above, can perhaps be thought of as a "game against nature" with more justification than in the suggestion discussed in section (4.5.3.1) because here Nature's move is unknown, though bounded. However it is still probably better to think of a minimax choice of model parameters as a modeller's conservative approach to his own ignorance. The real application of the game approach seems to come in theoretical studies and is discussed in the next section.

4.7.2. Strengthening a Theoretical Study

In the case where uncertainty must be combatted in order to establish the robustness of a theoretical result, slightly different questions are involved. The task is now to see what effect a different set of parameter values might have on the conclusion, in order to

satisfy a critical audience. The technique which is usually used is that called sensitivity analysis, in which each parameter is independently varied, often by +/- 10%, and the effects on the results of simulation observed. Specific conclusions are rarely drawn from these observations, unless some kind of "catastrophe" (Thom, 1975) is observed in model behaviour. As is pointed out by Shoemaker (1973b), sensitivity analysis is useful in that it gives an indication of which parameters are likely to have the largest effect on model output and are thus in need of the most intense empirical study. However sensitivity analysis has two important drawbacks:

(1) As the number of model parameters increases it becomes impossible to simulate all possible combinations of +/- 10% variations in each parameter, since there are 3^n of them where n is the number of parameters. Consequently, sensitivity analysis must often be restricted to a subset of all independent variations, usually the 2n one-at-a-time variations. This restriction excludes the possibility of observing interactions between parameter values.

(2) Sensitivity analysis is not at all global: it gives little indication of what effects may occur outside the 10% range. For a practical study, where experimental work has already confined possible parameter values within a range in which no qualitatively surprising results are expected this objection may not be too serious. For a theoretical study such as this one, which is undertaken in almost complete ignorance about many important parameters, a more global approach is desirable.

A possible alternative to sensitivity analysis is to regard the control task as a two-person game in which the payoff to the second player is a measure of disagreement with the null hypothesis. The first player is able to make choices of the control variables so as to minimise the expected cost. The second can choose parameter values (within agreed limits) so as to act as a voice of caution in the treatment of any new finding. It is the deliberately critical spirit of scientific enquiry that lends the game-like aspect to this procedure. It is not, in this case, irrational to think of a second player whose interest lies in finding parameter values that tend to contradict any novel finding of a study: such players actually exist.

The payoff function in this game may take various forms, depending on the manner in which the robustness of the study is being tested. Two forms of test and their corresponding payoffs are:-

(1) The robustness of the result is being tested by the suggestion that the parameter values represent an optimistic view of the real world, in the sense that real costs might be higher.
In this case the payoff function is the ordinary cost of the control, and the game is zero-sum.

(2) The robustness of the result is being tested by the suggestions that the parameter values are for some (perhaps unspecified) reason conducive to the rejection of the null hypothesis about the choice of control. In this case the payoff to the second player should be a measure of the divergence of the optimal control variables from the values prescribed by the null hypothesis.

These two payoff functions lead to radically different amounts of calculation needed to solve the game. The first, zero-sum game

simply requires the alternation of minimising and maximising steps in the calculation of one optimal control. By contrast, in the non-zero-sum case, each function evaluation corresponds to finding a complete optimal control. Accordingly in this study, because of time limitations, no further mention will be made of the second type of game, although its implementation is undoubtedly a desirable objective. Instead some attention will be paid to the zero-sum game, which is perhaps additionally justified by its resemblance to the procedure of the practical decision maker under conditions of uncertainty mentioned in section (4.7.1).

The scope and value of the results of such a game depend on the behaviour of the functions involved. There is an important theorem in games theory (Owen, 1964) according to which, if the payoff function of the zero-sum game is a convex function of the control variables for each set of parameter values, and a concave function of the parameters for each set of control variables, then a minimax solution exists. Finding such a solution would be interpreted to mean that there was no set of parameter values that could cause the given optimal control to yield a higher cost. If the null hypothesis about the control variables were still rejected in this situation, then at least the rejection could not be regarded as an artefact of an optimistic world-view.

If the conditions of the theorem do not hold, then the implication of the minimax process becomes correspondingly less global and may be restricted to ranges of parameter values around the initially assumed values but not covering the whole range of uncertainty. Even so, this would be a more definite conclusion than can be usually be drawn from sensitivity analysis. The final advantage of this process over sensitivity analysis is that the computation required to establish the game solution will increase roughly as n^2 rather than 3^n as in a

complete sensitivity analysis, so that above a break-even point peculiar to each study, the minimax procedure will be easier. In Chapter 6 a simple minimax study with some theoretical consequences will be introduced.

4.8. The Number and Type of Genes to be Introduced by the Plant Breeder

In the rest of this chapter no mention has been made of parameter values, since these will in several cases be varied in the experiments of Chapter 6. However there is one important "parameter" common to all the modelling done in this study which will be discussed here: the number of different genotypes of disease race and plant line to be considered.

In one sense, as soon as multilines are considered, it becomes impossible to consider introducing only one gene, because the multiline strategist has at his disposal, in theory, all the genes that have been incorporated in the crop in the past. Nonetheless most multiline theorists and producers have not envisaged using the full number of genotypes. Both the Iowa oat multilines (Browning and Frey, 1969) and the CIMMYT wheat multilines (CIMMYT, 1974) have usually consisted of about ten to fifteen lines each containing a major gene. The reasons for this choice have not been stated directly though the complexity of handing large numbers of lines presumably played a part.

Van der Plank (1968) theorised that since (in his view) the utility of multilines depends on the presence of stabilising selection, and stabilising selection is only reliably enforced by what he called "strong" genes (section 2.3.4.1), the number of lines in a multiline should correspond to the number of "strong" genes available. In these

lines the strong gene should be used one at a time, because their multiple use in a single line would reduce the action of stabilising selection against the superrace (it would have fewer unnecessary genes for virulence in such a multiline). Whether these arguments are valid or not, their usefulness is limited to cases where stabilising selection is enforced by a sufficient number of strong genes for resistance.

The attitude adopted in this study differs from both the practical and the theoretical positions described above, for reasons which are mainly to do with computational limitations but which have a sound theoretical basis as well. The model will deal with the introduction of two dominant major genes, A and B to a crop. Each gene will be considered to confer immunity to races that are not of genotypes aa and bb at the respective loci. Of the four possible genotypes only Ab, aB and AB (and their corresponding pathogen races) will be considered.

The computational reason for choosing such a simple form of multiline for study is what is known as the "curse of dimensionality" (Bellman, 1957). The amount of computing that is necessary to find an optimal control increases at least as the number of initial conditions considered, which is L^s, where L is the number of levels of each state variable considered and s is the number of state variables. Thus in order to get a reasonably sensitive response to different levels of each state variable it is necessary to restrict the number of state variables. In this study four state variables (three disease race abundances + food storage) are studied at five levels, giving 625 initial conditions to be covered. The inclusion of another disease race at the same level of resolution would require 3125 initial conditions.

The theoretical reasons for choosing a simple multiline are based on Theorem 1 of Chapter 3. Since any multiline effect depends on the presence of the simpler races, there must be a sufficient quantity of the simpler lines present to nurture the simpler races. The smaller the fraction of the crop that is susceptible to a given race, the greater the disadvantage to that race and a race virulent on, say, a tenth of a crop is unlikely to have any influence on the growth of the superrace. Consequently a simple multiline is more likely to be optimal than a complex one.

A second argument for the type of simple multiline used here is similar to van der Plank's argument, quoted above, that multiline components should each have only one major resistance gene. What will be argued here, instead, is that there should be no more than two levels of complexity in a multiline (and thus in this case the ab genotype plant should be omitted). The argument is as follows. Consider the Ab and ab subsystem in the absence of aabb and bb diseases. If in the subsystem there are multiline effects then the subraces in the subsystem will tend to retard the output of the superrace of the subsystem - the aa race. But this race should be more important in the main system than the simpler races growing on the ab plant because it can grow on a larger fraction of the crop. Thus the presence of the ab plant tends to work against the race that will be important in any multiline advantage for the main system. Any counteractive effect of this kind of the presence of more than two levels in the multiline would be expected to be slight because of the great disadvantages that would face the simplest races. However this argument suggests that the optimal number of levels is not greater than two. Nevertheless it would be desirable in principle to allow a control algorithm to decide whether a large number of genotypes is desirable in a multiline, though it is impossible to do so here.

Two further points that should be noted about the genotypic system chosen here are that:-

(1) It includes the conventional use of the pure superline as a special case, in line with the analysis of section (1.2.4.3). Thus the algorithm is put in the position of making a direct choice between the conventional practice and the multiline alternative.

(2) If the multiline chosen were any simpler, the simpler line would not be able to receive the benefits, if any, of cross-protection, since there would be no race present that could survive without attacking it.

Nothing has been said so far in this chapter about the simulation of genes other than major genes conferring immunity to their corresponding disease races. However it is easy, within the framework of the model, to implement a number of the different types of resistance discussed in Chapter 2. It is possible to

(1) change the rate of spore production/unit area of races on particular lines

(2) change the area which a spore takes up in colonising a plant

(3) change the parameters of the yield loss equation.

If the rate of spore production of all races in the system is reduced uniformly, this is equivalent to introducing general resistance as defined by Caldwell (1968), that is, resistance which, while not necessarily conferring immunity is not vulnerable to mutation. If, on the other hand, the immune reaction conditioned by the main genes in the model were exchanged for a reaction in which spores of the hitherto nonvirulent races could occupy a limited area, this would be a way of providing a form of general resistance vulnerable to mutation,

like that reported by Simons (1972). Finally, if the parameters of the yield loss equation are varied, this could be used to introduce an amount of tolerance to disease. In Chapter 6 an experiment on the effects of using 'non-vulnerable' general resistance will be reported as representing the opposite pole of breeding technique to the use of unaided major genes.

4.9. Summary

In this chapter three things have been done:

(1) A philosophy for approaching a complex problem such as the multiline problem has been set out. Its basis is the idea that the task of a theoretical investigation of the problem at this stage is to synthesize some of the different ideas and approaches that already exist and to give some direction to experimental work, rather than to attempt to answer the question directly in the absence of much of the necessary information.

(2) A particular technique for undertaking these tasks has been espoused and its advantages and disadvantages set out. The technique is the harnessing of a numerical simulation of the crop/pathogen system to an optimal control algorithm. Its advantages are mainly the advantages of any default option, but in principle at least it can integrate the factors considered by specialists in various fields to be important in the multiline problem. Its disadvantages are the disadvantages of complex numerical work undertaken without an analytical mathematical basis: cumbersome computation and restrictions on generality.

(3) A specific implementation of this technique has been introduced in the form of a simulation model of a crop/pathogen system and its economic setting. This model combines a number of standard functions in a simple fashion so that the synthetic portion of the task referred to in (1) above can be carried out without rendering the model unmanageable for the purposes of optimal control.

FIG 4.1: COMPARISON OF GROWTH MODEL (-----) WITH CV. SUNSET WHEAT (+ + +)

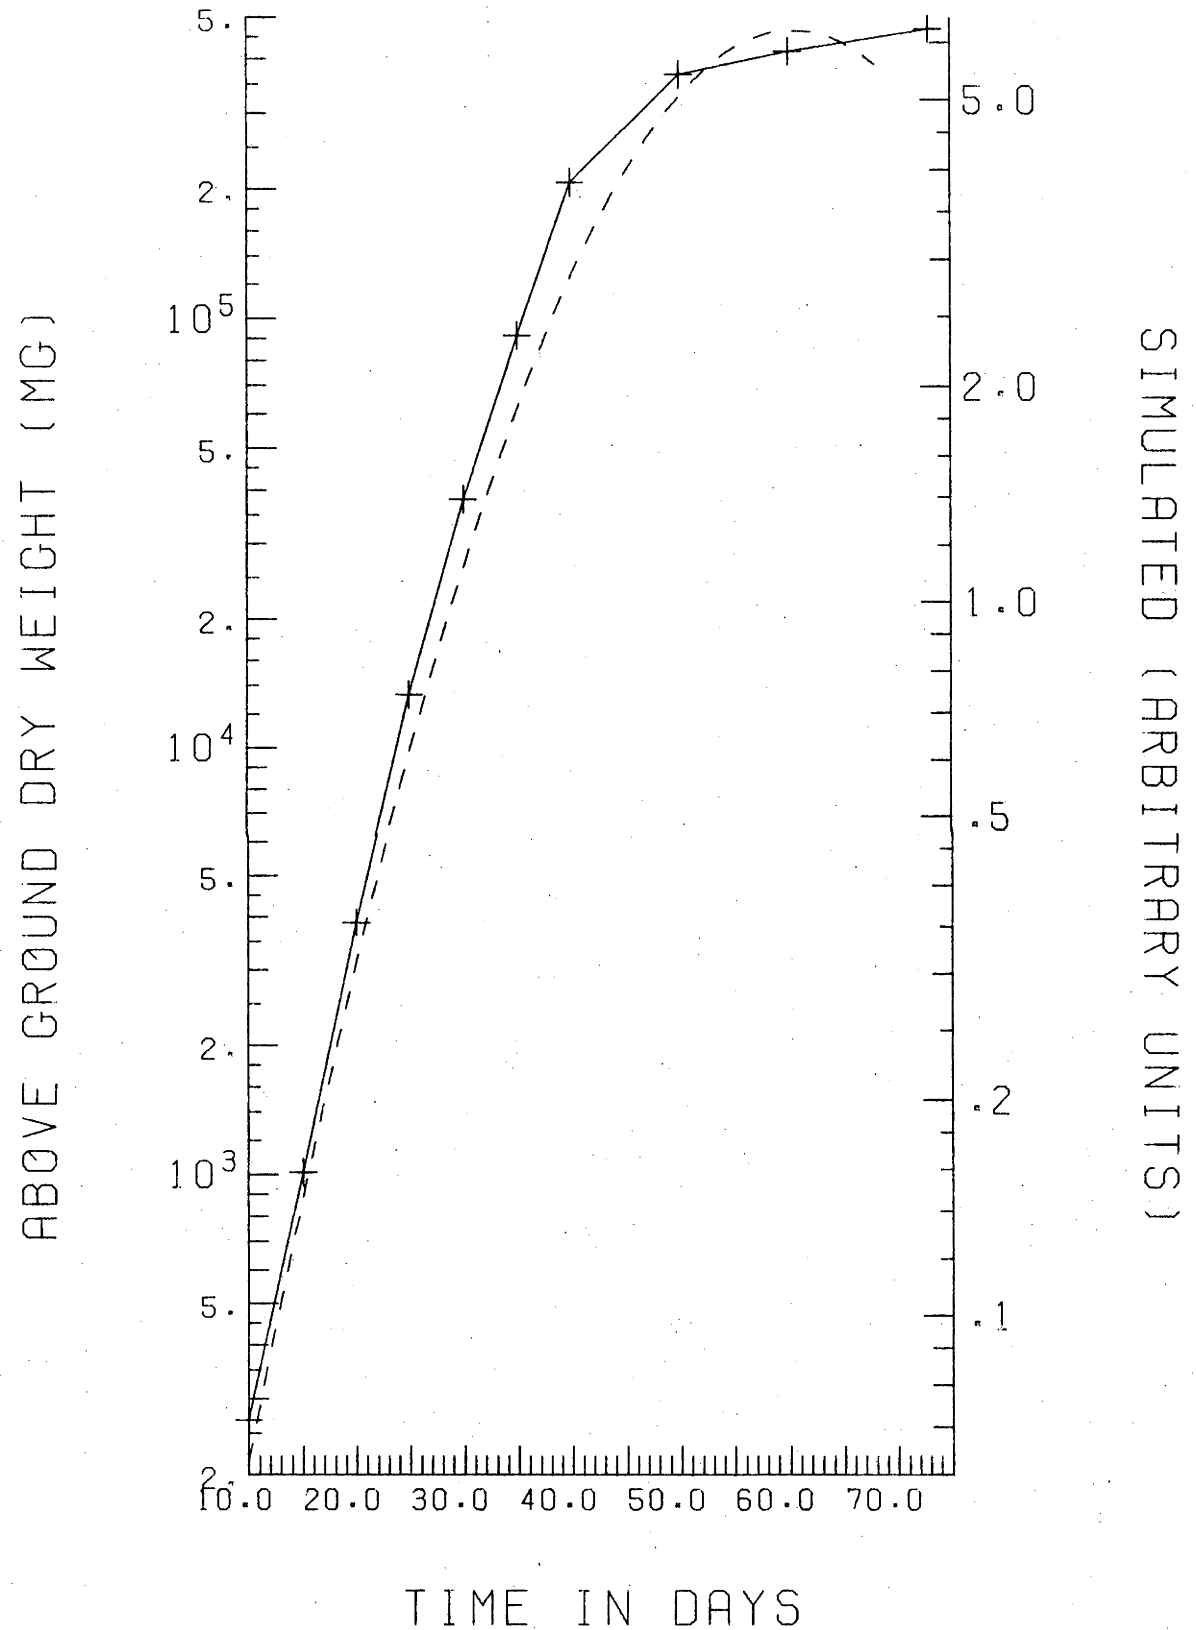

5. Computation of an Optimal Control

5.1. Introduction

The purpose of this chapter is to explain the computational procedure used in this study to find optimal controls for the model set out in Chapter 4, and to document the FORTRAN program by which the model and the optimal control procedure were implemented. Section (5.2) describes the computational procedure, and the subsections of section (5.3) describe the FORTRAN subprograms.

5.2. The Optimal Control Method

The technique that is used here for finding optimal controls for a closed loop, stochastic process is taken from Kushner (1971), who describes it as a "backwards iteration" using a Gauss-Seidel procedure. The iteration begins with a set of guesses at the expected cost of the infinite policy starting from each point on a grid of initial conditions spanning the region in state space which the system can reach, and with a set of guesses at the best control action for each of these initial conditions. During the course of the iteration the guesses at the costs are replaced by successively better approximations to the true costs. At the same time, the initial guesses at the appropriate controls converge on the optimal policy.

The sequence of improving approximations to the true cost of a policy is constructed in the following manner. For each grid point i, $1 \leq i \leq N$, on the nth iteration,

$$C^n(i) = \min_u \rho E_i^u(k(i,u) + C_t^{n-1}(X_1)) \qquad \ldots(5.2.1)$$

where $C^n(i)$ is the nth estimate of the cost of the policy starting from initial condition i. Now $E_i^u(C_t^{n-1}(X_1))$ is made up of estimates of

the costs of the policies starting from the various possible values of X_1, and the probabilities of occurrence of each of these values. Some of these costs were last estimated on the (n-1)th iteration, and some have already been re-estimated on the nth (current) iteration. Accordingly

$$E_i^u(C_t^{n-1}(X_1)) = \sum_{j=1}^{i-1} p_{ij}(u)C^n(j) + \sum_{j=i}^{N} p_{ij}(u)C^{n-1}(j) \quad ...(5.2.2)$$

where $p_{ij}(u)$ is the probability of state j given initial condition i and control action u.

The successive iteration over the grid of N initial conditions is the Gauss-Seidel procedure. In order to understand its significance, it may be useful to consider the simple case where there is only one grid point and only one possible control action. Then if the initial guess at the cost is C^0;

$$C^1(i) = \rho(k(i,u) + C^0)$$

$$C^2(i) = \rho(k(i,u) + C^1)$$

$$= \rho(k(i,u) + \rho(k(i,u) + C^0))$$

.....

$$C^n(i) = k(i,u)\ (\rho + \rho^2 + \rho^3 + ... + \rho^n) + \rho^n C^0$$

$$= k(i,u)\rho(1 - \rho^n)/(1 - \rho) + \rho^n C^0 \quad ...(5.2.3)$$

Thus the Gauss-Seidel procedure in this case is equivalent to adding the geometric series of discounted costs, and, whatever the initial guess C^0, as $n \to \infty$, its effect on the estimate of the true cost V(i) approaches zero.

Theorem 3 of Chapter 5 of Kushner (1971) proves that the more general procedure of equations (5.2.1) and (5.2.2) does in fact converge to the true optimal cost

$$V(\pi^*(i),i) = \lim_{n\to\infty} C^n(i) \qquad \ldots(5.2.4)$$

and its corresponding optimal control policy $\pi^*(i)$, whatever the initial set of guesses $C^0(i)$ and $u^0(i)$.

The calculation involved in equation (5.2.2) needs some further explanation, because, as X is a vector of continuous variables, its components do not usually fall on any of the grid points. Hence, in order to make the estimates of $C^{n-1}(j)$ it is necessary to interpolate between neighbouring points of the grid around X_1. It is also necessary to ensure that X_1 cannot lie outside the grid of points, either by choosing the grid in such a way that it covers all areas of state space accessible to the system, or by limiting X_1 to within the gridded region by some more or less *ad hoc* method (see section (4.5.2)). The interpolation method used here is piecewise linear interpolation (Schultz, 1973), which is easy to understand in one dimension, but less easy to understand in the four-dimensional state space of this study. For a p-dimensional space the interpolation formula is

$$C^{n-1}(X_1) = \sum_{j=1}^{m} \prod_{k=1}^{p} \beta_{kj}(X_1) C^{n-1}(j) \qquad \ldots(5.2.5)$$

where $m = 2^p$ is the number of points in the p-dimensional "box" immediately surrounding X_1

and $$\beta_{kj}(X_1) = 1 - |j(k) - X_1(k)| / w(k) \qquad \ldots(5.2.6)$$

where $j(k)$ is the kth coordinate of point j

$X_1(k)$ is the kth coordinate of point X_1

$w(k)$ is the distance between grid points along the kth dimension of the state space.

Thus $\beta_{kj}(X_1) = 1$ if $X_1 = j$

$= 0$ if X_1 is another grid point on the "box" and varies linearly between 0 and 1 along any edge of the box.

It is important to point out that this process, of necessity, introduces errors to the control process, because the true cost function is in all probability not linear between grid points. Thus wherever the cost function is strongly curved the 'twisted' hyperplane formed by the interpolation procedure will be in error, and the estimate of the cost will be in error. By contrast, when the cost function is flat or changing only slowly, the interpolation will give a better approximation. This problem applies also to the task of using an optimal control once it is calculated because,to find the correct control action in response to a given state, it will be necessary to interpolate in some manner between the optimal controls calculated for neighbouring grid points. This type of error can be combatted in two ways: by using a finer mesh of grid points, or by developing an empirical non-linear function for the purposes of interpolation. Within the limitations of this study neither approach was possible.

Thus the finding of an optimal control involves a nested series of operations. The outermost operation is the Gauss-Seidel iteration. Within this is the minimisation of the expected cost for each grid point in turn. Within this in turn is the evaluation of the cost for different trial control actions. At the inmost level there is the simulation of the events of each day of the epidemic. The models for the cost function and the epidemic have already been described in Chapter 4. It remains, in this section, to discuss the method by which each minimum cost in the sequence (5.2.1) is found.

5.2.1. The Minimisation of the Expected Cost

The method used in this study to minimise the value of the expected cost at each point on the grid of initial conditions is that of Davidon (1975), which belongs to the class of second-order

optimisation methods (Kowalik and Osborne, 1968; Walsh, 1975). In this class of methods, as in most optimisation methods, a function (in this case the cost function V) is minimised by iteratively evaluating its gradient with respect to the independent variables (in this case the control variables) and then taking a step down the gradient. The procedure then repeats until some preset criterion for stopping is met (for example a gradient less than a certain value). The distinguishing feature of second-order methods is that, at the same time as a step is taken, a store of information about the second derivatives of the function is built up. A PxP matrix, usually called H which is initially often the identity matrix is updated so that it becomes the inverse of the matrix of second derivatives of a quadratic approximation to the function. H is then used to modify the direction of the steps down the gradient and (in some methods, including the one used here) to set their length. In theory these second-order methods have the property of quadratic termination, that is, if the function being minimised is a quadratic in the P variables, its minimum should be found in P steps.

A detailed discussion of the theory lying behind the particular method used here and of its relationship to other second-order methods will be found in Davidon (1975). What will be discussed here are the specific features that led to its choice for this study and a number of modifications that were made to it to adapt it for the special requirements of this study.

5.2.1.1. Reasons for Choosing the Method of Davidon (1975)

Apart from the other virtues claimed for it in the reference, three reasons in particular led to the choice of this method.

(1) The steps actually taken during the course of the minimization can be varied from those specified by the algorithm (to comply with constraints) without affecting the quadratic termination property. The method in which such variations were made will be described in the next section.

(2) The PxP matrix is stored in a factorised form as a PxQ matrix J where Q is the dimension of a subspace of the space of control variables and $H = JJ^T$. This was desirable because in some cases it has been helpful not to use all the control variables, but particularly because the three components of the crop mixture cannot be chosen independently since they must sum to 100%. Thus in this study Q has been at most P-1, with a corresponding saving in gradient evaluations.

(3) Davidon specifically endorses the evaluation of function gradients via a finite difference method, which is, of course, the only way in which gradients of a function based on a simulation can be measured. A distrust of this approach has led to the development of optimisation methods which do not require gradient information (eg. Box, 1965), but Birta (1976) has given support to Davidon's point of view. Birta compared the use of finite difference and explicit gradient evaluations on the optimization of some standard test functions, and concluded that it was more economical of computer time to use finite difference methods than to use an algorithm that does not use gradient information.

5.2.1.2. Modifications of the Method

Three modifications to the published algorithm have been made in this study. The first adapts the recommended step length to the requirements of the constraints on the control variables. The second makes a change in the stopping criterion that was found to be of advantage in improving the sensitivity of the algorithm. The third adjusts the algorithm so that, optimally, it can find a minimax policy.

The modifications are discussed in turn.

(1) All the control variables used in this study take only non-negative values. In addition, most of them have a natural upper limit (e.g. 365.25 days worked per year) except the quantity of pathogen added (which has been given in practice a ceiling equivalent to a heavy epidemic). Consequently, when the proposed step cuts across the constraints it is necessary to shorten it. This is done in the following stages:

(a) the ratio $S = \max_j \Delta u_j/c_j$ is formed ...(5.2.7)

where Δu_j is the proposed step component in the jth direction

c_j is the width of the allowable region between constraints in the jth direction

(b) $S' = 1 + \phi S$ is formed ...(5.2.8)

(c) All x_j are scaled by dividing by S'. This has the effect of leaving small steps unchanged, but restricting large ones to at most $1/\phi$ of the constrained region.

(d) The series of ratios

$$\alpha_{jk} = \frac{u_j + \Delta u_j - d_{jk}}{\Delta u_j} \qquad \ldots(5.2.9)$$

is formed, where

d_{jk} is the lower (k=1) bound of the jth variable or the upper bound (k=2).

(i) If $0 < \alpha_{jk} < 1$, then Δu_j crosses the constraint

(ii) If $\alpha_{jk} < 0$ or > 1, the constraint is not crossed

(iii) If $\alpha_{jk} = 1$, the constraint is active, that is u_j is on the boundary and $u_j + \Delta u_j$ lies outside

(iv) If $\alpha_{jk} = 0$, the new point lies on the boundary.

In practice the activity of a constraint is assessed by $1 - \varepsilon < \alpha_{jk} <= 1$ for $\varepsilon = 10^{-8}$. If cases (ii) or (iv) apply for all j, then the proposed step is admissible and no further action need be taken. If case (i) applies for some of the j, then all the Δu_j are multiplied by $(1 - \max_j \alpha_{jk})$, so that the modified step just touches the nearest boundary. If case (iii) applies for variables other than those relating to the crop composition, then the component of the step in this direction is set to zero. If however one of the crop composition constraints is active a more complex adjustment is necessary, and a simple version of Rosen's (1960) gradient projection method is used. The problem is to keep the step inside the triangular region which defines crop composition

$$\sum_{j=1}^{3} u_j = 1$$

The edges of this region, including the active constraint, have the equations

$$u_i + u_j = 1, \quad i,j = (1,2,3)$$

so that the outward normal is one of (1,0,1), (0,1,1), (1,1,0), and the parallel vectors are (-1,0,1), (0,-1,1),

(-1,1,0). The component of Δu parallel to the constraint is found by taking the scalar product, and only this component is retained. If u is at a vertex, the procedure may have to be repeated for a second active constraint.

(2) The main stopping condition for Davidon's algorithm is that the mean squared value of the components of the gradient falls below a set error bound. In the constrained case, the feasible minimum is not usually at a place where the gradient is small, but at one where a constraint is active. Accordingly, another stopping criterion is needed, and this is supplied by a lower limit on the step size taken. In the original algorithm, if an improved function value is not found at the end of the suggested step, the step size is halved. In the program, an upper limit is set on the number of these bisections allowed before the current iteration of the minimiser is terminated. Because of the non-quadratic nature of the surfaces generated by the model, the limit on bisections is the one that usually operates, because the algorithm suggests steps that are usually too large. However 3 other stopping criteria are also needed. In the first place, an upper limit on the number of iterations is set. In the second place, the minimisation is terminated when the constraints do not allow any movement down the function gradient or when the suggested movement is less than 10^{-4} times the constraint width. Thirdly, since it has been found that the updating of the inverse matrix of second derivatives sometimes leads to wrongly directed steps, if the bisection limit is invoked after the normal operation of the algorithm, a steepest descent step down the gradient is attempted.

(3) The final modification made to the published algorithm is an adaptation in order to allow minimax values of the expected cost to be found. When the option in question is enabled, after each minimiser iteration, the second set of control variables is chosen, a correspondingly different J is used, and the sign of the function difference that is taken as an improvement is changed. The stopping criterion is now changed so that optimisation stops after failure to find improvements in two successive iterations or after the maximum number of iterations.

5.2.2. Optimisation as the Weak Link in Calculating an Optimal Control

It is important to observe that the weak link in the above procedure is the minimisation of the expected cost. If the simulation will run at all it will deliver the current cost k(i,u(i)). The Gauss-Seidel procedure will tend to converge simply because of the discounting factor. However the discovery of the true optimal control depends on finding the true minimum of the expected cost. Failure to find the minimum may be the result of inefficiency in the minimiser or the result of the existence of a number of local minima. If there are local minima the choice of the initial guesses at the control and the cost may cause the minimiser to become trapped in one of them.

While the effects of inefficiency in the minimiser are plain enough and can in general expect to be remedied by careful choice of method, stopping criteria and so on, the consequences of the existence of local minima require more thought. If there are local minima it is the result of system dynamics and the choice of the cost function, not of a fault in the optimal control procedure. It may be possible in such cases to locate the true optimum with some confidence by starting

the minimiser off from different guesses. However the utility of such a procedure is open to question because a system with multiple optima requires more precise control than one with a single optimum. There is also the possibility of catastrophic changes (in the sense of Thom, 1975) in which optima merge or vanish. The difficulties of making rational decisions under such circumstances may be so great that if crop/pathogen systems are found to have multiple optima, conventional methods of plant breeding are likely to continue to be used by default.

5.3. Documentation of the program OPTIPLANT

5.3.1. Overall Description

Abramsky and Van Dyne (1975) after surveying the literature on agricultural systems modelling suggested that there is a need for reports of computer programs to give

(1) the language used

(2) the type of computer on which and the institution at which the program was developed

(3) a verbal description of the code to amplify the comment cards

(4) the computer code itself.

The program OPTIPLANT which applies the method of section (5.2) to the model described in Chapter 4 is written in FORTRAN V. Apart from a Checkpoint routine useful for safeguarding long runs against system crashes and a call to an elapsed-time routine, both of which can be deleted if desired, it uses only standard FORTRAN V functions. It was developed on the UNIVAC U1110 computer at the Australian National University. As compiled it occupies 60 512-word blocks of core storage, 40 of which represent data storage of the results.

The minimum time for the program to iterate once over the 675 initial conditions is about 2200 seconds when the program has nearly converged. At earlier stages, when more minimisation is being done, time per grid point is up to four times longer. The results stored (for each grid point) are

(1) the current estimates of the components of the optimal control vector u. These are the amounts of each of the genotypes to be planted, the amounts of each pathogen race to be added, and the values of the economic control variables described in sections (4.6.1-2).

(2) the stochastic parameters of the system; the mutation rate and the proability of a pathogen irruption.

(3) the expected cost $V(\pi(i), i)$.

(4) the expected state output (storage level and pathogen spores of each race).

(5) the yield of the crop per unit area.

(6) the variance of the expected cost.

(7) $k(i, u(i))$ and $V(\pi(X_1),X_1)$.

These results are all stored in single precision, either in a (5x5x5x5x20) array during execution or in an unformatted mass-storage file between executions. Most operations on the results, however, are carried out in double precision.

The program itself consists of five subprograms. There is a main program which sets up the Gauss-Seidel iteration and calls the subroutine DAVMIN, which carries out the minimising procedure of equation (5.2.1). DAVMIN carries out its function evaluations by calling FUNCOS, which sets up each of the eight possible epidemic variations (see section 4.5) and calls EPIMUL. EPIMUL simulates the course of each epidemic and returns the state output to FUNCOS, which

then calls COST. COST calculates (k(i,u(i)), using the relevant model of section (4.6), and interpolates among the previously existing cost estimates by the method of equation (5.2.6) to estimate $V(\pi(X_1),X_1)$. FUNCOS takes the eight computed costs and forms the expected cost $E_i^{\pi}\ \rho(k(i,u(i)) + V(\pi(X_1),X_1)$ and its variance, which are returned to DAVMIN.

When each Gauss-Seidel iteration is completed the main program compares the set of costs calculated on the current iteration with the set calculated on the previous iteration, using chi-square as a measure of their difference. When the chi-square value drops below a pre-set bound, the iteration process is held to have converged sufficiently and the program terminates.

In the following sections a verbal amplification of the code of each subprogram is given. The program itself is reproduced in Appendix A. The amplifications are in the form of comments about groups of lines in the program.

5.3.2. The Main Program

Lines

20-35 VOPT is the array which holds the calculated results of the program (see section 5.3.1). RANGE contains the possible levels of each state variable's initial condition. The first subscript indexes the variable, the second the level. IEPARA and EPARAM contain integer and real parameters which will be referred to individually below. UMIRR contains a guess at the optimal values of the control variables. The first two columns of SLIM hold the values of the constraint limits on the control variables.

49-65 If IBYP = 0, the guess at the control and a guess at the cost function are put into VOPT and thence into the unformatted file 23 on mass storage. In later runs IBYP is set to 1.

67-75 The contents of F are the f_{ij} of equation (4.3.5). The first subscript indexes a plant line, the second a pathogen race. H is the area rendered immune by a cross-protecting spore.

82-110 REC stores a record of all the calculated costs for the present and last iterations. Initially it is filled with zeros in column 1 and the guess from line 56 in column 2. At the start of each Gauss-Seidel iteration, a chi-square measure of the difference between the columns is calculated. If this measure is greater than ERR, and IGSITR is less than or equal to the maximum allowed number of G-S iterations, another iteration is begun. Otherwise the internal subroutine UPDATE (see below) is called and the program terminates.

112-119 This is the beginning of the main iteration. Throughout the program I1 indexes the storage, I2 the aa races, I3 the aabb and I4 the bb. IGSTEP is usually set to 1 except for diagnostic purposes.

120-122 DAVJST is Davidon's matrix J (see section 5.2.1.1). Its initial version is stored in a separate file. Numerical errors prevent its carry-over from iteration to iteration, and so it is read in afresh at the start of the treatment of each point.

124 Subroutine DAVMIN that carries out the minimisation process described in section (5.2.1). IVPOS indexes the point on the grid of initial conditions being considered currently.

126-140 It will be observed that I4 runs only over the range 1 to I2. Because there is a complete symmetry between a state with amounts q_1 and q_2 of races aa and bb respectively in a crop consisting of C_1 %A and C_2 %B and one in which the corresponding values are q_2, q_1, C_2 and C_1, only 375 points of the full $625 = 5^4$ points in the grid need to be separately investigated. In lines 121-135 the current optimal control and output are used to give the control and output for their mirror image point directly.

141-145 After each minimisation the JOBSUP external function returns the time remaining to the scheduled end of the run, and updates the data file by calling UPDATE if more than five minutes have passed since the last call.

147-149 After each iteration the data file is updated and the contents of REC shifted in preparation for the next chi-square comparison.

153-181 When the program terminates, VOPT is investigated to see whether any of the optimal controls favour the use of less than 99% of the superline. If some do, the input/output relations for the first such grid point encountered are printed. If not, a corresponding message is printed out.

182-199 Subroutine UPDATE is an internal subroutine which updates the data file and sets IUPDATE = JOBSUP (1) to record the time at which this was done.

5.3.3. Subroutine DAVMIN

36 LAMBDA is Davidon's (1975) λ.

38 DIREC is used to measure the improvement in the function, represented in this case by YIELD. If the flag is set for minimisation (I=0), YIELD is unchanged. If the flag is set for maximisation, the sign of YIELD is changed.

40 If IENAB = 0 the subroutine minimises. If IENAB = 1 it maximises.

42-43 ISAD recalls whether the subroutine is in its minimising or maximising phase. ISADPR counts the number of minimising or maximising iterations that have been unsuccessful. It is reset to 1 after a success and augmented by 1 for a failure. If it becomes greater than 2, the minimax procedure terminates.

45-46 IPRINT regulates the number of intermediate results printed out. ITLIM sets the total number of iterations allowed before a default stop (usually 10).

58-62 The current value of the control vector U is read out of VOPT. DELTA was included to allow different sensitivities to different independent variables in the gradient

evaluation (line 106) but this option is not used in the current version. DAVREF is a record of the initialised version of DAVJST.

65-75 The initial function evaluation. If the output of an epidemic is similar to the input then successive function evaluations can allow the cost at the grid point in question to converge after the fashion of equation (5.2.3). In the arguments of FUNCOS, the function evaluator, X is the state vector, YIELD is the expected yield, EXCOS is the expected cost, VARRUS is the variance of EXCOS, and COSTX and COST are its components, E(k(i,u)) and $E(V(\pi(X_1),X_1))$. FUNCOS is called repeatedly and the resulting EXCOS inserted in VOPT until the effect on $E(V(\pi(X_1),X_1))$, if any, is small.

76-96 When the local convergence described above is complete, VOPT is up-dated, the state output is stored in temporary variables where the output resulting from later improved controls will be stored and a heading reporting the grid point and its input/output relations is printed.

97-98 IRENAC is a flag recording whether the last new step down the gradient was successful in reducing the cost. If so, IRENAC = 1. EPS is the lower bound of the mean squared gradient, used as a stopping condition.

101-137 The gradient evaluation procedure. IBLOCK determines whether the gradient evaluation is being made as a preliminary to making a step (IBLOCK = 0) or as a preliminary to up-dating DAVJST (IBLOCK = 1). The gradient itself is measured by taking a step along the direction in control space specified by the columns of DAVJST. NUA - NUB index these

columns. NUC and NUD are used to index the rows of DAVJST, corresponding to the variables in U being changed. For each column a vector TEM representing a perturbation of U in the given direction is filled by reference to DELTA and DAVJST. TEM is then scaled by the ratio $\max_j$ TEM (J)/SLIM (J,2,) where SLIM (J,2) is the width of the constraint in the jth direction. The result is multiplied by 10^{-4} and the Euclidean norm of TEM, AMOD, is calculated. The perturbed U (UDASH) is used in FUNCOS to create a perturbed EXCOS = FG, which is used to calculate the gradient.

140-163 Step One in Davidon's algorithm. A check is made that the default number of iterations has not been exceeded and a heading is printed. The step direction is set as being opposite to the gradient. Here and in the rest of this section there will be no attempt made to explain the details of the Davidon algorithm. The variables used in his updates are usually represented here with FORTRAN names chosen for their mnemonic similarity to Davidon's names, and an explanation of their significance is given in Davidon's (1975) paper.

166-172 Step Two of Davidon's algorithm. A new step Δu (TEM) is formed from the step direction set in Step One and its magnitude is set by DAVJST. IDIV indexes the number of bisections of this trial step.

173-189 IPROJ is a flag which is set if the step u has been projected into a constraint (section (5.2.1.2)). IFLAG is set if there is an active constraint. IMOVE is set if the gradient suggests a component of Δu that is more than its constraint width x 10^{-20}. Once these flags have been initialised, the Δu is scaled as described in equations (5.2.7-9).

191-199 Δu is tested to see if it is large enough to set IMOVE.

200-209 The α_{jk} of equation (5.2.9) are formed (ALMIN) and classified.

210-242 IREC records active constraints. When an active constraint is discovered for $1 \leq I9 \leq 3$ this corresponds to an attempt to go outside the triangular region of possible crop compositions, and the gradient projection method is brought into play. In lines 214-216 the number of such active constraints is counted. If there are none, then instruction 54 takes over the search for active constraints on the other control variables. If there are two, IFLAG is set at instruction 58 and all changes in crop composition are set to zero. If there is one, IFLAG is set at instruction 56. If IPROJ is set, this means that a single constraint was detected before and the gradient was projected, but that this created a new active constraint. Such cases are equivalent to two active constraints and are sent to instruction 58. If IPROJ is not set, the first three rows of column 1 of IREC consist of two 1's and a zero and form the outward normal vector to the active constraint. The first 1 encountered is multiplied by -1, which gives a vector parallel to the constraint. The scalar product of TEM with this part of IREC is then taken and used to project the TEM along this parallel vector. IPROJ is then set and the new TEM is recycled to instruction 38, via 54.

243-249 In these lines any active constraints affecting later components of TEM are treated by setting the corresponding components of TEM to zero. If an active constraint has been detected anywhere, the modified TEM is recycled to 38

to check that it now complies with the constraints and to calculate new α_{jk}.

251-272 The values of ALMIN stored in columns 3 and 4 of SLIM are sorted and the largest one is used to form the fraction of TEM that will finally be used in UTRY, which is a trial control. At this point numerical error sometimes results in the controls escaping from their constraints and attaining physically impossible values. Accordingly, values of controls close to the constraints are set onto the constraints and values that have overstepped the constraints are brought back to them.

274-275 If either no move is possible for the control along the function gradient, or Davidon's main stopping criterion is satisfied, the minimisation terminates.

277-290 Otherwise, the trial control is evaluated, and if an improvement is found control passes to instruction 78.

291-300 If no improvement has been found, Δu is bisected as many as five times and the resultant trial controls evaluated for improvement.

301-319 If no improvement is found as a result of the bisections, there are initially two possibilities. If IENAB = 1, one minimax phase has ended, and unless the last-but-one was also a failure the next should begin. If the subroutine is operating only as a minimiser, then if IRENAC = 0 (implying that there has not been a success since the last failure), the subroutine should terminate. If IRENAC = 1, DAVJST is reinitialised, and a complete new iteration is started with a gradient evaluation at the last successful improvements.

320-345 Instruction 78 is executed if an improvement has been found. ISADPR records this for reference in the next minimax phase and IDREC records which type of phase has been successful. Step Three of Davidon's algorithm is then begun. The previous best costs and controls are replaced by the improved versions, as are the temporary records of state output.

347-432 These lines are a direct implementation of steps Four through Seven of Davidon's algorithm and are best understood by consulting Davidon (1975).

433-441. In these lines the start of the next minimax phase is prepared. The flag ITEM is set if the minimax procedure is operating and either the past iteration is the first or a success has just been achieved. In either case there is a need for a new gradient evaluation and control passes to instruction 5 with altered values of NUA and NUB. Failing this, control passes to 20 for a step to be taken using the already existing gradients.

445-466 In these lines VOPT is updated and the final result of the minimisation is printed out.

5.3.4. Subroutine FUNCOS

31-32 If MUTMAT (I,J) = 1, then pathogen race 1 can mutate to pathogen race J. The convention used is that race aa is indexed by 1, aabb by 2 and bb by 3. Similarly VIR (I,J) = 1 means race I is virulent on line J.

35-36 PEXT is the probability of irruption; REALMU is the mutation rate.

45-50 The control vector U is split into the controls referring only to the epidemic - crop line proportions and amounts of added disease - which are put in UBAS, and economic controls which are put in ECU. ECU (4) holds the total amount of added disease so that this can be costed.

54 SPEXT is the number of pathogen propagules m^{-2} added to the system by the irruption of a single race.

56-75 An iteration is set up which covers the eight possible types of irruption and puts the number of spores to be introduced, besides those specified in the state variables, into SPAD. For each filling of SPAD, EPIMUL is called and returns a state output in X. X(1,2) contains the yield and this is immediately used in calculating expected yield, since subroutine COST will replace it with storage after levying consumption.

5.3.5. Subroutine EPIMUL

28-42 N is the number of days in the season. IMOD is 2 if the plant growth model is to be activated and 1 otherwise. ISTART is the date of first infection. LATEP is the latent period of infection. INFECP is the number of days for which a lesion produces spores. GAIN regulates yield compensation, which is absent if GAIN = 0. ICPTIM is the number of days for which a cross-protected area stays immune. VIGORS contains the growth constant of equation (4.3.5) for each of the races. YPAM1 and YPAM2 are used in calculating yield loss and IYCRIT is the date of the "critical period" at which yield loss is determined. FACT is C_2 of equation 6.4.1.

47-57 ALOSS will form a record of diseased area on each plant line, S will record the spores produced each day and SPOSUB will record the total area producing each kind of spore on each day. Y records plant area according to three indices. The first subscript indexes the type of area according to whether it is unoccupied or occupied by one of the three pathogen races. The second subscript indexes plant lines. The third indexes the day on which the area was occupied.

58-65 For the model without plant growth the total leaf area SPLAT is set to 1 and the unoccupied area of each race for all days is initially set at the fraction of 1 that each line occupies in the mixture. DELYT, which is a variable used in yield estimation for the growing plant, is here set at a constant 1.

67-71 SEED is the initial size of the plants. GROW is the growth rate of equation (4.3.13). FBACK is f of equation (4.3.14), regulating the efficiency with which an advantage is used by each line.

73-89 I1 indexes the day in the epidemic. If I1 = ISTART the full number of spores provided for by the state, the control and any irruption is introduced into S. If the total number of spores present is above 10^{-10} a flag (IDEP) is set to indicate the need to pass through the spore deposition sub-model.

91-105 At this point the denominator of the expression giving the probability of a spore landing on a plant is calculated (see equation (4.3.6)). For the non-growing (van der Plank) plant this is a constant. For the growing plant it is necessary to calculate the current value of SPLAT.

107-133 This is the spore deposition submodel. SPORAV is the average area taken up on a plant of line j by a random spore in the current population. OCCUP is the exponent in equation (4.3.10). Once the new spores have been deposited, the date of initiation of any cross-protection now just defunct is calculated, and the area that was cross-protected on that day is added to the unoccupied area.

135-156 The production of new spores is the most deeply nested loop in the whole program. Consequently SPOSUB(I), the area of tissue now producing spores of race I is calculated, at the risk of numerical inaccuracy, by adding the areas that have just become infectious and subtracting those which have just ceased to be infectious, rather than by direct addition of all infectious areas. At the length of infectious period ordinarily used in this study (20 days) the saving of computing time is roughly 90%. Once SPOSUB(I) is formed, S(I) is calculated by multiplication by the corresponding VIGORS(I), and the expected number of mutations is added to each race.

160-179 These lines implement the growth model of section (4.3.3).

182-200 These lines compute the yield loss. While Il is less than IYCRIT, the critical date at which yield loss is determined, ALOSS(J) is augmented by the area of newly infected tissue on line J. At IYCRIT this area is converted into a Cobb scale reading. Since the modified Cobb scale is linear, and sets its 100% upper limit at a 37% total cover (Melchers and Parker, 1922), the conversion factor is 0.37. The result of the conversion is then used as input to the FORTRAN version of equation (4.3.16).

205-213 In this loop the output of pathogens at harvest is converted to an input at the next cycle's of first infection according to the rule of equation (4.3.17).

5.3.6. Subroutine COST

24 NUMCRT is 1 if the subsistence farmer's criterion is being used, and is 2 for the agribusiness criterion.

28-49 These lines implement the subsistence farmer's criterion as set out in section (4.5.1). COSTX is k(u(i),i).

55-70 These lines implement the agribusiness submodel of section (4.5.2). In order to cast the problem in the form of minimisation of a cost, COSTX is set as -RETN.

75-88 The first task in interpolating to find $V(\pi(X_1),X_1)$ is to discover the 4-dimensional box of nearest neighbours in state space which encloses X_1. The bounds of this box are put in NNB as they are discovered by passing through the relevant row of RANGE until the first entry greater than the corresponding component of X_1 is found.

91-107 These lines implement the piecewise linear interpolation described in section (5.2). XINT is equivalent το β.

6. Experiments, Results and Analysis

6.1. Introduction

The purpose of this chapter is to report on a number of computer experiments that were done with the aims of

(1) determining for each optimality criterion, as a starting point for comparison, baseline optimal controls, for model parameter values that do not represent the presence of any of the special factors discussed in Chapter 2

(2) determining the effects of including a degree of general resistance in the model

(3) determining the effect of including a cross-protection effect in the model

(4) determining the effect of including stabilising selection in the model

(5) determining the effect of increasing the realism of the model by including plant growth effects.

In section (6.2.) the experimental design of the study will be described and the reasons for its choice discussed. Section (6.3) is a discussion of the methods of analysing the results of the kind of computer experiments made in this study, because the methods of analysis partly determined the kind of computations to be made. In section (6.4) the results of the experiments using the agribusiness criterion will be set out, and in section (6.5) the results for the subsistence criterion will be given. Finally in section (6.6) a summary of the conclusions of the chapter will be made.

6.2. Experimental Design

In Chapter 1 the task of choosing how to use disease resistance genes was analysed as that of answering the questions: which genes in which plants, when, and where. In Chapter 2 a number of biological factors were discussed which have been said to bear on the answers to these questions. In Chapter 4 a model within which the presence of these factors can be expressed was described, while in Chapter 5 the means for discovering an optimal control for the model were set out. The questions that must now be faced are: what are the optimal controls of the model for the different criteria, do they involve the use of multilines, and how do the biological factors affect these controls. This chapter describes some computer experiments intended to provide the data on which answers to these questions can be based. The design of these experiments has been affected in a number of ways by the limitations of this study, both biological and computational.

6.2.1. Biological Limitations

In this study the choice of possible genotypes has been reduced to three by the considerations of section (4.8). The choice of when they can be used has also been somewhat restricted because only one of the criteria, the agribusiness criterion, has the option of not planting anything at all. Even the choice of where various genotypes can be planted has been restricted by practical considerations (section 3.5.2.2) to the investigation of the possibility of using multiline crops rather than other spatial patterns (section (1.2.4.3)). These limitations are consistent with the main task of the study, which is to examine the rationale for the use of multilines. Thus the experimental task can be narrowed somewhat to the task of searching for any use of multilines and to explaining the effect of changes in the biological parameters on the underlying reasons for such use.

6.2.2. Computational Limitations

If the parameter values of the model were well established and computing resources were unlimited the search might be begun by fixing the values of those parameters unconnected with the allegedly significant biological factors and then exploring the region of parameter space spanned by all possible values of the remaining ("significant") parameters. Once this had been done, empirical laws might be discovered relating the optimal crop composition to the state variable values and the "significant" parameters. Explanation of the result would then consist of accounting for the form of these laws. However this study is restricted by the possible amount of computation to a very sparse sampling of the significant parameter sub-space, a sampling, moreover which cannot be related in any precise way to the region which would be of most interest if all the "non-significant" parameter values were known. Thus a series of "soundings" must be taken in parameter space. If one of these returns a multiline optimal control, the analysis then consists of, first, attempting to check that the result is not an artefact of computation, and then of attempting to explain why the multiline was chosen rather than a pure line.

In order to maximise the chance that this sparse sampling will pick up any cases in which multilines are optimal (or, conversely, to exclude as conclusively as possible the chance that they are ever optimal), the relevant biological factors must be strongly expressed in the model. For example in modelling stabilising selection, the simpler races must be given a large advantage. Such strong effects make it difficult to compare the results of different experiments directly, because it will not at first be clear in what way the cost function V has been changed. The first and most basic comparison that can be made is that between any optimal multiline controls that are

discovered and the corresponding suboptimal controls produced when the algorithm is contrained to use the pure superline. It is in the analysis of this comparison that the investigation begins, and it is this analysis that is the main topic of the next section. Once it has been carried out an extension to comparing the significance of the main biological factors can perhaps be made, through a comparison of the results of the analyses of different experiments.

6.2.3. The Experiments

In Figure 6.1 an experimental design for each criterion of optimality, based on these general principles, is laid out. The part of the diagram above and to the left of the dashed line represents the experiments that perform the sparse sampling just described. The part below refers to a minimax test for robustness (section 4.7.2) to be applied to any result discovered that involves multilines. The test will be described in more detail in section (6.3.3), but follows the same structural principle of pairing a constrained pureline result with a result where multilines are possible in theory.

The pattern followed by the experimental design is in fact a type of sensitivity analysis, since it was impossible to follow the effects of joint variations of biological factors. The uppermost box in the flow chart represents a run in which P_{ext} was set to zero and only the pathogen-free initial conditions were considered. This was done in order to discover the rational economic response to various levels of storage for each 'economy' when there was no risk of losses to disease. The columns of boxes below represent, from left to right,

(1) a baseline epidemic system where parameter values were chosen so that the range of possible pathogen initial conditions spread fairly evenly across the range from

zero to 90% yield loss, given a constant-size plant.

(2) an epidemic system where the same constant-size plant was given a degree of general resistance against all pathogen races.

(3) an epidemic system in which cross-protection is introduced to the constant-size plant. This brings the possibility that spores of the simpler races can be used as a kind of fungistat, and so there is an additional experiment in which spores may be added but only the pure superline grown, in an attempt to separate the fungistat effect from any multiline effect.

(4) an epidemic system where there is stabilizing selection, again for the constant-size plant.

(5) an epidemic system where the baseline pathogen is applied to a plant like the baseline plant except that it grows.

The manner in which each effect is implemented and the values of the parameters used will be discussed in the subsections of section (6.4) as the results of each experiment are discussed.

In summary, this section has outlined the computational and biological limitations on an experimental program based on the concepts and methods of the previous chapters, and has laid out the pattern of experiments to be done within these limitations. The approach is to sample the parameter subspace corresponding to the significant biological factors at wide intervals to maximise the chance of detecting any effect that leads to the use of multilines. If multilines are found to be optimal in some cases, analysis begins with a comparison of the optimal pure-superline policy for the same epidemic system. The methods of analysis are the topic of the next section.

6.3. Methods of Analysis

Each computer experiment takes the form of

(1) setting initial guesses at the grid of policy costs and control actions (these may be the results of a previous experiment)

(2) imposing any special constraints on the control variables (these are usually the constraints to use only the pure superline and not to add spores of the simpler races. Both are relaxed in the search for multilines).

(3) setting parameter values to those required for the experiment

(4) starting the Gauss-Seidel iteration and allowing it to run until the activity of the optimiser ceases, or the convergence error becomes small.

In each case where the controls are permitted to search for optimal controls involving multilines there are two possible types of result: either the procedure of the experiment results in the choice of a multiline for at least one initial condition or it does not. If it does, since accepting the presence of a multiline is equivalent to a rejection of the null hypothesis of the study, the following questions must be asked:

(1) does the result represent a significant usage of multilines? (The definition of "significant" in this context will be considered below).

(2) is the result robust, in the minimax sense of section (4.7.2)?

(3) why are the multilines being used?

(4) is the result locally or globally optimal?

The ways in which these questions can be answered will be considered in this order.

6.3.1. Judging the Significance of a Multiline Result

Whether or not a multiline result is significant depends on

(1) the definition of "significant"

(2) considerations related to the shape of the cost surface and the means available for discovering the optimal control

(3) the errors resulting from the process of computation.

In the following sections we will consider first how the shape of the cost surface may affect judgements of significance and, second, the sources of computational error that may affect the numerical exploration of the surface. Using the results of these sections we will then attempt a definition of significance suitable for the study.

6.3.1.1. Interactions between Controls

An important problem in attempting to explain or judge the significance of any use of multilines prescribed by the optimal control algorithm lies in the fact that the other control variables will interact with the crop composition in determining the optimal control. (It should be noted that the "area planted" variable in the agribusiness criterion has both direct economic significance and significance as representing a mode of gene use, though not the one of most interest here). For example, the use of a multiline may decrease the yield per unit area of the crop. In this respect it shares some of the effects, for the agribusiness farmer, of planting a smaller area - it increases the price at a constant fraction sold. Similarly the subsistence

farmers will see the use of a multiline as decreasing their bodyweight at a constant work rate. Such interactions are inevitable in any economic context and are certainly unavoidable here.

More generally, the difficulties introduced by such interactions result from the fact that the cost gradient even at the optimal values of the economic controls (proportion sold, area planted, work done) for a pure line crop may not be orthogonal to the economic control axes. Thus even if the pure-line optimum has been found, it need not be expected that a multiline optimum will differ from it only in the crop composition. If the pure superline optimum has not been found, relaxation of the constraint on the algorithm may cause it to use multilines to repair its failure, even if they are not in fact optimal.

If there are some analytical aids to discovering the optimum or if the problem is of low dimensionality then it may be possible to discover that this non-optimal use of multilines has occurred. A partial check will be carried out in later sections in the course of describing the results by plotting the surface of cost as a function of the economic variables, so that the policy chosen by the algorithm can be compared with the location of the optimum on the surface. This check has two limitations. First, because of the recursive nature of the equation for the optimal control, the activity of the optimiser changes the shape of the surface and so it is not possible to say from the surface what the true optimum is, but only that the current choice of the control is (or is not) near the current estimate of the optimum. Thus the check is most useful in seeking out cases where the optimiser has stuck at a sub-optimal decision. The second limitation is that the cost surface cannot be drawn for all control variables. Because of the number of dimensions involved, it can only be observed "statically" at one crop composition.

In cases where analytical or pictorial aids do not allow the interactions of the control variables to be disentangled from error it can be difficult to discover why multilines are used. This is an aspect of deciding whether their use is "significant", for if multilines have simply been used (suboptimally) to aid the deficiencies of the optimiser they should be explained away rather than explained. It is also difficult to check on the alleged optimality of the multiline. It is not legitimate to replace the allegedly optimal multiline by a pure line and evaluate the expected cost: the whole policy and in particular the economic control variables for the multiline point in state space will have altered in some degree to adjust to the presence of the multiline. Only if a pureline point is found in the multiline policy with a lower cost than the multiline point for the same initial storage and superrace conditions can the multiline be rejected.

Therefore in a study as complex as the present one, at some stage the decision of the optimiser that a multiline is optimal has to be taken as final. Some checks can be made but in the end the complexity of the system makes it impossible to discover whether the algorithm is correct without comparing its results with the results of another algorithm. In some cases plausible explanations of why multilines have been used by the algorithm can be made which support the likelihood that their use is optimal (see section 6.3.2), but if there are known to be sources of error in the algorithm, then these explanations are persuasive and qualitative rather than conclusive and exact. We now turn to considering the possible sources of error.

6.3.1.2. Computational Errors

There are four types of computational error to be expected in a study of this type.

(1) simple numerical errors

(2) errors due to the level of resolution of the optimiser

(3) convergence or residual error

(4) interpolation errors

6.3.1.2.1. Numerical Errors

Since there is no replication in most of the results it is usually impossible to estimate directly how much of the differences between the costs and between the recommended controls of neighbouring initial conditions on the state space grid is due to numerical error. In some experiments, however, where the control is constrained to the use of the pure superline and there is no cross-protection, there is a form of replication available through the comparison of initial conditions differing only in the amounts of the simpler races present. This form of replication was not available in all such cases because advantage was taken of the fact that it was 15 times faster to calculate pure superline optimal controls on a restricted grid of 5x5 initial conditions varying only in storage level or superrace abundance, rather than the full grid of 375 points. Where the replication was available, it was found that separate evaluations of the policy cost at the replicate grid points were the same to at least seven decimal digits. Inaccuracies were detected in only two aspects of the program, neither of them affecting the results. The first of these was that the variance of the expected cost (in cases where it would be expected to be small) was sometimes negative and of the order of 10^{-5}. This represented the result of repeated subtractions of numbers of the order of 10^{12} or 10^{14}, so that errors in the expected cost were actually only in the 16th decimal digit. The second error that was observed was that pure line crops sometimes produced small amounts of the simpler races (of the order of 10^{-7} spores), presumably because of the repeated

additions and subtractions in lines 140-147 of EPIMUL.

Thus, since the optimiser was restricted to detecting cost improvements in the seventh decimal digit, numerical errors in the function evaluation were not significant in this study. The other possible source of numerical error was accumulated error in the updates of the matrix J. However, since the optimising process was periodically restarted by a steepest-descent step, any errors in the updating process could have no effect.

6.3.1.2.2. The Resolution of the Optimiser

Because of the need to truncate the optimising process in order to save computer time, the optimising subroutine was not allowed to resolve function improvements less than 10^{-6} of the current function value, nor to resolve control variable changes whose largest component was less than about 0.1% of the constraint width in that direction. Consequently, controls that were initially similar might in theory have drifted apart to some extent during the optimisation process. It was also possible, if the function was flat, for marginally suboptimal, but noticeably aberrant, controls to be frozen into the final result, both because of the level of error permitted in the cost increment and because the optimiser tends to concentrate on the control variables that have the most effect on the cost. Finally, as a result of these restrictions and of the difficulties of optimising on the complex cost surfaces that were found, the optimiser sometimes tends to be "sticky", that is, an apparently settled control can sometimes be made to achieve considerable improvements by relaxing the restrictions on its resolution. In such cases the improvement comes about because the gradient changes sharply, and the only successful steps that can be

taken at any one time are very small, and are ignored at the lower level of resolution.

Because of the lack of replication, in most cases the significance of these effects in the study was hard to judge. However in some test cases apparently suboptimal points (as judged from the cost surface) could be brought close to the optimum, though only by allowing the optimiser to search for any improvement at all, thus showing that the cost surface as drawn was correct and that the resolution of the optimiser was a limitation on the accuracy of the result. Some policies produced during the course of the experiments apparently suffer from the low level of resolution, and these will be mentioned in the next sections, but to increase the level of resolution of the optimiser to its test level during the experiments would have increased the amount of computing many-fold. To put this shortcoming into perspective it should be remembered that the limitations on the resolution of the optimiser in any practical non-linear optimisation will always set a upper bound to the accuracy of the results.

In cases where there was replication, as described above, the null hypothesis that the optimiser was responding only to random, normally distributed errors in its choice of controls can be tested by using an analysis of variance and examples of this test will be given in section (6.4.1). However this test cannot be applied to the cases where there is no replication. In these cases a regression approach might be used, but it would be impossible to distinguish between errors due to the optimiser and errors due to lack of fit of the results with the regression model, and there do not seem to be any simple regression models appropriate to the results described in the next section.

6.3.1.2.3. Convergence Error

The third possible type of error is convergence error. Typically during convergence, the optimal control is selected while the chi-square measure of convergence is still large (i.e., significant at the 5% level). Once the control has been selected, the chi-square value decreases roughly exponentially with the number of Gauss-Seidel iterations. This decrease represents an average change in the value of the cost estimated from iteration to iteration that is governed by a nonlinear second order difference equation. There is thus some difficulty in estimating the 'true' value of the cost, though in theory by continuing the convergence for long enough, the magnitude of the change in the cost could be reduced as far as is desired.

The problem that arises as a result of convergence error is that of comparing the costs of two policies, the "parent" constrained to the use of pure lines, and the derived policy (which has been allowed to diverge from the parent) free to use multilines. The result of the differing amount of convergence of the two policies is that their true difference in cost is partly obscured by a kind of background noise of cost changes distributed unevenly among the grid points. Thus unless more computing is done than is necessary only to establish the form of the optimal control and to test for the existence of optimal multiline points it is not possible to test the hypothesis that multilines may cause a lowering of average cost over all initial conditions. However the changes in cost of individual points in the derived policy due to control changes may stand out against this background noise, and so one possible definition of a "significant" use of multilines is that the mean change of cost from the parent to the derived policy in the case of the multiline points is more negative than that for points where purelines are optimal in both policies. This can be tested quite simply with a t-test.

6.3.1.2.4. Interpolation Error

As was pointed out in section (5.2), the fact that linear interpolation is used to find the cost of the output of each control action will inevitably result in errors. If linear interpolation is then used again to calculate the appropriate control action for a state input intermediate between the grid points for which the policy was calculated, these errors will be compounded. In a more elaborate study these errors could presumably be reduced by an iterative process in which non-linear interpolating functions were developed that better fitted the apparent cost function as it was produced or by using a finer grid. Since these were not possible here, there is an inherent inaccuracy in the results of this study, whose magnitude is not easily measurable, though it can sometimes be seen to be small. (See for example section (6.4.1.2)). At present the existence of this type of error must simply be accepted. However there is no apparent reason why it should produce any systematic bias for or against the optimality of multilines, and so, given the general quantitative uncertainty surrounding the parameters important in the choice of modes of gene use, this type of error does not appear to be of great theoretical significance.

6.3.1.3. A Definition of Significance

The results of an optimal control of a complex system fall uneasily between the realms of a deterministic and a statistical analysis. On the one hand the present results are deterministic in the sense that if there is no cost gradient suggesting the optimality of multilines, the optimiser will never use them; all deviations from the parent policy are improvements. On the other hand, the sources of error described above tend to reduce the chance that the true

optimum is found either by the parent or the derived policy; it is not clear that all the changes in the derived policy result from the relaxation of the constraints on the parent policy.

There are thus alternative possible definitions of significance in this context. One, from the statistical point of view, might be that significant use of multilines occurs when there is a cost reduction that is detectable by an appropriate test against the background noise of convergence. Another, deterministic, point of view would be that any use at all of multilines is significant as showing, at the very least, that there can exist a cost gradient which suggests the use of multilines. This is a weaker test of significance because statistical significance implies deterministic significance, but not vice versa.

Both these types of judgement need to be considered from the point of view of the policy as a whole, for a stronger test of significance is whether multilines are likely to be used in practice. A point in state space for which multilines are prescribed but which is not visited by important control trajectories is not of practical significance. The question of what trajectories are important is to some extent open. Probably of most interest are those that start from the no-disease conditions, since the most obvious question to ask about multilines is whether there is any point in using them when the crop is clean. Another important set of trajectories starts from the region in state space where the simpler races are abundant, storage is low and the superrace is rare. This region corresponds to the recent failure of a resistance gene and the associated control decisions are those that must be made when a new gene is introduced. If statistically significant multilines are rarely used on these trajectories, then the other costs of maintaining the multiline option, mentioned in

section (4.6.1) but omitted from the criteria for the sake of simplicity, become important and multilines are not "significant" in the third and strongest sense. In the description of the results these questions will be investigated by using the derived optimal policy to respond to the simulated pathogen over a number of crop cycles. This will incidentally illuminate one of the oldest questions in the topic of optimal gene use (section 1.2.4.1), whether rotation of varieties can combat pathogen evolution successfully.

In summary, we conclude with three definitions of significance that can be applied to test any occurrences of multilines in the results. They are not mutually conflicting, and they are useful for characterising the results. Application of these tests does not resolve the problem of deciding whether the true optimum has been found and is a multiline optimum, but that decision can only be made on the basis of repeated studies.

6.3.2. The Reason for Using a Multiline

If a multiline is declared to be optimal for some initial condition by the algorithm, the declaration can be explained at three levels within a study of this type

(1) The economic level. Either $k(i,u(i))$ or $V(\pi(X_1),X_1)$ or both have been decreased by the use of multilines

(2) The state variable level. Some of the state variable outputs of the given initial condition must have been altered in an (economically) advantageous way.

(3) The epidemiological level. Interactions within the multiline epidemic must have caused the changes in state variables.

In accounting for the epidemiological interactions, either within or between experiments, we deal directly with the assumptions of the model, and thus in this study this is the most fundamental possible level of explanation.

6.3.2.1. Economic Interactions

An explanation at the economic level is the first step in explaining the use of multilines, since if there are no direct benefits, the optimiser will not use them. In this analysis the cost decreases at each multiline point are classified according to whether they result from decreases in $k(i,u(i))$, $V(\pi(X_1),X_1)$, or both. It is obvious that one change or the other must take place, and in this classification we are seeking evidence of a pattern of cost change than can summarise the purely economic reasons for using multilines. However a decrease of both cost components at once is a priori unexpected (section (3.5.3)) and in such cases it is necessary to examine the economic criterion itself to see how such a simultaneous decrease can come about.

6.3.2.2. State Variable Changes

The second step in analysis is to identify the changes in the state variables that have affected $k(i,u(i))$ or $V(\pi(X_1),X_1)$. Changes in $V(\pi(X_1),X_1)$ can come about, to a first approximation, only through changes in the storage output to the next cycle or the superrace input to the next cycle. (The latter is the superrace output from the current cycle transformed by passage through the offseason). If multiline points are equilibrium points or if the average cost over all points in state space has been changed by the use of multilines then the costs associated with a given X_1 will also have changed. In this

context the state variables representing abundances of the simpler races are the independent variables, because, by Theorem 1 of Chapter 3, they are the means by which the multilines achieve their effect. Changes in storage and superrace abundance are the dependent variables which, given a favourable interaction with the economic control variables, in turn influence $V(\pi(X_1),X_1)$ for good.

From a plot of the expected cost of the optimal policy against the initial values of the state variables it is possible to see the expected effect of the given change in the state variables. Because of the nonlinear nature of the cost function, the change in cost cannot be directly divided into components due to each state variable change. Instead a simple four-way classification can be made of the state variable changes according to whether they were movements up or down the cost gradient with respect to storage or superrace abundance.

The possible effects on $V(\pi(X_1),X_1)$ of an increase in storage, are, in the case of the agribusiness criterion, always favourable because there is no cost attached to storage, and so the possession of stored food represents only the possibility of selling food that does not have to be grown. In the subsistence criterion an increase in storage is not always desirable, and in this case a state variable affects $k(i,u(i))$ directly. Otherwise analysis at the state variable level bypasses consideration of the present cost. The significance of an increase in the abundance of the superrace cannot be told directly from the cost equations, but it is to be expected that (see section 3.5.3) that it will be adverse. Thus analysis at this level is mainly directed towards answering the question: which component of the change in the state variable output is responsible for the reduction of the expected cost of all future epidemics. In terms of an explanation at the state variable level, this analysis can be

expressed in such forms as "the use of multilines was indicated in this case because it permitted an increase in the storage output".

6.3.2.3. The Epidemiological Level

The basic epidemiological events that determine the divergence of the derived policy from the parent policy are the patterns of disease cover at the critical point when yield loss is fixed and at harvest, when the conditions for the off-season competition of pathogen races are set. There is only one basic interaction within the model by which the presence of the simpler races can affect these patterns; by infecting plant tissue they occupy it and make it unavailable for the superrace. This is the basic van der Plank property of the model. If only 1% of the total crop area is covered by the simpler races, then the superrace growth rate will be decreased by at least 1%, which will produce a cumulative effect on the superrace cover over several generations. At the same time, of course, the simpler races will be expected to affect the yield of the simpler lines.

In addition there are two possible "optional" interactions

(1) where there is cross-protection, the simpler races can occupy area (temporarily) without infecting it.

(2) where the plant grows, the simpler races can modify the total size of the plants in the mixture.

Finally, in this model the pathogen races compete during the off-season. This interaction, though it is determined by the pattern of the disease cover at harvest, tends to intensify any effect of the presence of the simpler races on the superrace during the crop epidemic itself.

Any particular multiline epidemic produced by this model can thus be described at the epidemiological level by its position in a 2n(n+1)-dimensional space, where n is the number of lines (and races) present. There are n+1 classifications of infected area on each line (including uninfected), n lines and two significant dates. The epidemiological change between the parent and the derived policies corresponds to a movement in this space. In this study 2n(n+1) is already 24 and so a simplified measure of the change is necessary. For this purpose we will use just two variables

(1) The total change in superrace input to the next cycle. This expresses the interference experienced by the superrace, and using the total rather than a proportional change allows a common comparison with the output from the growing plant where the superrace may have changed as a result of change in size of the plant.

(2) The change in yield of the crop. The yield itself seems to be the best way of summarising the joint effect of three different pathogen races on three different plants on the yield.

In a sense all possible multiline effects within the model are already "explained" because it is known what they must be based on. However it is not known in advance what the pattern of factors contributing to their use will be, and a description of this pattern is the principle contribution of the analysis of any multiline result. Later, the joint effect of crop composition and the variation of a "significant" biological parameter will be used as the basis for comparisons between the results of different experiments (section 6.6)).

6.3.3. Robustness of the Multiline Result

The minimax approach to the robustness of theoretical results can, as was shown in section (4.7), be applied to the testing of different types of robustness. In these experiments the method will be applied to the testing of robustness with respect to uncertainty about the two parameters that it will be most difficult to estimate in practice: the mutation rate μ, and the probability of an irruption P_{ext}. In the maximising phase of the minimax game option of the program, these will be used as control variables and allowed to range between, respectively, 1 and 10^{-20} and 1 and 0.

This procedure is intended as an illustration of what might be done in a real control situation where a conservative approach is being taken to the possible effects of parameters whose values are only known within an interval. It is also intended as an illustration of how the criticism might be met that multilines are perhaps only optimal for the initially chosen values of μ, and P_{ext}. However the procedure may also serve as an indication of the extent to which risk itself is important in the optimal control. If the maximising of the cost does not leave these two stochastic parameters at their extreme values it is evidence that not knowing whether (for example) the super-race is going to arrive contributes more to the cost than the damage due to an inevitable irruption.

6.3.4. Uniqueness of the Optimal Control

In section (5.2.1) it was pointed out that if there is more than one local optimum of the cost function for any grid point, then which one is chosen may depend on the initial guess at the optimal control, and the one which is chosen may not be the global optimum. In this section a more detailed discussion of the causes and resultant

problems of multiple optima will be given in order to show how the analysis of results is affected by their presence.

It is not necessary for the cost function k(i,u(i) to have multiple optima in order that $V(\pi(i),i)$ should have multiple optima. It is only necessary for k(i,u(i)) to be non-convex, because since $V(\pi(i),i)$ is the sum of a sequence of functions k_1, k_2, k_3, ...k_∞, the summing of non-convexities may produce multiple optima. Furthermore, it is not necessary for k(i,u(i)) to be non-convex for all i and u: occasional non-convexity is quite sufficient to induce multiple optima in $V(\pi(X_1),X_1)$. A reference which is instructive on the importance of convexity in optimal control is Halkin (1964), though Halkin was concerned with the use of the maximum principle in open loop control.

However, once multiple optima exist for even one grid point a complicated situation arises. Because the equation for the control is recursive (section 3.5.1), the choice of alternative optima affects not only the grid point i but the costs of other grid points whose state outputs are near i. Thus the choice of a local optimum will have an effect that propagates (usually with decreasing amplitude) through the whole policy. Consequently, in order to assess the effect of choosing one alternative over another it is necessary to confine the policy to optimise within each local optimum in turn. If there are multiple optima at several grid points, combinatorial problems obviously arise. Also, the choice of alternatives at one grid point may affect the existence of optima at other grid points.

Where such complex interactions occur as the result of the existence of multiple optima it becomes difficult to make comparisons between policies that have been allowed to converge from different starting points. Suppose that two unconstrained policies have converged

independently and one of them recommends multilines while the other does not. Evidently the two results represent alternative local optima, and it can be said that for one local optimum multilines are optimal and for the other they are not. The expected cost of the two policies can also be compared (see section 6.3.1.2.3.). However it is difficult to answer question (3) of the list in section (6.3) - why the multilines are being used - because the two results do not show what would be the result of using a pure line in the case where the pureline is allegedly suboptimal.

Thus any presence of multiple optima means that the analysis of multiline optimality (if it is found) takes on a local character corresponding to the local character of the result itself. Just as multiline optimality can only be asserted to be true for the particular optimum that has been found, so also the significance of the result, its robustness and its rationale all relate to the particular optimum. The implication of this restriction for experimental design in cases where multiple optima exist is that the appropriate technique is that used in this study: to test the null hypothesis of the study for each set of parameters by first producing a policy that is constrained to the use of pure lines and is locally optimal and then relaxing the constraint and giving the unconstrained optimiser an opportunity to produce a divergent result.

As will be seen in sections (6.4 & 5) the results of this study exhibit several types of multiple optima, including some that result from direct multiple optima in k(i,u(i)) and some that result from the effects of summing non-convex functions. These types of functional aberration are inconsistent, in the sense that they appear only for certain state values and parameter values. There are thus catastrophes in the sense of Thom (1975) present in the results, which appear to

result partly from the use of relatively complex and realistic (i.e., non-linear) functions in the criteria and partly from the complexity of pathogen dynamics. From the point of view of this study these catastrophes represent complications in analysis and restrictions of significance of the results. It seems quite likely that catastrophes will occur even in models using simpler criteria, and in section (7.4) the significance of this phenomenon from the point of view of practical applications will be discussed.

6.4. The Agribusiness Criterion

6.4.1. The Economic Parameters

The values of the parameters in the equations used to describe the agribusiness criterion (section 4.3) were chosen so that the natural condition of the farmer in the disease-free state would be a prosperous one. These values are set out in Table 6-1, and the following points about them are worth noting. Since the maximum yield is set at 400 units and $b = 1/Q_{opt} = 300$, it is possible for the farmer to grow more than the amount that will give him maximum return. (Because there are variable costs associated with the area planted (FERCOS) the optimum amount to produce is a little less than 300.) The fixed costs and the variable costs together are small compared to the attainable return, so that while the fixed cost is 365 units of return the maximum return to the farmers is 10,672 in any one year. The value of the fixed costs does not affect the optimal policy in any way; if the farmers were all paying back loans at a high (but constant) rate their actions would be just the same. The fixed cost here was thought of as representing chiefly the maintenance of farm families.

Table 6-1. Economic Parameters for the Agribusiness Criterion

FORTRAN Name	Source in Text	Value
BASRAT	p_0 (equation 4.6.1)	100 return units/unit crop
QOPT	1/b (equation 4.6.1)	300 crop units
FIXCOS	c_0 (equation 4.6.6)	365 return units
FERCOS	k_3 (equation 4.5.6)	1000 return units
SPOCOS	section (4.6.2)	0.1 return units/spore/m^2

Thus, with these parameter values, and the capability of storing 500 crop units, the agribusiness farmers in cooperation have a range of economic options. They can produce consistently, or they can hold the crop in storage in order to force the price up while planting a lower area. It is certain that they need not go bankrupt in the absence of disease, and their situation in fact resembles that of those farmers in the U.S.A. who until recent years were subsidised not to produce because their attainable production would have flooded the market. In the next section the actual behaviour of these prosperous farmers will be presented.

6.4.2. The Disease-Free State

The first step in reporting each result is to display the optimal policy itself, that is, the rules relating the optimal control variable values to the state variable input. As the number of state variables increases it becomes more difficult to display the results graphically and, in practice, interpolation in arrays of the same dimension as the number of state variables would be the way in which the control action would be determined. In the disease-free case, however, there is only one state variable operative, the storage, and thus the basic properties of the 'economy' can conveniently be displayed graphically.

Figures 6.4.1 and 6.4.2 display the control policy, which is qualitatively sensible. At storage levels below Q_{opt} (equation 4.6.3) the entire storage is sold since, with no disease there is no danger of not being able to grow the optimum amount the next year. At the higher levels of storage some of the total is retained in order to avoid depressing the market price. At low storage levels a substantial area is planted, but when the market can be fully supplied by selling from storage, nothing is planted.

The second step in reporting the result is to describe the consequences of the control action at each grid point of initial conditions in terms of the expected cost (or, in this case, return) and state variable output. The disease-free control is not a stochastic one and so the 'expected' state outputs are certain, but the procedure is the same. Figures 6.4.3 and 6.4.4 show the expected return for the infinite policy starting at each storage level and the storage output after the optimum amount has been sold. The effect of the optimum policy is to smooth out fluctuations in the return that might result from unwise sales. Also, at the 8% discount rate there is little present advantage in the possession of stored food, since the returns from the years after it is sold dominate the result. The state output pattern is that to be expected from the description above of the control laws.

Although the patterning of the economic control through time has no direct theoretical importance in the disease-free case, it serves as a useful basis for comparison with the pattern when the pathogen is present. Figures 6.4.5-8 repeat the information in the previous four figures, but displayed as a pattern in time instead of state space, starting from an initial condition in which the silos are full. Figure 6.4.7 shows the proportion sold in each subsequent year

and figure 6.4.8 shows the area planted. Figure 6.4.5 shows the return to be expected in each year and figure 6.4.6 shows the level of stored food. The policy quickly settles down to a state that is not only steady but stable, after the removal of the surplus, and because of this the present return estimated in the feedback process (115,480) is close to the total return actually accumulated during the 50-cycle period (113,715).

The question now arises of whether this optimal policy is the global optimum. To show that it is we display the surface of returns that the optimiser faces at each grid point in attempting to maximise the return in the final stages of convergence. Figure 6.4.9 shows the surface of $-k(i,u(i))$ (the return is the negative of the cost) for each storage value and figure 6.4.10 shows $-V(\pi(X_1),X_1)$, while figure 6.4.11 shows the total return $-V(\pi(i),i)$. The contours divide the interval between the lowest return and the highest return on each map into deciles, with "1" labelling the 10% level. The contour marked "*" encloses the returns > the 99th centile.

The first point to note about these contours is that the sets of contours, for $-k(i,u(i))$ and for $-V(\pi(i),i)$, are very similar, with the 99% contour widened by a flattening of the maximum of $-V(\pi(i),i)$. The reason is that since the optimal return does not depend much on the storage level (fig. 6.4.3), $-V(\pi(X_1),X_1)$ is almost constant for the various storage outputs produced by different combinations of area planted and proportion sold. What difference is made is the result of the fact that the return from subsequent years is maximised if a large amount of stored food is handed over from the first year, as can be seen in figure 6.4.10. The second point is that the form of the optimal control laws in figures 6.4.1-2 can be understood as a rotation and slight distortion of the contours of cost within the control

constraints, with the break in the slope of the curves describing the policy corresponding to the storage level at which the optimum point passes the bottom right hand corner of the return surface. Finally, it should be observed that the cost surface is non-convex, partly because of the dogleg in the contours that can be seen at intermediate levels of storage. This dogleg is the result of the limit of 500 units on the storage level. For a given proportion sold, as the area planted rises above the optimum, increased costs are incurred and the price falls. However when the total of what is grown and what was stored exceeds 500 units, the excess is lost, and although further increases in the area planted cause cost increases they do not affect the price. This dumping of the glut of the product decreases the effect of excessive production on the return and the flattening of the gradient of the cost along the vertical axis is the result.

In summary, the optimal policy for the agribusiness in the absence of disease is to plant the minimum area needed to ensure sale of the optimum amount of the crop and to dispose of surpluses without causing a glut. This policy leads to a stable steady state at zero storage, which is a convenient yardstick for comparison with cases in which the pathogen is present.

6.4.3. The Baseline Epidemic

6.4.3.1. The Biology of the Epidemic

The baseline epidemic was designed so that it would represent a serious disease problem. A list of the key parameter values is given in Table 6-2. The result of using these values is that on a fully susceptible crop,

(1) the introduction of the standard irruption of 50 spores per unit area reduces yield from 400 units to 360.

(2) the maximum spore input of 1000 spores per unit area reduces yield to 50 units. Figure 6.4.12 shows the dependence of yield on pathogen input.

(3) if 10^{-8} spores per unit area are introduced (corresponding to a mutation for increased virulence in a fairly light epidemic of a simpler race) 95% of the maximum possible pathogen level is reached in 9 crop cycles. Figure 6.4.13 shows the time course of this increase.

Thus the ordinary progress of a pathogen race in a fully susceptible crop, without interference from other races, is a rapid increase from negligible levels to a level which reduces yield by 88%.

Table 6-2. Baseline Epidemic Parameter Values

FORTRAN Name	Source in Text	Value
SPEXT	(section 4.5.2)	50 spores/unit area
F	f_{ij} (equation 4.3.5)	5×10^{-6} area units
VIGORS	s_{im} (equation 4.3.5)	10% spores/unit area/day
PEXT	P_{ext} (section 4.5.2)	0.1
REALMU	(section 4.5.1)	10^{-8}
ILATEP	T_l (equation 4.3.4)	7 days
INFECP	T_p (equation 4.3.4)	20 days

6.4.3.2. The Constrained Control Result

The optimal pure superline agribusiness baseline policy was converged without using the fact that the grid points differing only in abundance of the simpler races are equivalent. The result thus contains a degree of replication (section 6.3.1.2.1) which will be referred to below. Apart from variation between replicates, the policy

and its outputs should only depend on two state variables: the storage and the superrace abundance. We now proceed, as with the disease-free case, to describe first the control laws and then the outputs of the policy. In the following four figures, paired graphs are shown with the dependent variables in common within each pair and the two state variables as independent variables. In each case what is plotted is the mean values of the dependent variable +/- two standard deviations where the variance is calculated over all 75 independently converged grid points sharing the same value of the independent variable. The plotted range of variability is thus due partly to error and partly to the effect of the other state variable.

Figures 6.4.14-18 show that the control laws and their outputs are not as simple as their counterparts in the disease-free state. There are suggestions of downward trends in proportion sold and area planted with increasing storage though they are not as definite as in the disease-free policy. There are also a slight upward trend in proportion sold and a downward trend in area planted with superrace abundance (number per m^2). The gradualness and (in some cases) reversals of these trends make it hard to see how the policy achieves its ends, and there is in addition a wide range of variability in most cases.

The expected return for each state conveys a more immediately interpretable message: at high storage the return is independent of the pathogen and at low pathogen levels it is independent of the storage. The first surprise of these results is that the return is so high, despite the seriousness of the disease, and that it is so little affected by the amount of the pathogen present.

The storage input/output relationship shows rather similar behaviour to the disease-free system, except that on average, more is retained from a given level of storage, suggesting a provident approach to the next year's epidemic. As might be expected, the storage output is lower at the higher superrace levels. The second surprise of these results is the superrace output, which is consistently much lower than the maximum to which it would be expected from Fig. 6.4.13 to tend, and which decreases with the superrace input at high superrace levels.

There are thus three problems of interpretation posed by these results

(1) If there is error in the results, that is, if functionally identical grid points are assigned different controls, what are the sources of this error?

(2) Are the trends in the policy laws significant despite any such error?

(3) How do the policy rules operate to reduce the superrace level and maintain the level of return?

The following three sections deal with these problems.

6.4.3.2.1. Tests of Significance

The analyses of variance displayed in Table 6-3 contain part of the answers to the first two questions. It is evident in the first place that there is error in the policy, in the sense that the sum of squares for error is non-zero. It is also clear that the effects of the state variables on the policy are strong and significant (though a linear regression model would not be appropriate for their description), with the storage variable as the dominant factor.

Table 6-3. ANOVA for Effect of State Variables on Choice of Control

Control Variable: Proportion sold

	Sums of squares	d.f.	Mean Squared Error	F-ratios
Storage	11.038	4	2.7596	81.216**
Superrace	.778	4	.1946	5.727**
Interaction	2.468	16	.1543	4.540**
Error	11.893	350	.0340	
Total	26.178	374		

Control Variable: Area planted

Storage	5.860	4	1.465	30.913**
Superrace	4.315	4	1.079	22.764**
Interaction	2.033	16	.1270	2.681
Error	16.587	350	.0474	
Total	28.795	374		

(** significant at the 1% level)

6.4.3.2.2. Sources of Error

In order to illuminate the reasons for the occurrence of the error in the replicates, a series of cost surfaces like those of figures 6.4.9-11 are displayed in figures 6.4.19-23. Although they are produced by the same method as the previous figures, the interpretation of these cost surfaces is slightly different because of the stochastic element that is introduced with the pathogen. Thus the maps labelled "-k(i,u(i))" are not maps of E(-k(i,u(i)), but of the values that -k(i,u(i)) would take if the expected output of the pathogen actually resulted. Similarly, the maps labelled "$-V(\pi(X_1),X_1)$" are maps of $-V(\pi(E(X_1),E(X_1))$, and the surfaces drawn are not exactly the same as

the surfaces "seen" by the optimiser, though the differences will be small at the higher pathogen levels. This inaccuracy is tolerated here because these maps can be drawn by interpolation in the accumulated results of the policy computation, which include the expected state output. The maps are drawn on a 50x50 grid, and if the exact surface were to be drawn, 50x50x8 = 20,000 epidemics would have to be simulated for each map. (Each function evaluation requires the simulation of eight epidemics - section 4.5.2.)

In figure 6.4.19, representing the cost surface at the lowest level of the superrace, the contours of $-k(i,u(i))$ are like those of the disease-free case, because the expected pathogen output is not high. The contours of $-V(\pi(i),i)$ are, however, substantially different from the disease-free contours, especially at intermediate levels of storage, in that it is now never optimal to sell all the available crop. The reason for this can be seen from the contours of $-V(\pi(X_1),X_1)$, which exhibit two new features. These are, first, stronger sanctions against selling all the first year's storage and, second, a pronounced advantage, especially at lower storage levels, of planting only a small area if more than about half the crop is sold. There are thus two opposing tendencies in $-V(\pi(X_1),X_1)$: one requiring a large area to be planted, and the second requiring a small area to be planted.

The trouble which these opposing tendencies can cause even in the relatively benign, low-disease state is indicated in the bottom right hand corner of the map of $-V(\pi(i),i)$ for storage = 250. Here the 90% contour narrows and then widens into the corner, indicating that there is probably an unmarked local optimum in the corner. This is the first sign that there may be multiple optima in the agribusiness policy, and in fact one of the (replicated) grid points asserts that the optimal action for this state is "sell everything, plant nothing", rather than that shown by the starred contour.

Thus even for these simpler cases where the superrace in only beginning to increase, multiple optima are a possible source of error. In figures 6.4.20-23, which show the surfaces when the superrace is abundant, the opposing pressures on the system can be seen to have caused much sharper changes of regime, in which double and even triple optima are common. The fine detail of these more complicated contours for instance the presence of numerous cusps, are beyond explanation at present. The cusps do not seem to be the result of the grain-size of the grid on which the contours are interpolated. A possible reason for their occurrence is that they represent the overlapping of local optima each with a roughly circular area of attraction.

It is now clear that there are two sources of error responsible for the error terms in the analysis of variance. One is simple failure to find the nearest optimum, as for example in the control for storage = 125, superrace = 500, which has been set by the optimiser to a point on the 8th decile contour at zero area planted. The other is, as above, seduction by the wrong optimum. These two types of error cannot be distinguished from each other except by comparison of the printout of the policy with the contours of cost, and then not always. However it is possible to give an account of the factors tending to cause both types of error, and both types appear to occur in this policy.

Failure to find the nearest optimum can be partly explained because these surfaces have two characteristic features which cause difficulties for the optimiser: the cusps already noted and a tendency, obvious in the disease-free case, for the optimum to be approached by a curving ridge. On both these types of feature anything but a very short step along the gradient will tend to result in a cost increase, and so the optimiser tends to proceed in a slow and cautious zigzag. When the steps in the zigzag fall below the lower size limit allowed

by the optimiser a suboptimal result is accepted.

The multiple optima appear to arise, as far as can be judged from the surfaces, for two reasons

(1) the interaction of the non-convex shape of $-k(i,u(i))$ with the tendency for it to be reproduced in $-V(\pi(X_1),X_1)$ inverted from left to right because the former function is decreased by sale of storage and the latter by retention of storage.

(2) the opposing tendencies within $-V(\pi(X_1),X_1)$ with respect to the area planted, where a large area increases available storage but also the superrace abundance, coupled with the non-linear dynamics of the epidemic.

It is to be expected that a simpler cost function would reduce the tendency of the epidemic dynamics to produce multiple optima; however even if the cost function were a convex function of the control variables, the non-linearity of the yield and pathogen output functions would tend to introduce non-convexities which in turn might produce multiple optima when summed. The existence of multiple optima should therefore not be dismissed as an artefact of the particular economic criterion used here.

6.4.3.2.3. Policy Operation

Despite the problems introduced by the complexity of the cost surface, when the operation of the policy is followed through time, it can be seen to behave sensibly. Figures 6.4.24-29 show the return for each year, the control action taken and the outputs, including the yield.

The policy achieves its effects of maintaining return and suppressing the superrace via a rather complicated fluctuation of the area planted. The control pattern repeats itself almost exactly every twelve crop cycles, though there are signs of a slight drift away from exact periodicity. The figures show the evolution of the policy starting from a clean crop and full silos. Alternative initial conditions only affect the first few cycles, after which the pattern reasserts itself, shifted in phase. Storage is built up while the disease is low, and then the accumulation is sold when an increase in the disease makes it advisable to reduce the area planted. The accumulation is not great enough to prevent large fluctuations in the return, but in the real world these variations in area planted would correspond to rotation of crop land into a crop sufficiently different from the one being modelled to have no chance of susceptibility to the pathogen being modelled, and the fluctuations in the return would be somewhat reduced.

It is not clear, especially given the considerations of the last section, how many of the particular features of this rotation are optimal and how many are suboptimal or accidental. Some of the errors to be expected are, however, distinct from the errors introduced in the convergence of the policy. As with the plotting of the cost surface, in calculating the evolution of the policy the graphing program assumes that the state output from each crop cycle is the expected output. Thus the path that is followed in state space is not exactly the "expected" path since $E(X_1(X_1))$ is not the same as $X_2(E(X_1))$, even though X_1 and X_2 are calculated by running epidemics, not by interpolation, and are accurate. A more important source of error is that the program assumes that the optimal control for a given state can be found by linear interpolation among the controls for neighbouring grid

points even though the pathogen output changes non-linearly between neighbouring controls. This effect is more important here than in the disease-free case because of the pathogen and because of the rapid changes in state that result from the rotation. The presence of the multiple optima may also have the effect of deceiving the policy about the future behaviour of the system. The policy may assume at one grid point behaviour that differs from that which it assumes at another grid point. The net effect of all these actual and suspected errors is that while the expected return calculated in the course of deriving the policy for this particular initial condition was 11.0×10^4, the actual return calculated by summing the discounted $-k(i,u(i))$ was 7.55×10^4.

In summary, the sensible policy for the agribusiness confronted by the baseline epidemic is to rotate with another, non-susceptible crop, or failing this, to leave a fallow. The reason for this is that the superrace abundance is strongly, though temporarily, affected by the area planted. At the same time as the rotation is carried out, some use is made of storage to smooth out the effects of the rotation on the return, but major fluctuations can still be expected. The quantitative analysis of this conclusion is complicated by the existence of multiple optima in the cost function and by the limited validity of the linear interpolation used in implementing the policy. However, qualitatively, the concurrence of the policy with commonsense after an elaborate numerical excursion is quite impressive.

6.4.3.3. The Unconstrained Optimal Policy

For the baseline epidemic and the agribusiness criterion and the given level of resolution of the optimiser, the optimal unconstrained policy was the same as that for the constrained control and no multilines were introduced. Thus it can be concluded that, at

least locally, multilines were not optimal. Consequently the null hypothesis of the study was not rejected under these conditions.

6.4.3.4. The Constrained Baseline Game

The control laws for the constrained baseline game are shown in figures 6.4.30A-D. At first sight they are apparently very similar to those for the control with constant P_{ext} and μ. However an analysis of variance of the difference of the control values used in the control and in the game shows that significant changes have been made from the parent control in response to the changes in the stochastic parameters. Table 6-4 shows that the dependence of the proportion sold and the area planted on the storage and of the proportion sold on the superrace abundance have been affected. Inspection of the control laws shows that the changes take the form that less is sold in the game at low storage and disease levels, which suggests a more cautious approach in the game in the face of future disease and that the area planted falls somewhat at intermediate disease levels, which would have the effect of decreasing the superrace output.

Table 6-4. ANOVA for the Difference between Baseline Control and Game

Control Variable: Proportion sold

	Sums of Squares	d.f.	Mean Squared Error	F-ratios
Storage	1.232	4	.3081	51.613**
Superrace	.3616	4	.0904	15.145**
Interaction	4.989	16	.3118	52.237**
Error	8.672	350	.005969	
Total	8.672	374		

Table 6-4 (cont.)

Control Variable: Area planted

	Sums of Squares	d.f.	Mean Squared Error	F-ratios
Storage	.9556	4	.23891	28.154**
Superrace	.0000379	4	.00000950	.00112
Interaction	3.843	16	.24019	28.305**
Error	2.970	350	.00849	
Total	7.769	374		
Control Variable: P_{ext}				
Storage	20.320	4	5.080	83.860**
Superrace	8.682	4	2.171	35.831**
Interaction	8.837	16	.5523	9.118**
Error	21.202	350	.0606	
Total	59.042	374		
Control Variable: Exponent of Mutation Rate, $M(\mu = 10^{-M})$				
Storage	1866.7	4	466.7	45.070**
Superrace	347.32	4	86.83	8.386**
Interaction	769.42	16	48.089	4.644**
Error	3624.0	350	10.354	
Total	6607.4	374		

Figures 6.4.31 and 6.4.32 show the pathogen's side of the game. The F-ratios in Table 6-4 show that P_{ext} and μ have been significantly changed, as the figures suggest, with the storage level having the dominant role in determining the change. The most obvious effect is that P_{ext} has generally been increased, though not to unity in all cases, while for two storage levels it is unchanged. These storage levels are those at which the value of P_{ext} can have little effect - 250,

where only a small area is planted, and 500, where the farmer is self-sufficient. The greatest average change in P_{ext} from its initial value of .1 occurs at low storage and low pathogen abundance where an irruption would have the greatest effect on the return.

The effect of changes in μ on the cost, as revealed by gradients calculated during optimisation, was always much less than the effect of similar proportional changes in P_{ext}, and this is reflected in the pattern of changes in M, where $\mu = 10^{-M}$. The results here seem counter-intuitive at first: where M has been changed from its initial value of 8 it has been increased, thus decreasing the mutation rate, and this has happened at high pathogen abundance and low storage. The explanation is that in the pure superline a high mutation rate is a high rate of back mutation to avirulence, which decreases the effect of the super-race. Accordingly the mutation rate has been decreased in that case, the case where a substantial area is planted, where the loss of a few spores would be most noticed.

Thus the pattern of action of the farmer's adversary in the game is apparently to adjust P_{ext} and μ in the direction which will decrease the immediate return of the policy. The fact that the stochastic parameters have not been set uniformly to their extreme values may suggest (as was suggested above in section (6.3.3)) that there is some extra disadvantage for the farmer in the risk of extra disease as distinct from its certain arrival. In fact, since μ is most harmful when reduced in the pure line and may be most harmful when increased in a multiline, this suggestion may well be sometimes true. In the present case, however, the explanation is to be found in the workings of the optimiser. Inspection of the printout showed that there were in fact no intermediate values of P_{ext} and μ: they were always either at their extreme values or their starting values. The reason

that they were sometimes at their starting values was that where the 'farmer' aspect of the optimiser chose to plant no crop at all, the 'pathogen' aspect could discover no effect of altering P_{ext} and μ, and left them at their starting values. The spread of results is thus partly accounted for by the presence of the local optima in the cost function, which fool the farmer's adversary as well as the farmer. Figure 6.4.33 shows that the adjustments involved in the game have in fact made the cost surface more complex at the low disease levels at which they are most important.

If it were not for the interactions between the control variables, the most effective policy for the pathogen would have been found to be to maintain the extreme values of P_{ext} and μ, and the correct policy for a pessimistic farmer is thus to assume when growing a crop that the superrace will arrive that year, even if there are no reports of its occurrence. He would thus gain by being able to calculate his policy more easily (as a deterministic problem) but would lose on the average by creating a depressed market with higher-than-planned-for yields.

The failure of the optimiser, unaided by human insight, to set the stochastic parameters to their extreme values is a failure of the method and points out the need to guard against control variables which interact as P_{ext} and the area planted do. However the failure does not appear at first to be of much consequence because the values of P_{ext} and μ have been reset in every case in which they could matter. This appearance is false as can be seen in Figures 6.4.34-9. which show the evolution of the game through time, from the same starting point as the control. The playing of the game has been sufficient to break down the regular limit cycle of the control policy into an apparently chaotic oscillation, though the essential qualitative approach to disease

control via rotation remains. The unfolding of the adversary's moves is shown in figures 6.4.38-39, which show that the failure to find the true 'worst' values of P_{ext} and μ can have an important effect, for intermediate values of P_{ext} and μ predominate in the pathogen's choice - or at least the farmer's expectation of that choice (section 4.7). The reason is once again the fact that the chosen control depends on interpolation, and the interpolation causes the farmer to imagine that intermediate values of the stochastic parameters are likely - an optimistic form of pessimism.

In summary, the worst-possible values for P_{ext} and μ in this case turn out to be their extreme values: 1 and 10^{-20} respectively. Adopting these values has a definite and significant effect on the policy chosen, and on its expected cost and behaviour through time, and also introduces extra difficulties in maximising the expected return (though the complications in the cost surface may be partly due to the interpolation problem). The policy is changed in the direction of greater caution, as is desirable, given the purpose of the game.

6.4.3.5. The Unconstrained Baseline Game

The relaxation of the pureline constraints on the baseline game did not result in the introduction of any multilines. The failure to reject the null hypothesis in this case means that the pureline result for the control of the baseline epidemic can be regarded as a locally robust policy.

6.4.4. The Agribusiness System with a Plant with General Resistance

6.4.4.1. The Biology of the System

As was pointed out in section (4.8), general resistance can be conveniently represented within the epidemic model as a decrease of the growth rate of the pathogen. A decrease in the constant sim of equation (4.3.5) represents a decrease in van der Plank's r which cannot be reversed by mutation on the part of the pathogen. The converse of this representation, which has not been discussed before, is that it is impossible to conceive of a plant that does not have some degree of general resistance, since a zero level of general resistance would correspond to $r=\infty$. From the point of view of the modeller (and the farmer, in the short term) at least, it is of no consequence whether the upper limit on the pathogen's reproduction is set by the pathogen or by the plant.

Thus all the plants modelled in this study have a degree of "general resistance". However for the purpose of determining the effect of increasing the level of general resistance of a crop, the vigours of all three pathogen races was uniformly halved from their baseline values so that

$$s_{im} = 5x10^4 \text{ spores/unit area/day}$$

for all i and m. The result of the change on the natural dynamics of yield and pathogen increase are shown in figures 6.40-41. The yield for a given spore dose is now higher because the pathogen cannot increase so rapidly in the epidemic season, and the final equilibrium of the pathogen is at a lower level because it cannot saturate the wild hosts so fully. Moreover, the original mutation which leads to the increase through time takes longer to generate a serious disease problem than in the baseline epidemic. When the pathogen does become established it increases almost as fast as in the baseline epidemic

because at this stage the limiting factor on the increase between cycles is growth on the wild hosts whose general resistance has not been augmented.

6.4.4.2. The Constrained Control Policy

The control policy and its outputs are displayed as for the baseline control (Figures 6.4.42-46), but because the policy was converged without replication the analysis of variance procedure cannot be repeated. Comparison of the policy with the baseline policy shows a basically similar structure with some notable differences. The most important change in the policy is the much more definite minimum on area planted and maximum on proportion sold at storage = 250: however this would have been optimal (according to the cost surface) for the baseline epidemic as well, and so the change represents improved convergence of the policy.

The other notable difference is that the area planted now increases at higher levels of the superrace and lower levels of storage. From the cost contours it can be seen that as far as immediate advantage goes, a lesser area of the crop needs to be planted to maximise $-k(i,u(i))$. As far as maximising $-V(\pi(X_1),X_1)$ goes, there is still a conflict between low-planting and high-planting options. On the one hand the lower pathogen vigour means that higher amounts of storage can be handed over to the next cycle without a disease penalty and on the other it is no longer so valuable to hand on this storage because, as figure 6.4.44A shows, the return is even less dependent than before on the possession of storage. The contours of $-V(\pi(i),i)$ show that the incentive to plant less than the baseline tends to win out at low superrace levels and the opposite tendency tends to prevail at high superrace levels, which is consistent with figure 6.4.44B.

As a result of these changes the output of the "general resistance" policy has also changed from the baseline output. The variance of the return has decreased, partly because of the lack of replication error and partly because of the lesser effect of the standard irruption of the pathogen. The mean return is also higher. As the superrace input increases, the storage output does not decline so much, and the superrace output increases more than it did - which the economy can now afford. Both these changes can be understood as a result of the tendency to plant more at higher superrace levels.

These differences in the cost function and the policy combine to make the evolution of the policy (Figs. 6.4.52-7) radically different from that of the baseline policy. Instead of the complex 12-cycle rotation there is now a simple (and also stable) two-point limit cycle in which the fluctuations of all variables, particularly the return, are much smaller. The superrace now never reaches the proportions of a serious problem. The starting point of this evolution is the no-storage, no-disease level and the relative simplicity of the system's behaviour shows up an interesting point. The storage of the crop is initially built up before the superrace becomes a problem (an action which would have no purpose if no disease were expected) and then declines to a low level before the rotation is set up but not to zero. This shows that the policy has in fact predicted correctly the course of events over the first several years, and has planted a sufficient but not excessive area, given the risk of an irruption.

In summary, the system with general resistance shows a strong qualitative resemblance to the baseline system, but also shows a number of quantitative differences. Several of these can be explained in terms of the decreased vigour of the pathogen on the crop, but there is a residue of puzzling features. Foremost among these is the

relatively simple appearance of the cost surface, which lacks the complicated fine structure of the baseline cost surface. There are two possible explanations for this, which are not separable on the available evidence. One is that the complexity of the baseline is an artefact of the optimiser resolution, that the policy was "unstuck" from incomplete convergence by the increase in general resistance and that the "general resistance" policy represents a better-converged result. The other is that the greater vigour of the baseline pathogen requires that the policy must think further ahead: not only must the economically contradictory effects of u(i) on -k(i,u(i)) and $-V(\pi(X_1),X_1)$ be resolved, but the effects on $-V(\pi(X_2),X_2)$ must also be given comparable weight. The correct explanation could be identified by further convergence of the baseline policy with a less restricted optimiser. Unfortunately this was not possible in the time available. The qualitative similarity between the two policies - the control of the pathogen by fallowing or rotation with another crop - remains.

6.4.4.3. The Unconstrained Control Policy

The relaxation of the constraints against using multilines did not provoke the algorithm into varying the crop composition. Thus the change in the level of general resistance in the plant did not cause rejection of the null hypothesis and the use of the pure superline was concluded to be locally optimal.

6.4.5. The Agribusiness System with a Cross-protection Effect

6.4.5.1. The Biology of the System

Implementing cross-protection in the model is somewhat more complicated than changing the level of general resistance because there are two parameters that can be varied: the area immunised by

the non-virulent spore and the length of time for which the effect lasts. In the spirit of section 6.2.1 a strong cross-protection effect was selected by using the following parameter values

Table 6-5. Cross-protection Parameter Values

FORTRAN Name	Source in Text	Value
H	f_{ij} (non-virulent, equation 4.3.5)	5×10^{-6} area units
ICPTIM	T_c (equation 4.3.5)	20 days

Thus the cross-protecting spore was taken as immunising an area of the same size as would be infected by a virulent spore, or as reducing the probability of infection over a larger area. Early runs of the program suggested that a protected period of the order measured by Littlefield (1969) had little effect on epidemic dynamics, and so a time period three times as long was used in the converged policy, to create a strong effect.

Figure 6.4.57A shows the interesting effect that would be expected on multiline yield as a result of cross-protection. It would be expected that the simpler races would interfere to some extent with the superrace on the superline. What is not clear is whether the improvement of the yield of the superline would be outweighed by the lower yield of the simpler lines that are supporting an extra pathogen's cover. The figure shows that, at least with this model, it is possible for the multiline to outyield the superline as a result of cross-protection. What is shown is a replacement series in which the AB plant is replaced by the A plant, under a constant dosage of 750 spores of each of the aa and aabb races. The pure A crop is completely devastated by this dosage, and the pure AB crop loses 86% of its yield:

the initial protecting effect of the infall of aa spores on the AB plant is small. However at mixture compositions where the A plant is able to escape most of the aa spores while still protecting the AB plant by generating some aa spores through the season, there is a definite though small overyielding of the mixture. To achieve this effect it was necessary to increase the area immunised by cross-protection by a factor of ten over the value used in the experiment. At lower values the effect completely disappeared. Thus while the parameter values used here are conjectural it seems that only a very strong cross-protection effect could cause overyielding, though the theoretical possibility remains open. Even without overyielding, though, cross-protection may tend to raise the yield of multilines over its value without cross-protection, and, with the concurrent interference with the superrace, this increase in yield may be conducive to multiline optimality.

6.4.5.2. The Constrained Control Policy

The starting point for convergence of the cross-protection policy was the converged baseline policy. When the cross-protection effect was introduced, the optimiser "noticed" it and made a number of small changes to the area planted and proportion sold variables for several grid points. Thus in the deterministic sense of significance there was a significant difference between the two policies. However the changes made were so small that a t-test was applied to test the hypothesis that the change in cost for those points where alterations had been made in the policy was significantly different from the change in cost produced by convergence. Since t=.047 (373 d.f.), no significant change had been made. The policy itself is reproduced in figures 6.4.58-62 as a matter of record, the differences not being detectable

to the eye. Figures 6.4.63-67 shows the contours of the return, with the difference that the contours are shown for a grid point at which 1000 spores per unit area of both the simpler races were present. Not even the highest level of natural infall of cross-protecting spores produces a noticeable difference in the cost surface.

The surprising feature of these results is contained in figures 6.4.68-73 showing the evolution of the policy through time. The surprise is that the extremely small changes in the policy (together with the effect of the natural irruptions of the simpler races) are sufficient to make a definite change in the evolution of the policy. The rotation is still complex, but it now repeats every nine cycles instead of twelve, and the wider fluctuations are somewhat damped. This change does not appear to have any biological significance, but is an indicator that the baseline rotation is fragile, in the sense that the details of its operation should not be regarded as having a very high importance. A practical version of the baseline control might well be produced on the model of the control for the general resistance system, by considering only the upper and lower limits of the area planted as control variables.

6.4.5.3. The Fungistat Policy

Allowing the control the opportunity to add spores of the simpler races as a cross-protecting "fungistat" did not provoke the optimiser into taking this measure. Since there is no doubt that the spores would have a beneficial effect on the yield, it must be concluded that the cost placed on the use of the spores (0.1 units of return per spore) made them too expensive for the benefit conferred. Thus any addition of the simpler races to a multiline with cross-protection in this system would have, as it were, to cross a 'cost barrier' before being optimal.

6.4.5.4. The Unconstrained Control Policy

Relaxing the constraint on the optimiser against the use of multilines did not result in the use of any multilines by the algorithm. The use of the pure superline could thus be concluded to be locally optimal, and also stable to the perturbation involved in introducing a strong cross-protection effect.

6.4.6. The Agribusiness System with Stabilising Selection

6.4.6.1. The Biology of the System

Like the implementation of a cross-protection effect, the simulation of stabilising selection requires some thought. Because the stabilising selection effect is often spoken of (e.g. Van der Plank (1963)), see section (2.3.4) as though it represents a deficiency of the superrace itself, rather than a property of the crop or the result of an interaction between races on the crop. It is probably not sufficient in modelling stabilising selection simply to reduce the vigour of the superrace on the crop. It is probably also necessary to alter the dynamics of the off-season so that the superrace is less effective then as well, and this is the approach that has been adopted here.

Reduction of the vigour of the superrace on the crop was achieved in the same way as it was for all races on the crop with general resistance: by halving s_{im} for all m and for i = 2. There was, by contrast, some difficulty in deciding how to change the off-season dynamics in the required fashion while remaining consistent with the approach previously adopted. The method used was to incorporate competition coefficients C_j into the summations of equation (4.4.1) so that

$$s(in)_{j(n+1)} = \frac{C_j s(out)_{j(n)} K_1}{K_2 + \Sigma C_j s(out)_{j(n)}} \qquad \ldots(6.4.1)$$

where $C_j = 1$ for $j = 1$ or 3

$= 0.5$ for $j = 2$

Two important properties of equation (4.4.1) are preserved by this approach;

(1) the maximum output from the wild hosts is still K_1 spores

(2) the growth rate of small populations on the wild hosts is still K_1/K_2 for the simpler races.

The further properties of the relation as far as the superrace is concerned are that now

(3) it can, in the absence of competition, still saturate the wild host and produce K_1 spores the next season. However

(4) a larger initial dose of spores from the crop is required to achieve a 95% saturation of the wild hosts, because

(5) the growth rate on the wild hosts in the absence of competition is half that of the simpler races, and

(6) when the simpler races are present the proportion of superrace output to the next crop is lower than the proportion in the input to the crop.

Thus the superrace on the wild hosts can be considered (roughly speaking for the model here is crude) as capable of producing fewer spores per unit area than the simpler races, though able to cover any host area given sufficient time.

Figure 6.4.74A shows the result of this modification of epidemic dynamics on the growth of the superrace in the absence of competition.

The appearance of the pathogen as a factor that can depress yield takes longer than in the general resistance case. (The dependence of yield on superrace abundance is the same as in the general resistance case.) Also, the final equilibrium level of the superrace is still further reduced.

The question now arises of whether this form of stabilising selection does stabilise the pathogen population and in what way. Since on a pure crop of a simpler line the simpler race will drive the superrace to extinction while on a pure superrace crop the simpler races cannot exist, there ought to be an intermediate crop composition at which they coexist. What is not clear is at what level they will coexist: whether there will be a range of neutrally stable states as in Leonard (1969b), in whose model any initial relative proportion of races persists, given the correct crop composition, or whether there will be a single attractor state to which others tend, or whether some more complicated pattern will arise. It is also not clear whether there will be only one stable crop composition, as with Leonard's model, or whether there will be a range of compositions for which there is stability.

The difficulty of imagining the likely state of affairs at equilibrium comes about because two types of saturation effect have been superimposed on Leonard's simple parallel increase of races. The first to affect the increase of the pathogen from low levels is the saturation of the wild hosts, which will have a different effect on the dynamics from that of the later saturation of the crop. It would be possible to adjust the crop composition so that the superrace would be the numerical superior within the season while remaining the inferior on the wild host. There is thus the possibility of an oscillating process occurring.

In fact the system of this study can be stabilised by the choice of crop composition, but its behaviour is interestingly different from a Leonard equilibrium. For the superrace with halved vigour in a mixture of A and AB plants, there turns out to be a very restricted range of crop compositions in which a stable state occurs. This range lies in the approximate region $.595 < F(AB) < .599$, and the equilibrium is a single stable point where there is saturation of the system with about 350 aa spores and 293 aabb spores. (This value is varied by the exact crop composition.) This state attracts all trajectories passing through points where both aa and aabb are greater than zero. Along the line joining the equilibrium state to the origin the two races increase at the same rate. Points off this line initiate trajectories that at first go away from the equilibrium and then turn to it over a period of many hundred crop cycles. The length of time taken to reach equilibrium is the reason that the estimates above of the stable range in crop composition are not more accurate. It takes considerable computation to test for the existence of a stable point and to find its position. Figure 6.4.74B shows the trajectories of a number of populations started at various initial compositions on a crop of mixed A and AB plants where $F(AB) = .59835$. The arrows on the trajectories show the direction of motion through time. The first arrow is placed at the end of the first cycle: the last arrow marks 100 cycles. Most of the progress towards equilibrium takes place in the first five cycles.

This equilibrium also has a quite unexpected property in that it corresponds in the long term to a higher yield than the pure superline. Figure 6.4.74C shows the result of a simulated replacement series in which each point marks the yield in the twentieth year of a crop of the given composition, grown under the stabilising selection regime with

the ordinary values of P_{ext} and μ. At the right hand end the yield is that given by long-term use of the superline alone. Small additions of the simpler line make very little difference to the yield because the simpler race cannot increase its population on such small fractions of the crop. At the left hand end there is the equilibrium yield under the simpler race, which has excluded the superrace and, being more vigorous, depresses the yield much more than the superrace does in pure culture. At the composition corresponding to the equilibrium of the pathogen, however, the yield of the crop exceeds the yield of either of the pure lines, showing that the steady population of the simpler race interferes so markedly with the superrace on the wild host that it more than offsets the losses it causes in its growth on the crop. However at a slightly lower proportion of the superline, at which the superrace is driven to extinction, the yield reaches a maximum. Yield is apparently maximised over this time-span by the use of the minimum amount of the simpler needed to ensure suppression of the superrace. This effect occurs only when the races have settled down to their equilibrium population. In the first few years of the multiline the yield is lower than that of the pure superline. As the advantage of the simpler race decreases, the peak of yield decreases in size and moves towards the left hand end of the graph. A trace of it is still present in the baseline epidemic, where a 90%A : 10%AB crop yields about .6% higher than an 80%A : 20%AB crop after 20 years (but well below the pure superline yield).

The stable state for the system is both more and less fragile than a Leonard equilibrium. It is more fragile because slight changes in the composition of the pathogen population cause wide swings away from equilibrium. It is less fragile because all trajectories eventually reach a common equilibrium point and because there is a

range, although only a small one, of crop compositions for which an equilibrium point exists. It is also worth noting that this equilibrium point is on the boundary of the region reachable by the pathogen. Thus the equilibrium is reached by saturating the wild hosts and the crop.

It is not clear how general the practical and biological significance of this result is. The general point should be taken that without some mutual inhibition of the pathogen races on a multiline, an equilibrium will only be attained at a saturation level. Beyond this there is the suggestion that real systems which exhibit saturation may be more robust than the Leonard system in the sense that equilibrium is attainable for a range of crop compositions. A Leonard equilibrium would be quite unattainable in practice as the crop composition would have to be continually adjusted or else small errors in estimating the natural growth rate of the pathogen would result in some races going to extinction. The form of the saturation seems to be critical, and it may be that only the interplay of the two types of saturation - on the crop and in the off-season - allow an attractor to exist or a long-term rise in yield to occur in a multiline.

6.4.6.2. The Constrained Control

The policy for the system with stabilising selection is displayed in figures 6.4.75 to 79. At first sight the laws appear to be the same as those for the general resistance case. In fact there are small differences which can be discovered by close inspection: the average fraction sold at low storage has been very slightly increased and that at high superrace abundance has been slightly reduced. Besides this, the area planted at low storage has been increased. There does not seem to be a test to show whether these changes are statistically significant.

It seems more likely that they represent slight adjustments to the sticky optimiser caused by the changes in off-season dynamics. This interpretation is supported by the contours of return which, although they show quite noticeable changes in the relative merits of many sub-optimal controls, put the optima in the same positions as for the general resistance case.

The results of these small changes and the change in the off-season dynamics (presumably mainly the latter) are a further slight increase over the general resistance case of the expected return at the zero storage level and a reduction of its variance. The variance of expected return is also reduced somewhat at high superrace abundance.

Inspection of the contours of the return surfaces shown in figures 6.4.80-84 reveals that there has been a further slight shift of the regions of higher return in the direction that was observed in going from the baseline to the general resistance case. At low levels of the superrace the tendency to plant less and to sell a higher proportion of the available product in keeping with the lesser likelihood of loss of yield at these levels. At the intermediate levels of the superrace these tendencies gradually weaken, especially at the higher levels of storage. But the overall trend can be accounted for by supposing that the need to provide against the risk from further epidemics has been further decreased by the further weakening of the superrace.

Figures 6.4.85-90, which show the evolution of the policy through time, reveal, however, a marked qualitative change in behaviour from what seemed above to be a policy rather similar to that for the general resistance case. The most striking feature of this change is that the policy evolution is no longer an oscillatory one and that after some

initial adjustments which even out the variations in the return caused by the reduction of the initial storage level and the appearance of the disease, the system reaches a steady state which is then maintained indefinitely. There is some doubt about the reason for this abrupt qualitative change. The oscillations in the previous results may have been partly induced by the use of linear interpolation to derive the evolution of the optimal control. This could come about because controls covering a limited range in state space tend to be treated as linear or bang-bang controls. Thus some of the oscillations in the evolution of the controls might be spurious. However in the range of systems spanned by the disease-free case and the baseline epidemic we can form a ranking of the seriousness of the disease problems. The baseline epidemic is a more serious threat (marginally) than the cross-protection case which in turn is more serious than the general resistance case, which itself is more serious than the epidemic where there is stabilising selection with, finally, the disease-free state having the smallest problem. In this ranking there is a clear trend from more widely oscillatory to less widely oscillatory evolutions of the optimal policy. This is in accordance with common sense because the slower recovery rates of the less serious epidemics in the initial phase of each rotation allows storage to be built up before the super-race can have a marked effect on the yield. Thus we can have some confidence that at some point between the system with general resistance and that with stabilising selection a threshold was passed on one side of which rotation is necessary and on the other side of which the disease problem has decreased sufficiently that the farmer can coexist with the pathogen in a steady state.

This steady state can with advantage be compared to the multiline equilibrium described in the previous section. In the pure superline

equilibrium the area planted is about seven per cent lower than the area which would be planted to the superline in the equilibrium multiline. This reduction is much greater than the effective reduction in area experienced by the superrace in the multiline as a result of the presence of the simpler race. Consequently, the abundance of the superrace in its steady state is only about half its abundance in the multiline equilibrium. The further consequence of this is that the yield per unit area of the superline crop never drops below about three hundred and seventy units per unit area. The total yield of the superline is about one hundred and ninety units compared with the total yield of the multiline equilibrium of three hundred units. However in the multiline case the variable costs would have been about twice as great. A comparison of the two steady states is thus reasonably favourable to the multiline case.

6.4.6.3. The Unconstrained Control

The relaxation of the constraint to use only the pure superline did not result in the optimiser introducing any of the simpler lines into the crop. The question thus arises of why one of the apparently more favourable multiline compositions from figure 6.4.74C was not used. In particular, since there has been so much discussion of questions of stability in multilines, it is of interest to see whether there is some unobvious disadvantage attached to the use of the equilibrium composition.

The results of an attempt to shed some light on this problem are shown in figures 6.4.90.1A-D. Strictly speaking a comparison between the use of the equilibrium composition of the crop and the superline line optimal policy can only be made by deriving the optimal policy in terms of the area planted and proportion sold for the equilibrium crop

composition but this was not possible in the available time. Instead the figures show the result of growing the equilibrium crop for 50 cycles given the standard probabilities of irruption and mutation. For this simple policy it is assumed that the total area is planted and that everything harvested is sold at once.

It can be seen that the yield is much more depressed than in figure 6.4.88, and that the two pathogen races increase in step with each other, reaching equilibrium much faster than in figure 6.4.74B. The equilibrium which they reach is higher than that of figure 6.4.74B in the superrace and lower in the subrace. The reason is that in figure 6.4.90.1 there is a continual influx of equal amounts of the two races from outside the system, due to the ordinary stochastic pathogen behaviour, and this reinforces the superrace which has an advantage on the crop to balance its disadvantage during the off-season. The influx also drives the system to this dynamic equilibrium faster than would otherwise be the case. The critical comparison comes in figure 6.4.90.1D, where the return collected at each cycle is shown. The initial phase for the equilibrium crop is relatively poor in return because the inflexible control floods the market. However in the final state the yearly return of the multiline equilibrium crop is slightly higher (compare figures 6.4.87 and 6.4.90.1D).

The most obvious conclusion from this result is that the equilibrium crop is in the long run a better crop to grow than the pure superline. However this conclusion would be premature, for reasons connected with the behaviour of the trajectories in figure 6.4.74B and with the stochastic behaviour of the system. What is shown in all the graphs of policy evolution is the behaviour of the system if the expected amount of the pathogen is the input to each stage (see section

6.4.1.2.2.3). In the case of the equilibrium this trajectory, which passes along the direct line from the origin to equilibrium, is very improbable. As figure 6.4.74B shows, if at low pathogen abundance an irruption of one race arrives but not the other, the system will be temporarily destablished. On half these divergent trajectories the superrace increases suddenly and the yield drops. Thus the expected crop return is probably lower than the return along the trajectory shown in figures 6.4.90.1A-D. Paradoxically, the equilibrium crop will yield a less stable return than the conventional pure superline.

The most balanced judgement is probably that in introducing stabilising selection the seriousness of the disease problem has been so reduced that there exists a wide range of only marginally sub-optimal policies, and that although the pure superline policy is locally optimal there may be other optima that are not greatly different in expected cost. Even the yield advantage of the highest-yielding multiline from figure 6.4.74C would be counterbalanced somewhat by this crop's sensitivity to irruptions of the simpler races as well as the superrace.

6.4.7. The Agribusiness with a Growing Plant

6.4.7.1. The Biology of the System

Because there is little guidance available from the current state of debate on the modelling of plant growth (section 4.3.3), the problem to be faced in introducing this particular kind of realism into the simulation are greater than in the previous experiments. The key parameters available to be varied are the initial size of the plant at Day 1 of the growing season and its relative growth rate (k in equation 4.3.13). The non-growing plant used in the previous experiments

corresponds to values of, respectively, 1 and 0. The question therefore, is how to vary these values in such a way as to increase realism while maintaining a system which is comparable with those used previously.

An initial attempt was to use the relative growth rate used in Figure 4. to mimic the data of Rawson and Hofstra (1969) and to choose the initial plant size so that the size at the critical point for yield determination was the same as that of the non-growing plant. However this choce of parameter values led to a plant that was totally unaffected by disease since almost all of the initial infall of spores died on the ground after missing the small area of plant tissue. This result emphasised the importance of choosing parameter values that allow empidemics on growing and non-growing plants to be compared.

Several criteria for making this comparison were considered. The two that received most attention were:

1. choosing parameter values so that the integral over time of plant size was the same in the growing and non-growing cases, thus in a sense normalising the size of the plant across the different models.
2. the functional criterion of choosing parameter values so that the damage done to the yield by the standard irruption was the same as for the baseline epidemic.

With two parameters available to be varied, it at first appeared that there would be a point satisfying both these criteria which would be a natural choice. Figure 6.4.91A is the curve of parameter values that satisfy the functional criterion 2. Unfortunately, the only point that this curve shares with the curve of points satisfying criterion 1 is that defining the non-growing plant. Accordingly, as

the functional criterion is strongly bound up with the stochastic nature of the problem, it was decided to take this curve as parametrising the growth of the plant and, using the Rawson and Hoffstra (relative) growth rate of 0.23 units per day, an initial plant size of .0112 was chosen. The maximum size which this plant attains in the absence of disease is 6.52 units at the critical point for yield determination.

Figure 6.4.91C shows the dependence of yield on pathogen input that results from the use of these parameter values; the effect of the pathogen on the yield is somewhat more severe than in the baseline case. The net result for the pathogen of a host whose size increases as the pathogen reproduces is to permit a much greater total area to be occupied by the pathogen, hence the further depression of yield at higher pathogen inputs. This acceleration of the pathogen can also be seen in figure 6.4.91B where the onset of the disease problem is more sudden than for the baseline epidemic.

6.4.7.2. The Constrained Control

The control laws for the agribusiness with a growing plant are shown in figures 6.4.92-96. These laws are similar in type to those described for the previous experiments. One distinctive and explicable difference lies in the reduction of area planted at low storage compared with the baseline epidemic. This appears to be caused by a greater probability of loss from disease in this system as compared with the baseline system. Inspection of the contours of return in figures 6.4.97-101 shows that the greater effective vigour of the pathogen detected in the previous section has resulted in steep

decreases in return as the area planted increases away from zero at low storage. Otherwise the contours resemble those of the baseline system in their complexity and in the presence of fine detail rather than the general resistance/stabilising selection systems.

As might be expected the process of calculating the evolution of the policy through time which has in the other experiments been the most sensitive detector of the effects of small changes in the return surface and in the control laws has produced a result which is noticeably different from the baseline case. Figures 6.4.102-108 show an evolution which is of the apparently chaotic type found in the baseline game rather than the regular limit cycles encountered elsewhere. In keeping with the hypothesis of section (6.4.6.2), this most vigorous of pathogens provokes a rotation in which the area planted and the return from each crop fluctuate more widely than in the other experiments.

6.4.7.3. The Unconstrained Control

When the constraints against using multilines were relaxed the algorithm responded by introducing small amounts of the simpler lines at 9 of the 625 grid points. In no case was the proportion of the superline decreased below 98 per cent and the cost reductions were not significant. The use of multilines was therefore significant in the deterministic sense though not in the statistical sense. Since in each case their use was associated with small changes in the area planted (of the order of 3 per cent) it was open to question whether their use was a result of the kind of interaction between control variables discussed in section (6.3.1.1). The multiline points occurred

exclusively at storage values of 125 units and in cases where at least one of the simpler races had an abundance of 1000 spores. It was observed that in each case the yield of the crop had increased by an amount of the order of .1%.

There was thus a novel factor in the response of the growing plant to the pathogen, not possible in the non-growing plant, which would undoubtedly make it favourable to use multilines, since while yield was being increased the output of the superrace was being decreased by similarly small proportions. Upon investigation, it was found that the yield of a replacement series, such as that shown in figure 4.2 can increase slightly, according to this model, for small inclusions of the simpler lines. The effect decreases as the growth parameters of the model move along the curve of figure 6.4.91A towards the non-growing plant and increases as the parameters are changed in the opposite direction. However, within the range shown in figure 6.4.91A the effect only occurs at very high densities of the simpler races, is never greater than a few tenths of a per cent, and appears only in almost pure superlines. It apparently represents a slight response of the superline to the depression in size of the simpler line which more than cancels the loss of the simpler line. In real terms, this might correspond to a crop that was being planted at too high a density so that the relatively lightly affected superline is partly released from competition by the heavier effect of the double inoculation on the simpler line. Whether this effect can be expected to occur in real multiline crops is debatable; in section (4.3.3) it was remarked that the evidence for compensatory growth in crop/pathogen systems is sparse. However it is significant that an attempt at the introduction of realism in relation to plant growth into the current level of debate

has been the only factor of those "significant" biological factors studied (section 6.2.3) which has induced the optimality of multilines for the agribusiness. It is therefore likely that relatively simple arguments based on a non-growing plant are not adequate as a basis for deciding on optimal gene use.

6.4.7.4. The Constrained Game

Although the multilines discovered in this experiment were not significant statistically or practically, the test for robustness in the minimax sense was applied. Figures 6.4.108-109 show a similar pattern for the adjustment of the stochastic parameters to that observed in the baseline game. For the same reasons as in the baseline game, intermediate values of P_{ext} and μ appear both in the "moves" of the pathogen aspect of the game and in the evolution of the game (Figs. 6.4.111-116). The significance of the evolution of the game is therefore subject to the same limitations as in the case of the baseline game.

6.4.7.5. The Unconstrained Game

When the constraints against the use of multilines were relaxed in the game, similar statistically non-significant introductions of the simpler lines to those in the unconstrained control were made. These occurred at the same grid points as in the unconstrained control, and since at these points an area greater than zero was planted this choice of multilines would have been exposed to the effects of the variation in the stochastic parameters. Thus although the multilines

used in this experiment were not statistically or practically significant, their use was locally robust in the minimax sense, which further supports the conclusion reached in section (6.4.7.3).

6.5. The Subsistence Criterion

6.5.1. The Economic Parameters

As with the agribusiness criterion the economic parameters for the subsistence criterion, listed in Table 6-6, were chosen so that in the absence of disease the subsistence farmers would be in a state of abundance. Abundance here, because of the utility function used (Equation 4.6.7) is defined by the possibility of maintaining the preferred body weight and no account is taken of the amount of work needed to achieve this. In the next section it will be shown, however, that the subsistence farmers need only work a small fraction of the year to attain their preferred body weight, even after a period of starvation.

Table 6-6.

FORTRAN Name	Source in Text	Value
FIXCOS	c_1 (equation 4.6.10)	365.25 crop units
WORMUL	k_3 (equation 4.6.9)	5.0 crop units/day worked
HAFWOR	k_4 (equation 4.6.9)	100 days
CONRAT	k_5 (equation 4.6.10)	1.5 crop units/day worked
SPOCOS	(section 4.6)	0.1 crop units/spore added/unit

Thus the subsistence farmers have the capacity to produce much more than they eat in any year. However, because of their utility function, there is a cost on the carryover of storage from year to year, and the upper limit on what they can store (section 4.6.2) makes it impossible for them to exist for more than a year if they do not work, even if they start from that upper limit.

6.5.2. Subsistence Farming in the Absence of Disease

Figure 6.5.1 shows the simple rule that governs the optimum work rate for the subsistence farmers in the absence of disease. The consequence of the abundant environment described in the previous section can be seen in this rule. While at low storage the farmers must work sixty days a year in order to maintain their preferred body weight, at high storage levels they stop work almost entirely and live on their stored fat. This rule is so effective that the expected cost for any state is zero and so is the state output (i.e. the preferred body weight is always attained) and so these variables are not plotted.

There is only one explicit control variable mentioned in the subsistence criterion: the work done in days per year. However there was an unintended side effect of the introduction of a cost for the

addition of spores of the simpler races. The use of the same "small" cost for spores introduced in the agribusiness case in effect introduced another control variable. An overweight subsistence farmer can move more rapidly to the preferred body weight by adding spores that have no epidemiological function to the system. Thus, in this one-product economy, the spores play a role not unlike sports or crafts. The square-law utility function magnifies the effect of using spores non-linearly and the result can be seen in figures 6.5.2 A-C. These figures show the cost surface for the subsistence farmer in the same manner as the return surface for the agribusiness was shown in section (6.4). The starred contour now encloses the regions where the cost is less than one per cent of the distance between the lowest and the highest cost above the minimum.

Figure 6.5.2A, showing the cost in the present year of the control action, indicates that at low storage there is a single optimum work rate at the sixty day level indicated by figure 6.5.1. At lower levels of work the cost mounts rapidly to the maximum level. At higher rates of work body weight increases more gradually and reaches a maximum at about 250 days worked. After this point, diminishing returns set in and the farmer expends more energy than he receives for every additional unit of the balanced diet produced. At work levels above the optimum it is possible to decrease the cost by adding spores to the system, but this was not permitted in the convergence of this constrained control law.

At higher storage levels the optimum amount of work moves steadily to the left and at the highest storage level is almost against the left hand end of the axis. The surface for storage equal to zero shows that at the preferred body weight itself the optimum work range is very sharply defined, as the cost contours change their curvature

from being convex towards higher work levels (corresponding to excess body weight) to being convex towards the lower levels where weight is being lost below the optimum. The change corresponds to the work level at which the addition of spores ceases to become a means of reducing body weight and becomes a drain on energy.

Because of the great effectiveness of the control law on the disease free state the surface of cost for all subsequent years (Figure 6.5.2B) is almost completely flat, but the contour gradations that can be seen form a kind of mirror image of those for the present cost. In this mirroring there is a resemblance to the agribusiness criterion. Figure 6.5.2B shows that the maximum future cost is usually located at the same control values that provide the minimum present cost. This appears to betray a preference that was not explicitly built into the model for achieving the preferred body weight by the least possible amount of work, and may represent residual convergence error only, since the magnitudes involved are very small. The flatness of these contours is evident from considering the contours of total expected cost, which are in this case indistinguishable from those of the present cost.

It should be observed that though the contours of cost are simpler than the contours of return for the agribusiness, they are also non-convex. Because of the possibility of adding spores there is only one global optimum, but if the farmer is not allowed to add spores he has a choice between two local optima in most cases, one at the maximum work rate and one at a lower work rate. Because the policy for the disease-free state was converged quite gradually from an initial control of zero days worked, the global optimum at the lower work rate was found in each case. However, in later experiments, the presence of the alternative optimum had a significant effect.

Figures 6.5.3 and 6.5.4 show the evolution of the simple control through time from an initial condition of maximum body weight. It can be seen that the initially low work rate rises after one season to its constant final value and that, correspondingly, the preferred body weight (scaled to zero) is reached and then maintained.

In summary, the optimum behaviour for the subsistence farmer in the absence of disease is to work as little as is necessary to attain the preferred body weight in any given year. Because of the diminishing returns to agricultural work and the consumption of the balanced diet in that work, the farmer may find himself in a situation in which he must work harder to reduce his body weight. This fairly paradoxical effect has its parallels in our own society, and should not be dismissed as totally unrealistic, but it did not affect the calculation of the control law displayed in this section.

6.5.3. The Baseline Epidemic

6.5.3.1. The Constrained Control

The biological parameters of this epidemic are the same as those of the agribusiness epidemic. Figures 6.5.5 A-B show the control policy. It can be seen that the optimum amount of work has been increased at all levels of storage, but especially at the higher levels, which is only to be expected, given the average reduction of yield. Figure 6.5.5B shows that most of the variability in work rate is now explained by the disease abundance, and that storage only plays a role in determining work done at low superrace levels.

The shape of the curve in figure 6.5.5B can be explained in the following way: at low superrace levels the expected yield is determined by a relatively small loss due to irruptions in a clean crop, and the

work needed to achieve the preferred body weight is only slightly above that needed in the disease-free case. At the second level of superrace abundance the yield decreases (to about 170 units) and it is no longer possible to reach preferred body weight from the lower body weights: working about 145 days gives the minimum net loss of 80 units of the balanced diet. Farmers at all levels of storage work this amount; thus the farmers at high storage levels are beginning, unlike those in the disease-free case, to take thought for the morrow by retaining as much fat as possible. At higher superrace levels the minimum net loss of weight becomes larger and larger, the farmers run into diminishing returns sooner and the optimum level of work falls.

Figures 6.5.6A-B show the dependence of the expected cost on the state variables. As might be expected, the possession of storage does reduce the cost somewhat, and also reduces the variance of the expected cost. However most of the variation in expected cost is determined by the superrace, which causes a steady rise in cost very different from the minor changes in return caused by the superrace in the agribusiness. Figures 6.5.7A-B show the storage output which is a smooth function of the two state variables. The fact that the expected storage output is almost at its upper limit at low superrace abundance indicates that the subsistence farmers attempt to reach their storage limit, as an insurance against loss, wherever possible. Figure 6.5.8 shows the expected input-output relation of the superrace (the control cannot affect this output and so the superrace cannot be made a function of the storage).

Table 6-7. ANOVA for Subsistence Baseline

Control variable: Days Worked Per Year

	Sums of Squares	d.f.	Mean Squared Error	F-Ratios
Storage	26795.	4	6698.7	81713**
Superrace	6132600.	4	153320.	1870200**
Interaction	106730.	16	6670.5	81370**
Error	28.69	350	.08198	
Total	746810	374		

** = Significant at the 1% level

When the shape of the cost surface which has been dictating these results is examined an interesting point becomes evident. Despite the relatively simple form of the surface compared to the surfaces of section (6.4), a catastrophe has occurred. As the amount of the superrace increases at storage = 300 from zero to 250, the hump of cost at higher work levels which results from extra weight accumulated has been converted into a depression and the local optimum at 365 days worked has been lost. This catastrophe is so much simpler than those of the agribusiness surface that it can be identified as a cusp catastrophe (Zeeman, 1976). In the terminology of Zeeman, the dimension of superrace abundance is a control axis and the dimension of work done is a behaviour axis. (N.B. The contourless plots are those in which the plotting routine encounters a surface in which the variation in cost is less than 10^{-6} of the minimum cost. Economically these correspond to cases where the storage output is at the lower limit whatever the amount of work done).

In economic terms the catastrophe occurs when the yield drops below the level at which it is possible to attain the preferred body-weight from a given level of storage, and for the baseline epidemic

this occurs between superrace = 0 and superrace = 250. When the catastrophe occurs the farmer's pattern of work changes from an attempt to increase production in response to the pathogen to a (futile) holding action against starvation, and this is illustrated by the evolution of the policy in figures 6.15.4-18. Storage begins at a high level and the work rate at this time is fairly low, though more than enough to keep body weight above the preferred level. As the yield decreases the work rate is increased and the high storage maintained - causing the early peak in the cost record - but when the yield drops and the catastrophe occurs the work rate drops and the storage falls steadily until the starvation level is reached in cycle five. As the storage drops through the preferred level there is one cycle of low cost, but by cycle five the cost is the constant (arbitrary) cost associated with a year's starvation at the lowest level of storage. During the course of the evolution the superrace follows its natural pattern of increase unchecked.

In summary, the optimal pattern of work for the subsistence farmer facing the baseline epidemic is a brief flurry of activity which defers starvation for two or three cycles. After this active period they lapse into a semi-apathy that conserves body weight as long as possible. The contrast with the power over the disease given by the agribusiness' option of rotating the crop is very sharp.

6.5.3.2. The Unconstrained Control

Relaxing the constraints on the optimiser to permit the addition of spores and the use of a multiline did not result in the use of either option. Thus apparently for the baseline subsistence case the use of a pure superline is locally optimal, and the requirement of maximising net food output in the face of the severe epidemic does not

permit the downward adjustment of net output via the energy-consuming use of spores.

6.5.3.3. The Constrained Game

The result of the constrained game can be described quite simply since, in the absence of the interaction of area planted with P_{ext}, the maximising phase of the optimiser set P_{ext} uniformly to 1. At the same time μ was slightly decreased in many cases, presumably for the same reason suggested in section (6.4). However the effect of changes in μ was so slight compared to the effect of changes in P_{ext} that in a number of cases it was left unchanged. Thus once again the correct outlook for a pessimist was to assume that the superrace is bound to arrive at once.

6.5.3.4. The Unconstrained Game

The relaxation of the constraints to the pure superline in the baseline game was responsible for two interesting phenomena. The first of these was that multilines were introduced. The second was that the convergence of the policy abruptly became very much slower. These two events proved to be connected in the following manner.

When it was observed that the rate of convergence was extremely slow, an examination of those grid points where the game was proving most resistant to convergence showed in a number of cases that there was anomalous behaviour. It appeared that the cost to the farmers increased without the stochastic parameters having changed. Further investigation showed that in fact P_{ext} was shuttling between 1 and 0 in successive maximisations, which was in contradiction with the result of the previous section.

The explanation began with the discovery that the converged solution of the constrained game had not been properly applied to the starting condition of the unconstrained game. As a result some initial values of work done were excessive, and for these values the most harmful effect of varying the stochastic parameters was to decrease P_{ext} since when this was done the yield rose and so did the farmers' body weight (at a constant work rate). For some such grid points the response of the minimiser was to introduce multilines, apparently as a means of reducing the yield, or to drive the work done to its upper limit and to add spores to reduce body weight directly. For a proportion of these responses the position on the cost surface was now such that the minimiser again increased P_{ext}, and a cycle began.

While this illuminated the reason for the increase and decrease of P_{ext}, it did not explain why there was no sign of convergence in the cycle. The explanation for this turned out to be one that depended in a quite basic way on the manner in which the stochastic introduction of epidemics was modelled. For a case in which a pure line (including a pure simpler line) is planted the cost is given by

$$E(V) = P_{ext}\, V(+) + (1 - P_{ext})\, V(-) \qquad \ldots(6.5.1)$$

where $V(+)$ is the cost if an irruption occurs and $V(-)$ is the cost without an irruption. This is a linear function of P_{ext}, and if a cycle of the type described above begins, its linearity means there is no unique solution for the game.

This can be seen by considering the two settings of the minimiser aspect of the control that are involved. Let u_1 be the control that minimises $V(+)$ and u_2 be the control that minimises $V(-)$. If $V(+, u_1) < V(-, u_1)$ and $V(-, u_2) < V(+, u_2)$ the cycle will begin. There will be an intermediate control u_3 for which $V(-, u_3) = V(+, u_3)$, and the line of cost connecting them will be level. This line is a set of "mini-

neutral" solutions of the game, but because it is not strictly concave for all P_{ext}, the system has no minimax point, and the cycle will continue indefinitely unless the convergence of other parts of the policy changes the cost surface.

Thus it was a matter of chance that the constrained game had converged, and though the presence of the multilines in the unconstrained game was intriguing they were unanalysable because of the very complex pattern of interactions with P_{ext} and the other control variables, and because of the uncertainty about whether the game would eventually converge. When it was found that the constrained games started for some of the other experiments exhibited the same behaviour even more frequently than the baseline unconstrained game, it was decided that the further application of the minimax test to the subsistence controls should be ceased.

In summary, the application of the game technique to the baseline policy suggested that the pessimist should, as in the agribusiness case, assume that the pathogen is certain to irrupt. However, in this case, if he has not decided on the correct amount of work to do he may be the loser by his pessimism. The result of the unconstrained game suggested tantalisingly that the pessimist should also plant multilines, but for the reasons set out above this suggestion could not be checked. The process which hindered the convergence of the game was in itself a rather interesting and unforeseeable interaction, which could only come about because of

(1) the diminishing return of yield for work which creates the local optimum at high work levels

(2) the utility law which makes a low P_{ext} disadvantageous to the farmers at these high work levels

(3) the way in which the stochastic introduction of new races is modelled.

It was especially the excessive simplicity of (3) that made it impossible in some crops to find a minimax solution and future workers should consider carefully the way in which they handle the stochastic element of optimal gene use.

6.5.4. Subsistence with General Resistance

6.5.4.1. The Constrained Control

The biological parameters for this epidemic are the same as for the agribusiness case. In the absence of rotation their effect in this criterion is to convert the insupportable baseline epidemic into an epidemic which can be lived with. One way of looking at the reason for this conversion is simply that the yield in the general resistance epidemic never drops so low as in the baseline case. A more abstract way of looking at the change is that the line in state space along which the catastrophe in the cost surface occurs has moved a considerable distance.

Figures 6.5.19A-B show the control laws for the general resistance policy. The first point to notice is that the mean work done for each storage level has reverted to something like the disease-free pattern, but at a higher level. This is because at the higher levels of storage it is once more possible to live almost entirely on stored food with the shortfall being made up by a small amount of labour. The second point to notice is that the work done now increases at higher storage levels instead of falling away, and that the work done at these levels is no longer wholly determined by superrace abundance. This is because diminishing returns no longer set in so early in the rise of

the pathogen, but a better explanation will be made below, in terms of the path of the catastrophe.

Figures 6.5.23-27, showing $V(\pi(i),i)$, illustrate the path of the catastrophe. (The effect of $V(\pi(X_1),X_1)$ is once again negligible and so only the total cost is shown.) Whereas in the baseline case the transition between the two types of surface came between superrace = 0 and superrace = 250, there is now a transition path which begins at zero storage at a point between zero and 250 superrace and ends at 1000 superrace between 150 and 300 storage. Thus at all superrace levels at least some of the higher storage levels permit the farmer to reduce his cost by working harder. These are the levels at which the surface has a cost maximum in the region of work = 180, and it can be seen that for constant storage the starred contour on these surfaces moves slightly to the right as the pathogen increases.

The result of the ability to make an effective response to the epidemic can be seen in the expected cost and state outputs of the policy. Figures 6.5.20A-B show a much reduced expected cost (note the different scale on the cost axis) which only begins to rise at high superrace levels and which at high storage levels is independent of the pathogen. Figures 6.5.21A-B show that although at high superrace levels preferred body weight is not attained on the average, the depression is comparatively small, and that at high storage levels preferred body weight can be attained and the need is no longer felt (at these levels) to lay by storage for the next year.

The evolution of the policy through time (Figures 6.5.28-31) confirms the impression that the general resistance epidemic is tolerable. The work rate in the absence of the pathogen is low because the initial condition includes high storage. As storage falls the work

rate rises and there are signs that the disease-free equilibrium would be established for a time, but when the yield begins to drop the work rate rises, and the system reaches an equilibrium at a slight positive level of storage. Why this slight positive level should be maintained is not clear. Another unclear point is why the brief peak in storage in year 4 that caused the cost to peak is produced. This peak may represent the tendency of linear interpolations among controls to produce oscillations in the control as it evolves.

In summary, the lowering of the vigour of the pathogen produces a qualitative change in the behaviour of the control; it is now possible to maintain preferred body weight by working harder than in the disease-free case. The subsistence farmers are in a state not of abundance but of sufficiency. The final equilibrium is, however, in the region of state space where an increase in the vigour of the superrace would provoke a vicious circle of reduced work and lowered storage, so this state of sufficiency is not structurally stable.

6.5.4.2. The Unconstrained Control

The relaxation of the constraints against using the simpler lines resulted in the introduction of multilines at 89 of the 625 grid points. Figures 6.5.32A and B show that these multilines were used at the lower levels of storage and superrace abundance. Examination of the results showed that the introduced multilines fell into two classes, one in which the amount of the superline used was decreased by less than a tenth of a per cent (twenty cases) and those where at least 30 per cent of the crop was made up of one of the simpler lines. The former category of multiline was not associated with a statistically significant cost decrease and in most cases apparently represented an adjustment by the optimiser of the form discussed in section (6.3). Those multilines in

which less than 70 per cent of the superline was planted were distinguished from the pure line parent points by the fact that

(1) at the economic level both $k(i,u(i))$ and $V(\pi(X_1),X_1)$ were reduced

(2) at the state variable level the storage was decreased and the superrace output was reduced

(3) at the epidemiological level the yield was also reduced.

In a number of cases spores of the simpler races were added to multilines containing plants susceptible to them.

Figure 6.5.33 shows that the tendency in the constrained control to increase storage at the superrace = 250 level has now been removed and figure 6.5.34 shows that the presence of the multilines has quite noticeably reduced the average superrace output. Figure 6.5.35 shows that the work rate at the superrace = 250 level has decreased on the average.

Within this pattern of adjustments in the derived policy there is a further sharp distinction between the use of multilines at the lowest storage level and at storage = 150. At the lower level the cost reductions are much larger (about 10^4) and storage is reduced by approximately 115 units. At the higher storage level the cost reductions are smaller (about 10^3) and storage is typically reduced by three or four units. At both storage levels the use of multilines reduces the superrace output by, on average, about 100 spores below the level for the constrained policy.

The principle that links these separate phenomena is that the multilines, by reducing the future abundance of the superrace, not only reduce the expected cost for future years, in years where the superrace has not yet saturated, but also removes the need to

accumulate large storage against the future effects of the disease. At the lowest levels of storage where the constrained policy tends to accumulate more storage as a response to the lower, more manageable, disease levels the effect is most marked, and the work rate drops from the upper limit by about fifty per cent. At storage = -150 where such large accumulations were not made in the parent policy, the work rate more often increases than decreases, and the storage output is maintained close to zero. At this higher level of storage the decrease in cost is dominated by the decrease in $V(\pi(X_1),X_1)$, which is almost wholly due to the change in the superrace output. At storage = 0, with one exception, the multilines are non-significant and the reason for the exception is not clear. For one grid point, where storage = 0 and the superrace = 250, a six per cent reduction in the cost has been achieved by increasing the work rate and planting sixty per cent of the crop to one of the simpler lines. Most of this reduction in cost was the result of the reduction in superrace output, which was independent of the change in work rate. However, other points with similar amounts of the simpler races did not attempt this method of reducing the cost, presumably for some reason connected with the resolution of the optimiser. The category of non-significant multilines mentioned at the beginning of this section consists mainly of these neighbouring points. There is thus the probability that a significant reduction in cost could have been achieved by the use of multilines in more cases than were actually found by the algorithm.

The most difficult result to analyse in these results concerns the manner in which the simpler races interacted with the simpler lines in the multilines. Amounts as high as 230 spores per unit area were added to the system at various points. In one case at least, the addition was made for purely economic reasons because there were no simpler lines susceptible to the spores added. However in most

cases the added spores would have interacted with the crop itself. Figure 6.5.37 shows what appears to be the nature of the interaction. The trend of the regression line ($p < .01$) shows that as the total of the spores of a simpler race naturally present and artificially added increases, so the proportion of the simpler line decreases (if it is the major simpler element of the multiline). This relation shows that a balance is being attempted between the depression of yield and the depression of superrace output that can be attained at a given level of the simpler race. When the relationship between the food produced and the work done is calculated for the intrinsic yield (y of equation 4.6.9) of each multiline, it was found that the multiline had been chosen so that the maximum net production of food would bring the farmers close to their preferred body weight and that no trade-off between present hunger and future reduction of the superrace had been made. At the levels of output produced by most multilines even a small deviation from the preferred body weight (e.g. ten storage units) would outweigh the benefits gained by further reduction of the superrace.

In the analysis of the multilines the return cost surface cannot be displayed as for the pure line policies, because of the extra control variables present and because of the interaction of the "spores added" variable with the yield of the crop. Instead, we pass directly to considering the evolution of the policy through time. Figures 6.5.38A-B show an example of the evolution of a policy whose initial condition is one of the multiline points, where a large component of one of the simpler lines is used. It can be seen that by the second year of the policy a pure or almost pure superline is being planted and there has been no obvious delay in the increase of the superrace. A survey was made of all the points in the multiline policy to determine what was the greatest length of time for which multilines might be

expected to be used. This was done by following the evolution of the policy from every grid point until the proportion of the simpler lines used dropped below 10^{-4}.

A survey of this kind was not completely exhaustive because of the relatively wide spacing of the grid points. However it was found that there was no point in the policy for which multilines were used for more than one season. In particular, for the initial conditions that were suggested in 6.3.1 as being important, multilines were not used at all and the system passed to a pure superline equilibrium.

In summary, the introduction of general resistance to the subsistence criterion made it possible for the control algorithm to choose multilines as optimal control actions. These multilines were optimal because they reduced the output of the superrace and reduced the cost incurred by storage in the face of increasing disease. However the improvements made to the cost were fairly small - of the order of ten per cent in most cases - and although the average reduction in cost produced by the use of multilines was highly signficant ($t = 517$, 373 d.f.), their use was not practically significant in the sense defined in section 6.3.1. The general character of the use of multilines in this policy is transitional: they are used briefly in response to initial conditions that would not normally appear during the operation of the policy but which could only be the legacy of a sub-optimal policy. Thus although it has now been shown that the idea of using multilines is not self-contradictory, the result of this experiment does not encourage their development and use.

6.5.5. Subsistence with Cross Protection

6.5.5.1. The Constrained Control

The cross protection effect in this experiment was implemented in the manner described in section (6.4.5.1). In this experiment, as in its agribusiness counterpart, the natural infall of spores of the simpler races was "noticed" by the optimiser. The effect on the yield of the crop was an increase of as much as 1.38 units over the yield without cross protection for the case where all three races were present at their maximum abundance. At the same time, the superrace output decreased but the decrease was small, being less than one spore per unit area in the most favourable case. The response of the optimiser to these effects was to increase the amount of work done slightly, the largest change being 2.04 days in the case of the greatest change in yield. No pattern was discovered in these increases in the work rate and they apparently represent a combination of stickiness in the optimiser and adjustments to the change in storage produced by the slight increase in yield. Because the changes in the policy were so small, the control law, and the cost surface and evolution of the constrained policy, are not displayed here graphically as they are indistinguishable from figures 6.5.5-18. Thus while the subsistence farmers clearly derive a benefit from living in a situation where there is cross protection, even if for some reason they are constrained to use the pure superline, the benefit under this constraint is small.

6.5.5.2. The Subsistence Fungistat

The relaxation of the constraints on the algorithm to allow the subsistence farmers to add cross-protecting spores of the simpler races to the pure superline did not result in any such spores being used.

Thus, as in the agribusiness case, the use of these spores without the multiplier effect of the presence of simpler lines in the crop was too expensive for the benefit conferred.

6.5.5.3. The Unconstrained Control

When the constraints against the use of multilines were relaxed the algorithm introduced multilines at 210 of the 375 independently converged points. Many of these multilines involved substantial use of the simpler lines and there were a number of cases in which pure crops of a simpler line were chosen. Economically, these mutlilines correspond in almost all cases to an increase in the present cost k(i,u(i)), which was more than compensated for by a decrease in the future cost. The difference between the mean change in cost between the pure line and multiline points, and between pure line points in the parent policy that remained pure line in the derived policy, was highly significant ($t = 257$, d.f. = 373). Thus the multiline points in this policy were both deterministically and statistically significant.

At the state variable level the use of multilines was marked by a decrease in storage at all levels except the lowest, where a downward change would have been impossible. At the same time, the superrace output was greatly decreased by in excess of 600 spores per unit area at the highest input level. The yield, however, was uniformly decreased, despite this interference.

Figures 6.4.39-44 show how these effects were achieved. The use of the simpler lines was not affected by the storage level but increased very sharply at the higher superrace levels. At these same levels the work done decreased, on average, compared to the baseline and a substantial number of spores of the simpler races were added. Because of the lowered work rate the decrease in yield caused by the use of

multilines did not affect the storage output at the superrace levels in any major way. (At these levels the storage output was often -300 for both parent and derived policy.) The most striking effect is that the combination of the use of multilines and the addition of spores of the simpler races caused the average output of the superrace at the two highest levels of superrace input to be lower than for some lower inputs (Fig. 6.5.41). Thus there appears to be the interesting prospect of an oscillation in crop composition in which multilines are used periodically to reduce the superrace abundance.

In parallel with the procedure in the experiment on subsistence with general resistance, an attempt was made to interpret the interactions of the control variables at the multiline points. The significant relationship between the spores added and the composition of the multiline was not found in this case. Instead (Figure 6.5.44), the dominant relationship appears to be one between the amount of work done and the composition of the multiline. The regression in Figure 6.5.44 is highly significant ($p < .001$). The explanation of this difference appears to lie in the fact that in the general resistance case the cost catastrophe has not occurred and the farmers are able to adjust their body weight around the preferred level, while in the cross protection case the catastrophe has occurred and the farmers are engaged in the holding action described in section (6.5.3).

Figures 6.5.45A-D show the evolution of the multiline policy through time, beginning with a point where only a small fraction of the crop is planted to the superline. The initial depression of the superrace from its high level can be seen but the promise of an oscillation in crop composition has not been fulfilled. The multilines are only optimal where there is already a substantial level of the

simpler races but their use delivers the system into a state where the pure superline is optimal and the preconditions for the use of the multiline cannot occur. Thus in this case also, the multilines play a transitional role. A survey like that for the general resistance multiline policy shows that the transition period for the cross-protection policy is much longer with periods of two and three years of use of multilines being common and the example shown in the figure is one where multilines are used for five years. In the last three years of this period, however, the fraction of the simpler races used is less than two per cent, which makes the crop only technically a multiline

In summary, the introduction of a cross-protection effect which, although intended to simulate a strong effect, has not previously seemed strong enough to be of great importance, has had a more profound effect on optimal crop composition than the halving of the vigour of the pathogen. The result of this experiment suggests that cross-protection can in some cases be a factor conducive to the use of multilines, even in the absence of over-yielding. However, as in the general resistance case, the use of multilines reported here is transitional and therefore, though deterministically and statistically significant, probably not significant in practice.

6.5.6. Subsistence with Stabilising Selection

6.5.6.1. The Constrained Control

The implementation of the stabilising selection effect in this experiment was the same as in 6.4.5. As in the agribusiness case the closest comparison of the results is with the general resistance case. Figures 6.5.46A-B show the control laws for the stabilising selection system. These turn out to be indistinguishable from those for the

general resistance case. However from a comparison between figures 6.5.47A-B and figures 6.5.20A and B the lesser vigour of the superrace on the wild hosts in the stabilising selection case can be seen to have decreased the expected cost, especially at low storage levels and high superrace inputs. The comparison between figures 6.5.21A-B and figures 6.5.48A-B show that this lowering of cost is not a result of a change in storage output but is caused exclusively, as can be seen in figure 6.5.49, by a further depression of superrace output.

The cost surface for this policy, shown in figure 6.5.50-54, repeated the pattern observed in the general resistance case that $V(\pi(X_1),X_1)$ was only weakly affected by work done in the current year and so only the total cost $V(\pi(i),i)$ is shown. Just as in the general resistance case, the cost catastrophe never occurs at the higher storage levels so that the apathetic response to the pathogen need not take place. The change in the off-season dynamics of the superrace has in fact, in most cases, only affected the cost surface in detail. (The easiest way to see this is to note that the labelling of the contours has sometimes been shifted by the plotting program, even though the contours appear similar.) However at levels of high superrace and intermediate storage there has been a noticeable flattening of the cost contours. Figures 6.5.55-59 show that the changes described above, which are changes in the behaviour of the pathogen rather than in the policy, result in a stable state in which less work is done and less cost incurred in the stabilising selection than in the general resistance case. The key change is shown in figure 6.5.57, which shows that the steady state yield of the crop is about 235 units instead of 210 units in the general resistance case.

In summary, the change from general resistance to stabilising selection with the associated decrease in the effectiveness of the

superrace has meant only that the farmers need work less hard in the equilibrium state for the same result. Their situation is still one of sufficiency rather than abundance.

6.5.6.2. The Unconstrained Control

Relaxation of the constraints against using multilines in the stabilising selection case resulted in the introduction of multilines on a similar pattern to that in the general resistance case. Figure 6.5.60A shows that the multilines were introduced over a slightly wider storage range but that on the average a smaller proportion of the simpler line was used. As a result of this use and concurrent reductions in the work rate, which were on a similar pattern to the reductions in the general resistance case, the storage output for the stabilising selection case assumed the same form (Figure 6.5.61) as in the general resistance multiline. However, as can be seen from figure 6.5.62, the effect on the superrace was more marked, as indicated by a wider range of variability than in figure 6.5.34. No significant difference could be detected between the patterns according to which spores were added to the multiline in this case and in the general resistance multiline, and the same type of relationship between the amount of the simpler line planted and the abundance of the corresponding simpler race (Figure 6.5.65) was found. However, the absolute amounts of simpler lines planted, the slope of the regression line and its F value were all smaller than in the general resistance case.

When a survey was carried out of the persistence of use of multilines in this policy, it was found that by contrast with the general resistance case it was quite common for small amounts of the simpler lines to be planted for several crops in succession. Figure 6.5.66A shows the most striking of these persistent uses. The slight

depression in the proportion of superline between cycles two and six marks the only case known in this study in which multilines were introduced after the use of a superline. At other points in the policy much larger proportions of the simpler lines were used, but only for one or at the most two crop cycles. The effect shown here, despite its small size, is thus of some theoretical interest but it is not clear whether it is an artefact; the cost of this policy turned out to be .5% higher than the cost calculated during the evolution of the corresponding pure superline policy.

In summary, the multiline observed under conditions of stabilising selection are also transitional and although there are some slight indications of a tendency towards their more persistent use, it is not clear whether this is because of the presence of stabilising selection itself or because of the low absolute level of vigour of the superrace in this experiment. A comparison with the final equilibrium level of yield in the stable equilibrium of figure 6.4.90.1A shows that the subsistence farmers in this experiment were tolerating a yield depression greater than was necessary. However they (in the form of the optimiser) had no incentive to seek this higher yield, since they were able to maintain preferred body weight without it. Future investigation in this direction should centre on the use of multilines where there is stabilising selection and where a desired level of yield is only marginally attainable in the equilibrium state.

6.5.7. Subsistence with a Growing Plant

6.5.7.1. The Constrained Control

The means by which the growth of the crop was simulated in this experiment have already been described in section (6.4.6.1). As was observed in section (6.4.6), the epidemic on the growing plant provides

the most serious disease problem of those simulated, and this is also evident in the response of the subsistence farmers. Figures 6.5.67A-B show the control law for the subsistence farmers responding to this epidemic and it will be seen that the law is of the same type as that for the baseline epidemic. The differences between the two can be explained in terms of the greater effectiveness of the pathogen on the growing plant. The mean level of work done has been reduced slightly at low storage levels and decreased to zero at the highest superrace level. This indicates that the yield is so low that the farmers lose more than they gain in terms of the balanced diet by doing any work at all. The expected value of the cost has not been altered by these changes because it is the product of starvation in both cases which has the same cost, whatever pathogen produces it. The storage output for the new policy is also the same as for the baseline policy because for the states at which the yield depression is greater with the growing plant, storage is either independent of yield because of previous reserves or at its lower limit. However the superrace output now plateaus more rapidly and at a slightly higher level.

The contours of the cost surface shown in figures 6.5.71-75 are almost indistinguishable from those of the baseline epidemic at the superrace = 0 level, but by the superrace = 250 level a slight tendency for the optimal amount of work to decrease below the baseline optimum can be seen in both $k(i,u(i))$ and $V(\pi(X_1),X_1)$. At the superrace = 500 level, the range of work that does not drive $V(\pi(X_1),X_1)$ to its highest level has become more restricted and at the superrace = 100 level the optimum level of work shifts decisively to zero in accordance with the control law. The catastrophe in the cost surface occurs along a line in state space close to the corresponding line in the baseline policy. When the evolution of this policy through time is followed (figures 6.5.76-80), the increased severity of the disease becomes

evident once more. The response of the farmers is of the same form as for the baseline epidemic. However, since the yield decreases in three cycles to about half the yield of the baseline epidemic in its steady state, the farmers almost immediately stop working and the starvation level is reached in five years. (Note that in the final equilibrium state, the variability of the cost is reduced because irruptions can now have scarcely any effect on the yield.)

6.5.7.2. The Unconstrained Control

When the constraints requiring the use of the pure line superline were relaxed, the algorithm introduced small quantities of the simpler lines at all grid points where the superrace was not either at its lowest or its highest level and where at least some quantity of the simpler races was present. No spores were added artificially. Upon inspection, it was found that the average cost reduction resulting from the introduction of the multilines was significant, with $t = 25.7$ (d.f. = 373). This cost reduction was achieved by reducing both the present and the future costs. The superrace output was slightly decreased by .5 spores per unit or less, and where the storage was above zero level it was increased by, on average, about .05 units.

It was found that as in the corresponding agribusiness experiment, the use of multilines was associated with slight increases in yield, of the order of .01%. Since the use of multilines was so much more common in the subsistence case, however, it was possible to analyse their pattern of use. Figure 6.5.81A shows a regression of the small quantities of whichever simpler line was used in the greater quantity against the superrace, for the multiline points in the policy only. Within the superrace range for which multilines were used, the relationship is extremely strong ($p < .001$). Thus in this system the presence

of the superrace can actually increase the yield of the crop over a very narrow range of compositions. Presumably the mechanism is that in this range of disease intensity the simpler lines with their double disease load are more severely affected by each increment of the superrace, and the compensatory growth mechanism in the model expands the yield of the superline.

At the same time as the multilines were introduced, a number of small variations (less than one day per year) were made by the optimiser to the work done at each of the grid points where multilines were used. These adjustments were both positive and negative, and were quite variable, but there was an underlying pattern which is shown in figure 6.5.81B. Here the change in the amount of work done from the parent policy to the derived policy at each multiline point is regressed on the corresponding changes in yield. Despite the high variability, the regression is significant ($p < .01$) and the change is in the expected direction, namely that as the yield is increased the farmers take advantage of it by working harder. This phenomenon is interesting because it represents an adjustment which is evidently near the limit of resolution of the optimiser, and gives an impression of the importance of stickiness in the optimiser's operation.

A survey of the persistence of use of multilines in this policy indicated that multilines where the proportion of the simpler lines was greater than 10^{-3} were never used for more than one year at a time. Thus the multilines discovered in this experiment were also of a transitional character and not of practical significance. Because the evolution of the multiline policy was numerically so similar to its parent policy, it is not plotted here.

In summary, it was found that, as in the agribusiness case, a particular feature of the growth model used to simulate the crop caused

multilines to become optimal. The use of multilines and their effect on the evolution of the policy was small, consistent with the principles derived from earlier experiments with the subsistence criterion. While the over-yielding that was the basis of multiline use here may or may not occur in real crops, its occurrence here is a reminder that the inclusion of premises about plant growth in the debate on optimal gene use may radically alter the conclusions.

6.6. Comparison of the Results of the Experiments

A number of lessons concerning the likelihood that multilines will be optimal in practice can be drawn from the results of these experiments. The factors which may affect their optimality in practice can now be ranked in order of importance. The first of these is the nature of the optimality criterion. It cannot be said, on the basis of the simple models used here, exactly what type of criterion is most conducive to multiline use, but it is possible to say, for example, that criteria in which future yield is sharply discounted are less likely to induce optimal multilines. This conclusion is based on the fact that all of the multilines used except those for the model with a growing plant traded off a loss of yield in the current year against a saving in the future.

The second most important factor appears to be whether there is the possibility of a crop rotation. It was noticeable in the experiments on the agribusiness criterion that the gradient of cost with respect to the "area planted" variable was almost always several orders of magnitude greater than the gradient with respect to the crop composition variables. It is probable that no multilines were used in most of the agribusiness experiments simply because the trade-off between present and future return was made so much more effectively by

the rotation than it could be by the multilines. Of course, the rotation used in the agribusiness system was an extremely effective one in the sense that a single year of zero planting completely exterminated the disease. As Stackman and Harrar (1957) point out, many rotations in real crops are not as effective as this, either because the pathogen can last through a fallow or because it has a host range which covers more than one species. However, it seems likely that unless interactions that can cause over-yielding are found in real multilines, or unless rotation is for some reason impossible, multilines are unlikely to be optimal in practice.

The two factors just discussed were not originally intended to fall within the scope of the experimental study, which was intended to investigate the biologically significant effects mentioned in section (6.2). These biological factors evidently rank in importance below the nature of the criterion and the possibility of controlling disease through rotation. However, where they were given their best chance of influencing multiline optimality - with the subsistence criterion - their introduction in ways that tended, apart from the growth factor, to reduce the seriousness of the disease problem caused, in each case, the use of explicable and apparently optimal multilines. The question thus arises of whether there is some factor common to all these biological effects which is the critical biological factor in determining the optimality of multilines.

The first step in answering this question is the observation that multilines used with the non-growing plant were apparently all chosen by balancing a loss in current yield against a diminution in superrace output. Figures 6.6.1-3 were constructed in order to show how variation of the significant parameters considered in the experiments might affect this trade-off. (The growing plant, which induces

multiline optimality for different reasons in this model, is not considered.) In each case, what is plotted is the output in terms of the yield and the superrace, as the parameter is varied for a range of crop compositions. A single point in state space is taken at which all three pathogen races are abundant (superrace = 500, simpler races both = 1000), and for each parameter value a replacement series is seen in which the crop composition changes from 100% AB to 100% A. The output of each replacement series forms one rung on the ladder in the diagram. The superrace output is itself subdivided, with one axis showing the output from the crop before competition on the wild hosts and the other showing the input to the next year's crop after this competition. Thus effects which take place in the crop itself can to some extent be distinguished from those outside the crop/pathogen system. A projection of the surface onto each pair of axes is shown as an aid to judging perspective. (The visible side of the surface is the underside.)

The right hand end of the line representing each replacement series corresponds to the 100% AB crop composition. In figure 6.6.1, where the trade-offs for general resistance are shown, the replacement series on the extreme right corresponds to the pathogen vigour used in the baseline epidemic. For each replacement series going from right to left, the growth constant of the pathogen, s_{im}, is successively reduced by 10%, so that the final replacement series corresponds to a pathogen growing at only one tenth of the baseline rate.

For the baseline parameter value, the trade-off that is made is almost exclusively a trade-off in superrace output with little yield loss and saturation on the wild hosts plays an important role. For pathogen vigour equal to 0.2 of the baseline value, the trade-off is largely to yield, with saturation on the wild hosts playing no part because superrace output is so small. At the highest level of general

resistance, the pathogen does not reproduce at a significant rate; there is no trade-off, and yield is not depressed.

Thus at low levels of general resistance, the trade-off is exclusively in terms of benefit to the next year, owing to a reduced superrace input to the next crop, while at high general resistance, the ratio of yield loss to depression in the superrace increases and then becomes unimportant as the total size of the yield loss decreases. This relationship shows that at high levels of general resistance there is no point in using multilines, while at low levels even the large possible trade-off in the superrace input to the next crop does not bring a noticeable benefit in future yields. (What constitute "high" and "low" levels of general resistance will depend on the economic criterion.) At intermediate levels, however, it is possible to make trade-offs in which a sacrifice in this year's yield can be expected to bring a noticeable benefit in the next year. The parameter values used in the experiment on general resistance correspond to the sixth rung from the right on the ladder, and are in such an intermediate position.

In figure 6.6.2, which shows the trade-offs that can be gained by cross-protection, a much smaller range of outputs is covered, although the cross-protection parameter (in this case, the area immunised for a constant 20-day period) runs from zero to twenty times the value used in the experiment. The upper left hand side of the range of outputs corresponds to the highest level of cross-protection. It can be seen that even where over-yielding does not occur in the replacement series, the tendency of an increase in the cross-protection effect is to reduce the yield loss that must be suffered for a given depression in the superrace. Since, for the baseline value (zero) of cross-protection, the loss in yield through the use of the multiline is

very small, the cross-protection effect observed in the experiment, although also small, could have an effect on the economic value of the trade-off.

In figure 6.6.3, which shows the trade-offs that can be expected from stabilising selection as modelled here, the intensity of the stabilising selection effect is parametrised by varying the coefficient C_j of Equation 6.4.1 and the vigour of the superrace on the crop, s_{2m}, so that they decline in equal steps from 1 to 0.1. The result of this parametrisation is initially, where the selection against the superrace is small, similar to that at low levels of general resistance. However, it can be seen that the transition to unfavourable trade-offs, where much yield is lost for a small depression in the superrace, is much sharper. The intermediate levels of the stabilising selection effect, where substantial trade-offs in both yield and superrace depression are possible, have moved somewhat to the lower disease output levels, and this may explain why, in the experiment on stabilising selection, a more sustained use of multilines was optimal than in the general resistance case.

In summary, the figures in this section help to show why qualitative arguments for and against the use of multilines have not served to show whether they are optimal or not. It appears that the absolute level of vigour of the pathogen, and especially its superrace, are more important than has been supposed in determining whether multilines are optimal. This absolute level of vigour controls whether any other epidemiological variables can have a useful effect in making the trade-off between yield loss and superrace depression more favourable. At high levels of vigour, multilines are of no use, at low levels they are unnecessary. Where the absolute vigour of the pathogen permits, the presence of stabilising selection and cross-protection tends to

increase any favourable effect of multilines. The effects of including more realism by the modelling of plant growth do not appear likely to disturb these conclusions, except where, as in the experiments reported here, an interaction between the components of the multiline allows an advantage for the multiline in the present as well as in the future.

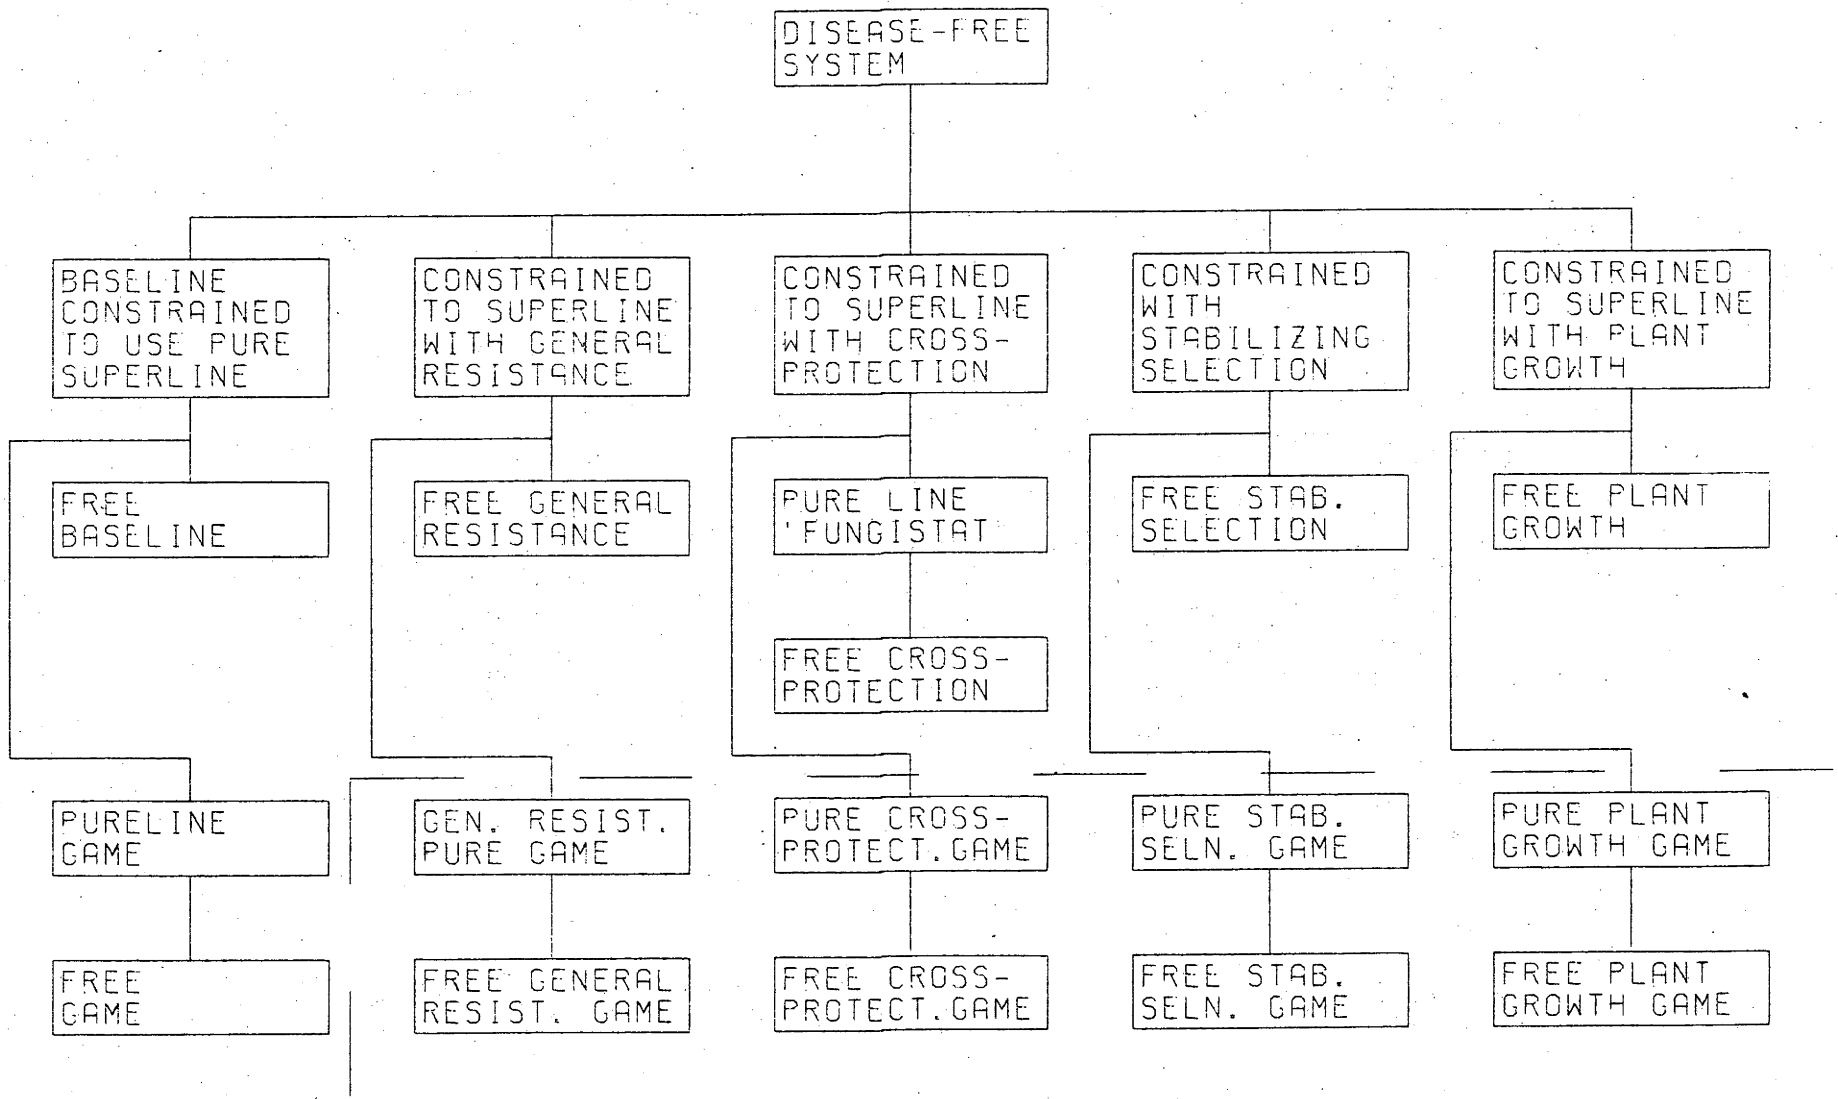

FIG. 6.1. EXPERIMENTAL DESIGN FOR BOTH CRITERIA

FIG. 6.4.1: DISEASE FREE AGRIBUSINESS
FRACTION SOLD VS STORAGE

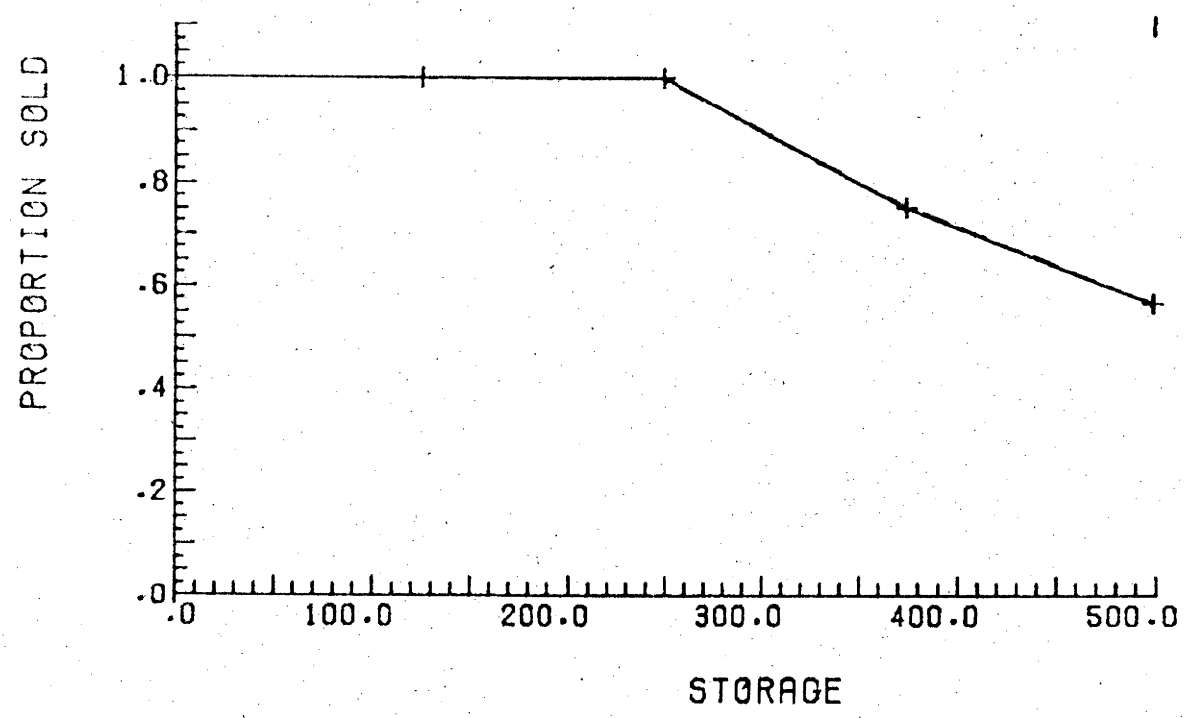

FIG. 6.4.2: DISEASE FREE AGRIBUSINESS
AREA PLANTED VS STORAGE

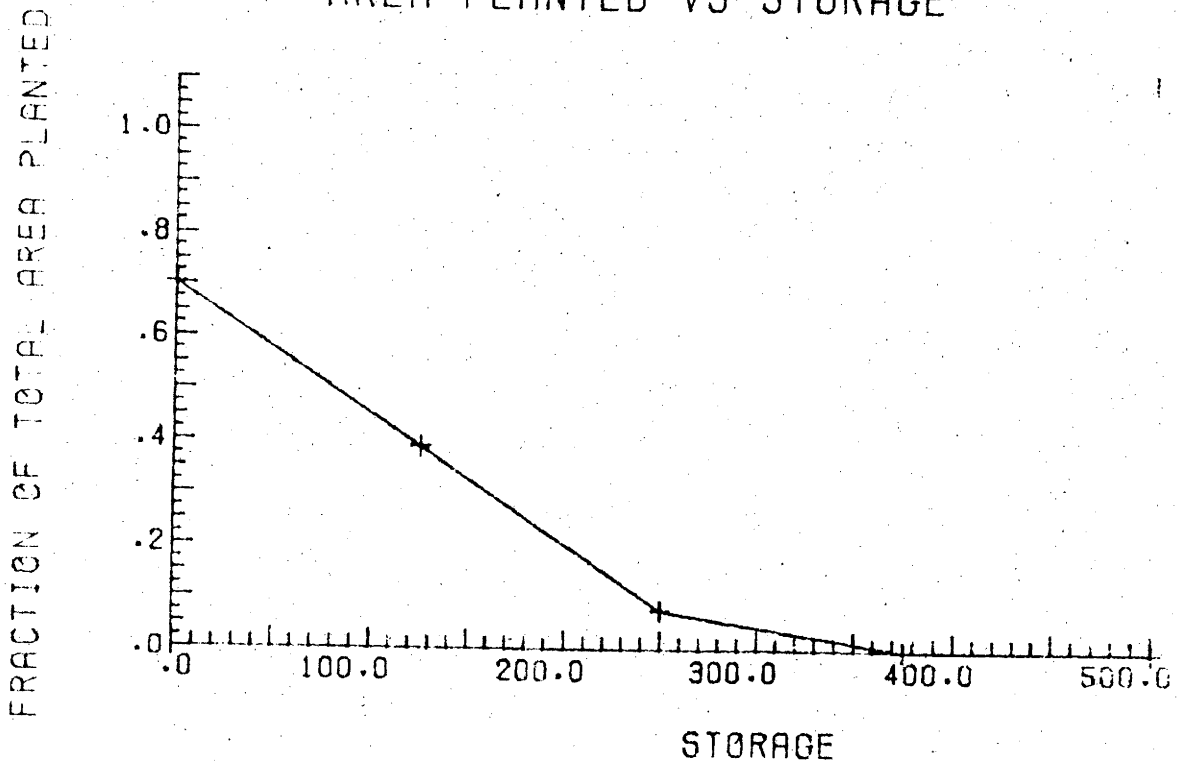

FIG. 6.4.3: DISEASE FREE AGRIBUSINESS
E(RETURN) VS STORAGE

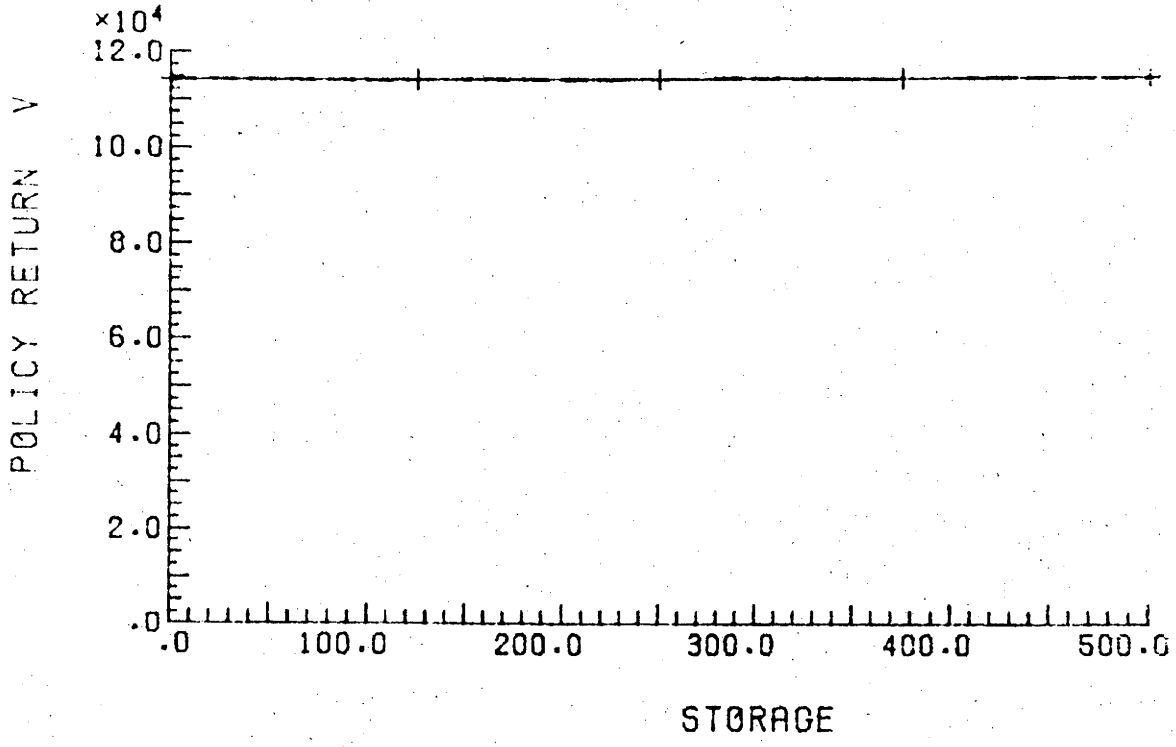

FIG. 6.4.4: DISEASE FREE AGRIBUSINESS
STORAGE INPUT/OUTPUT

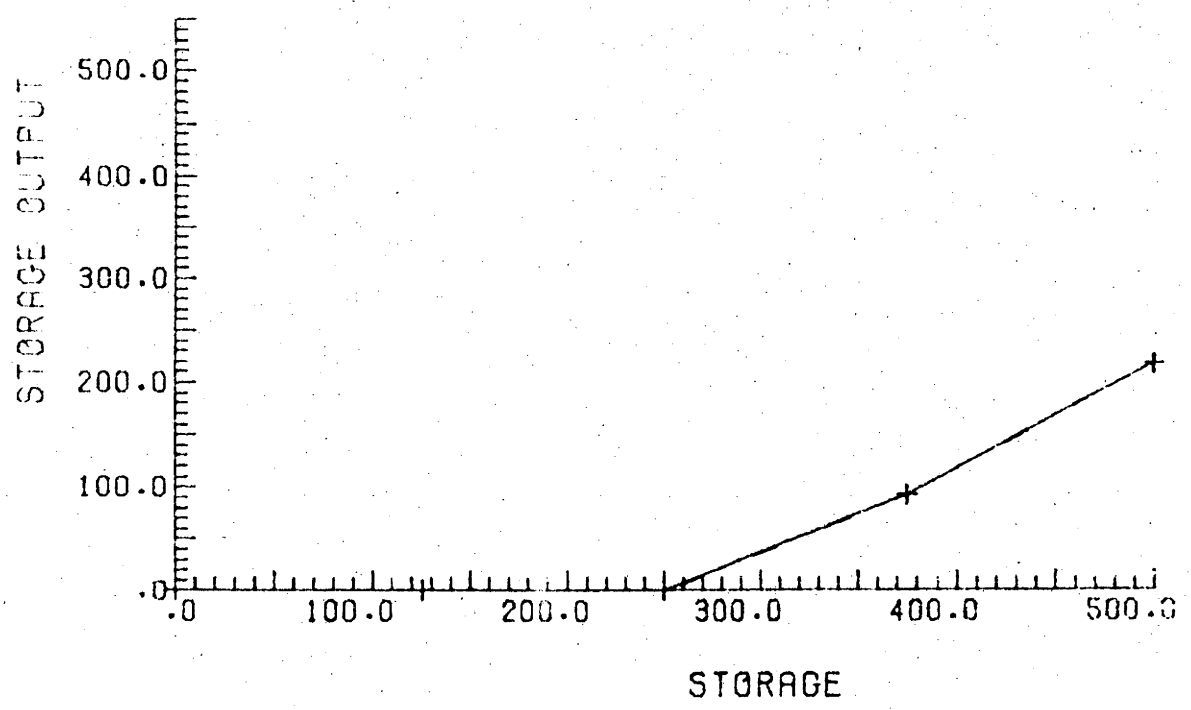

FIG. 6.4.5: DISEASE-FREE RETURN THROUGH TIME

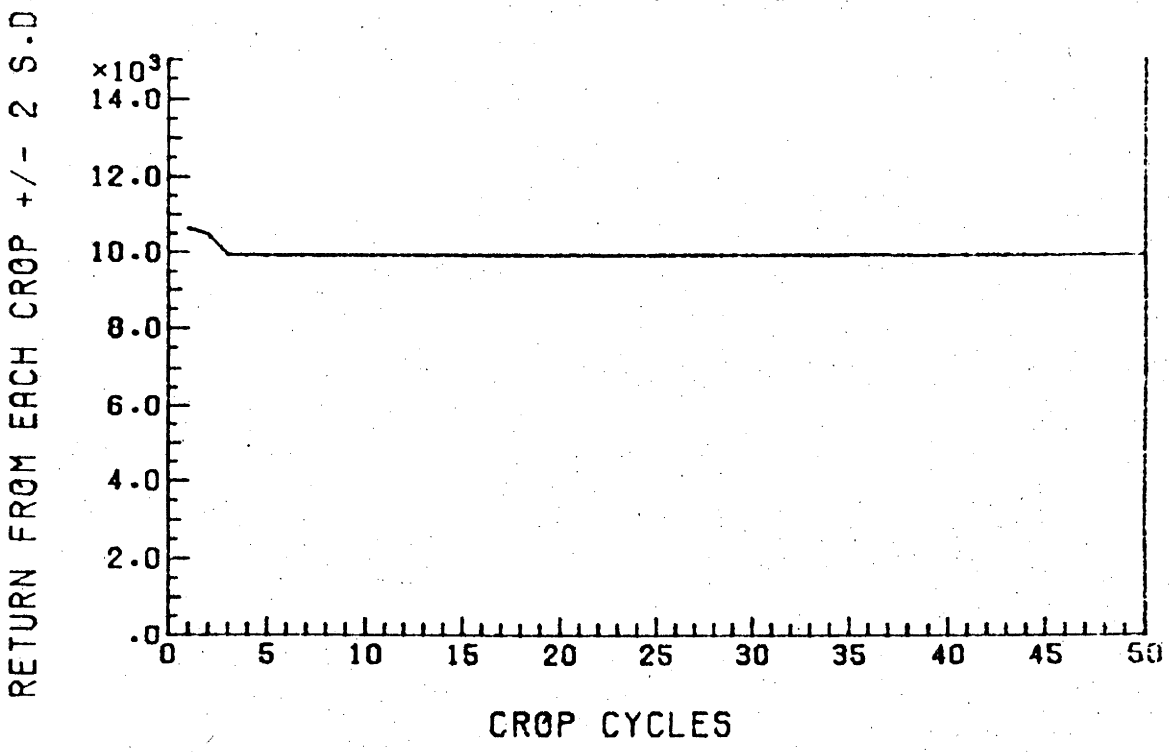

FIG. 6.4.6: DISEASE-FREE STORAGE THROUGH TIME

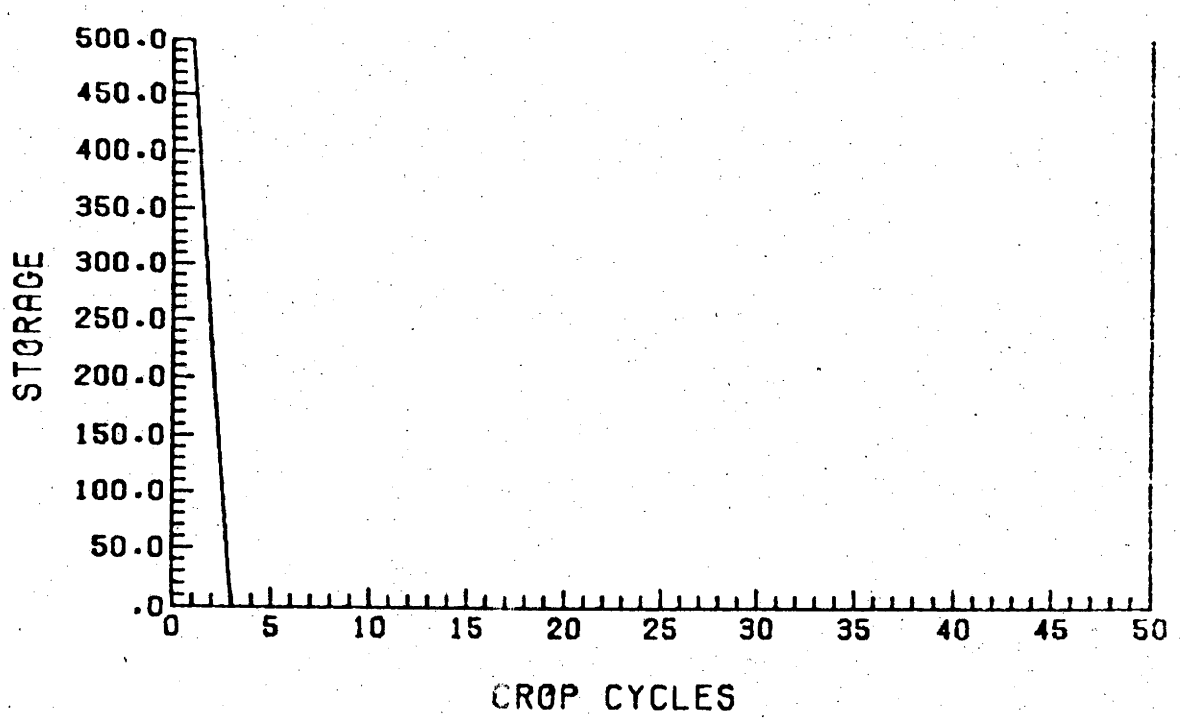

FIG. 6.4.7: DISEASE-FREE - FRACTION SOLD

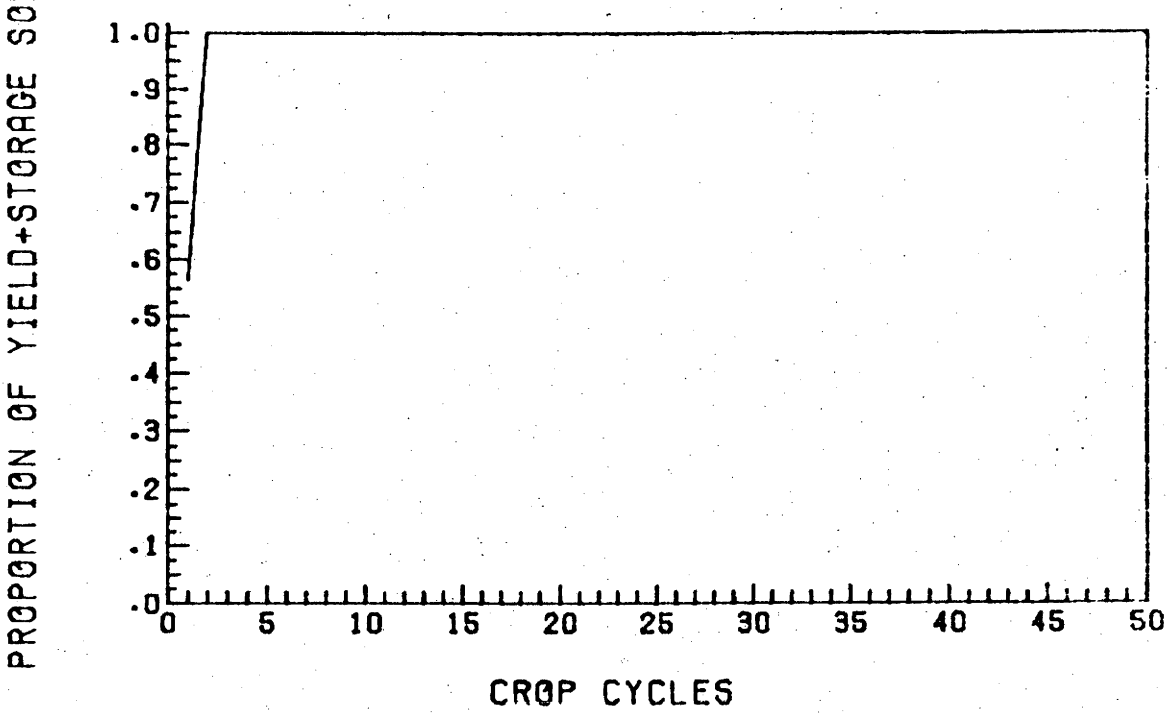

FIG. 6.4.8: DISEASE-FREE - AREA PLANTED

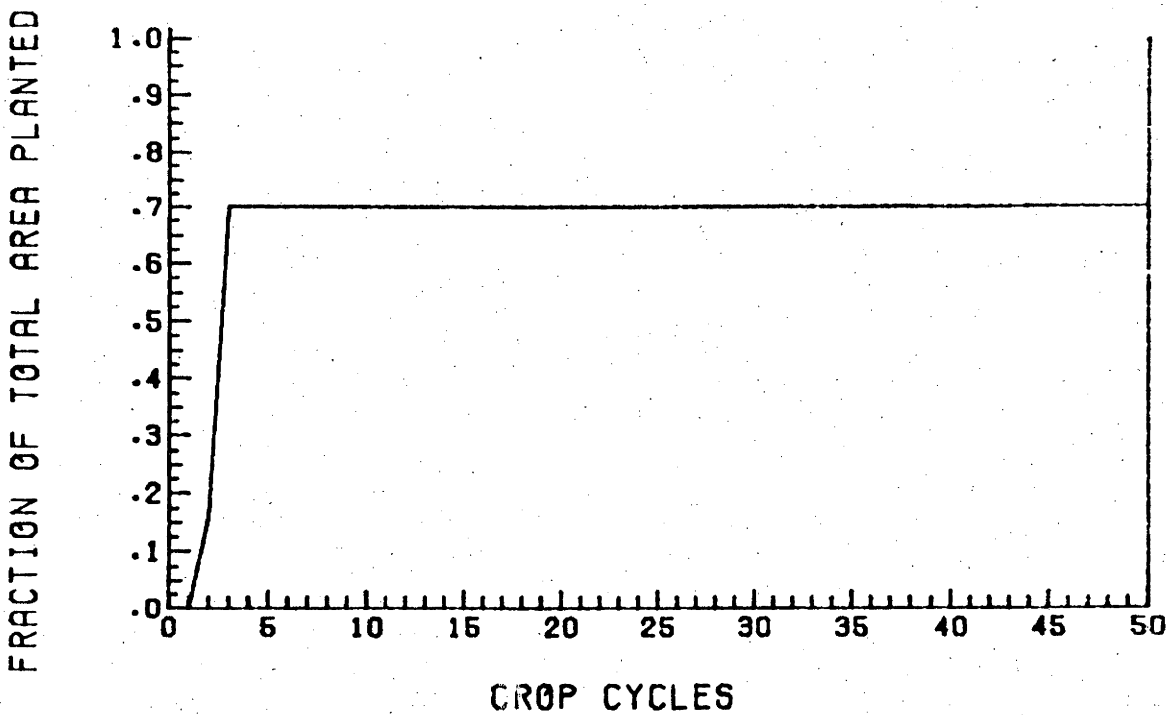

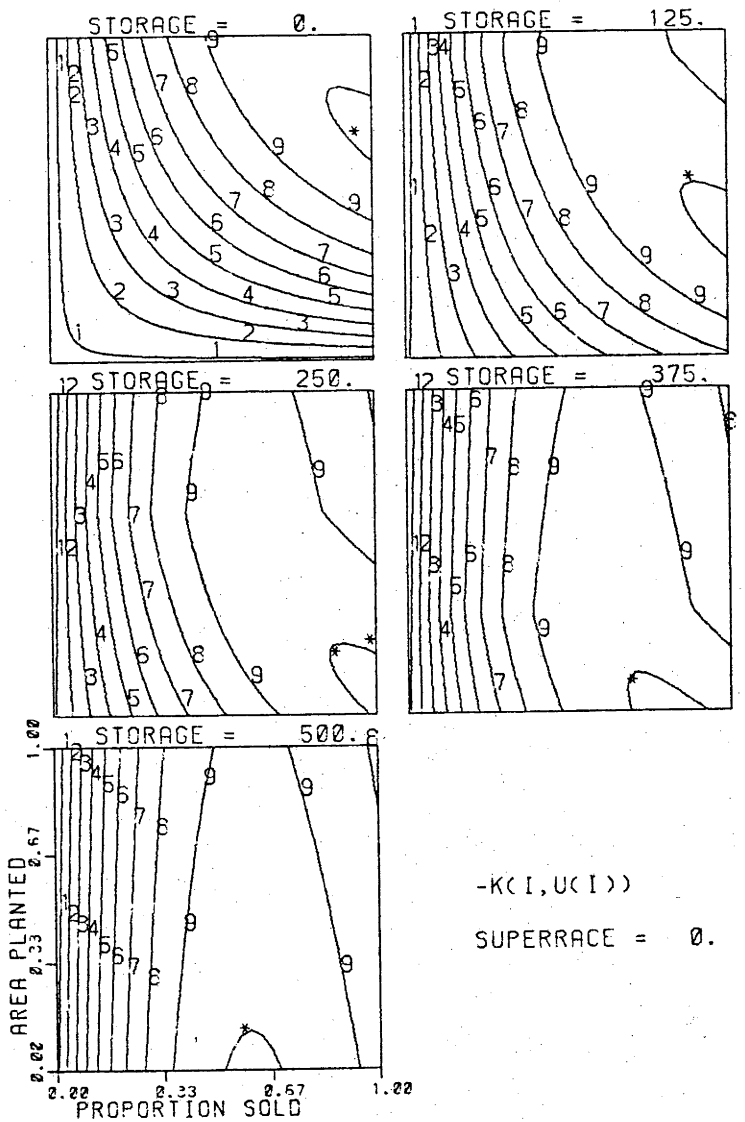

FIG. 6.4.9: DISEASE FREE RETURN

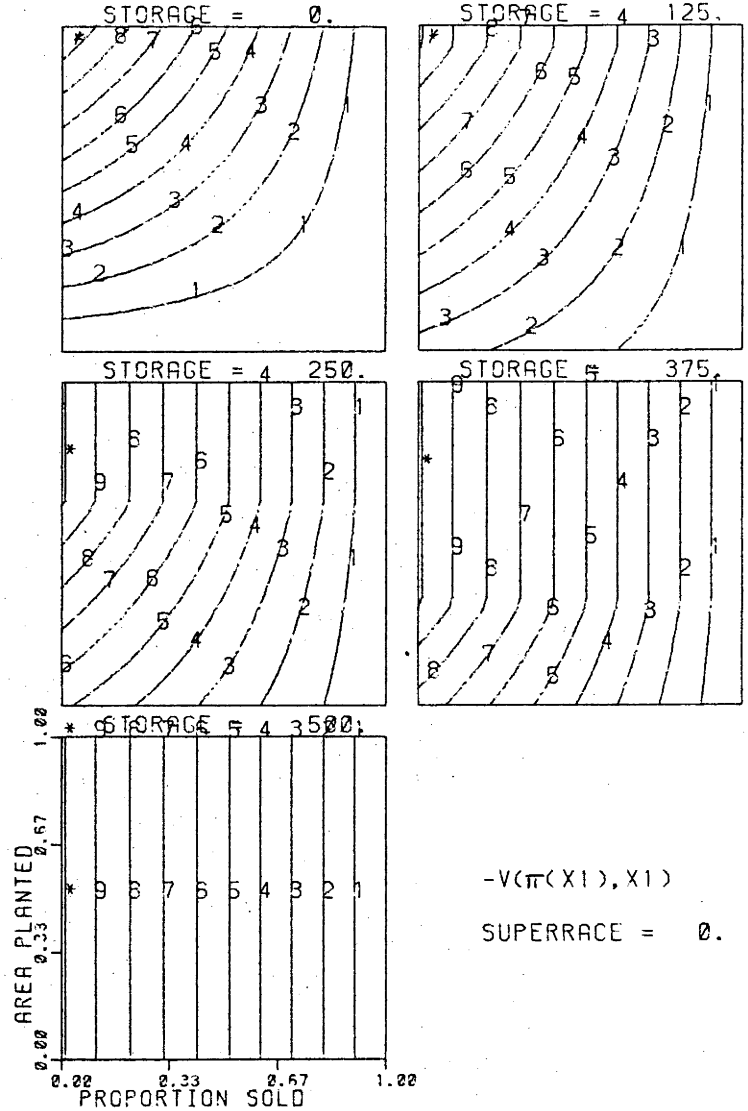

FIG. 6.4.10: DISEASE FREE RETURN

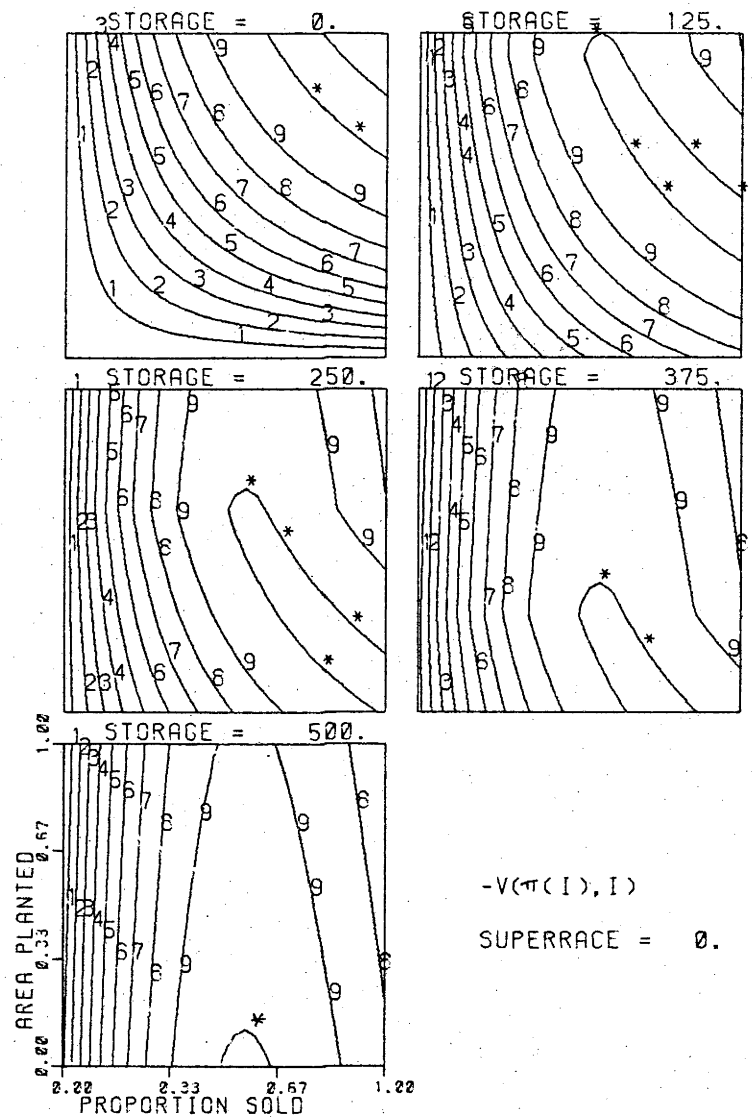

FIG. 6.4.11: DISEASE FREE RETURN

FIG. 6.4.12: BASELINE EPIDEMIC
DEPENDENCE OF YIELD ON PATHOGEN INPUT

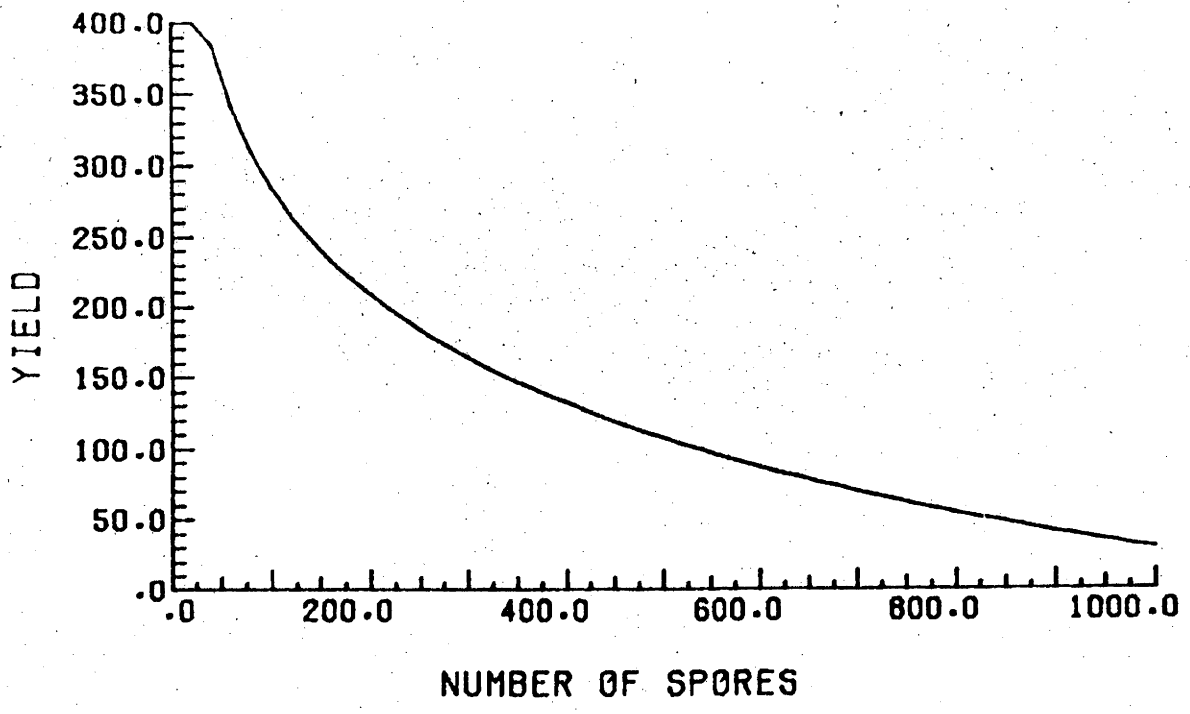

FIG. 6.4.13: BASELINE EPIDEMIC - NATURAL INCREASE OF PATHOGEN

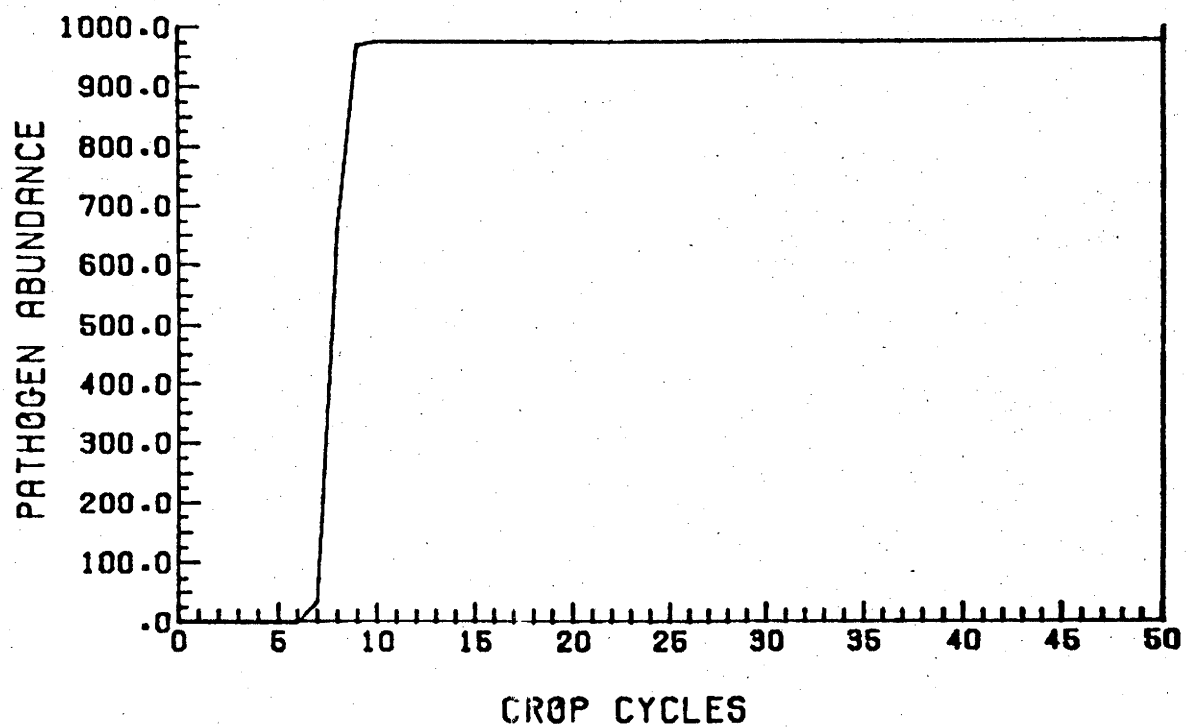

FIG. 6.4.14(A): AGRIBUSINESS BASELINE
FRACTION SOLD VS STORAGE

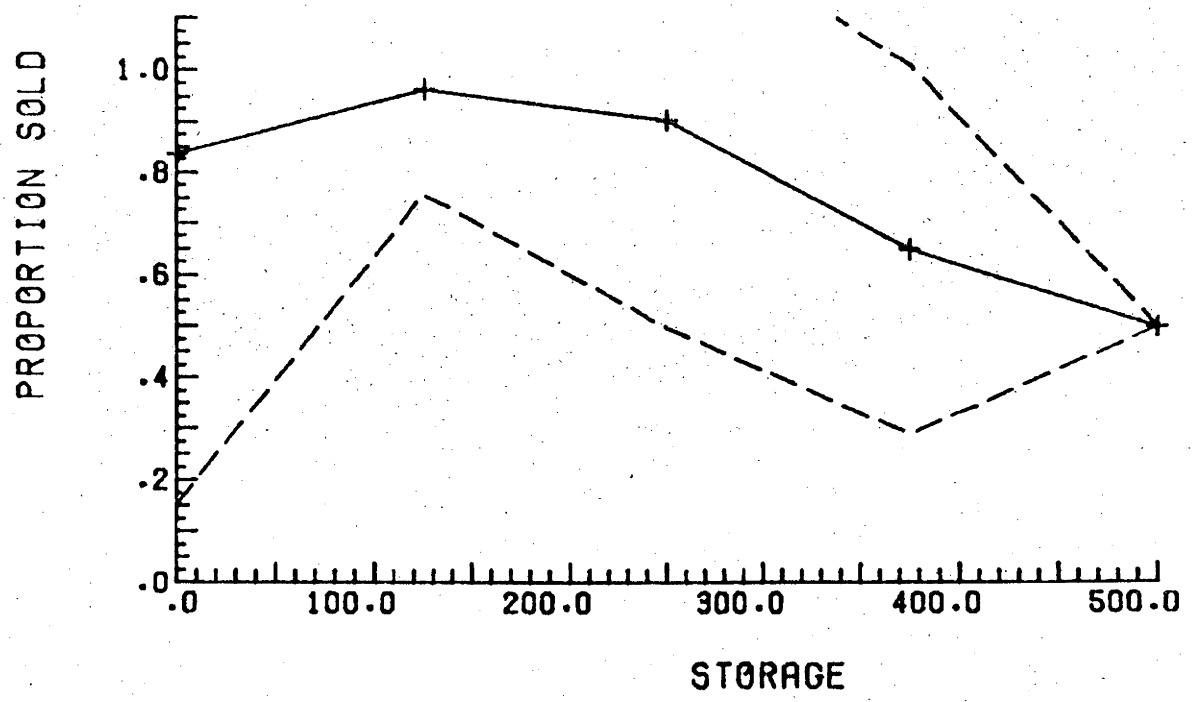

FIG 6.4.14B: AGRIBUSINESS BASELINE
FRACTION SOLD VS SUPERRACE

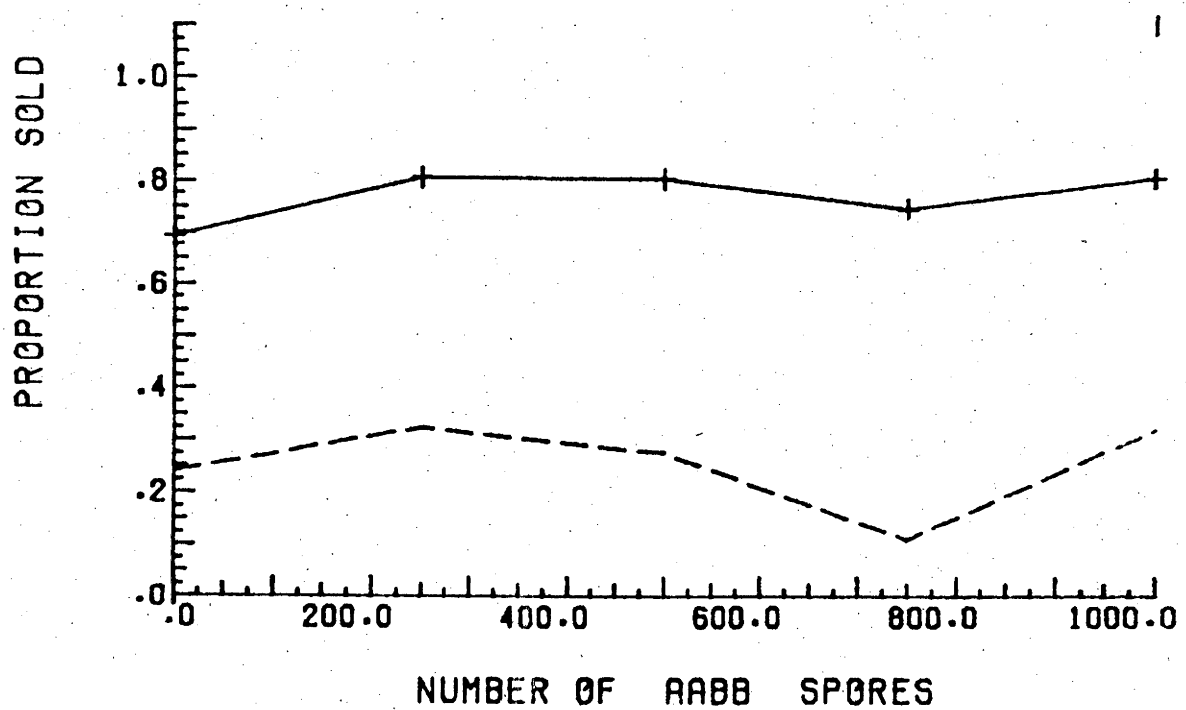

FIG 6.4.15A: AGRIBUSINESS BASELINE
AREA PLANTED VS STORAGE

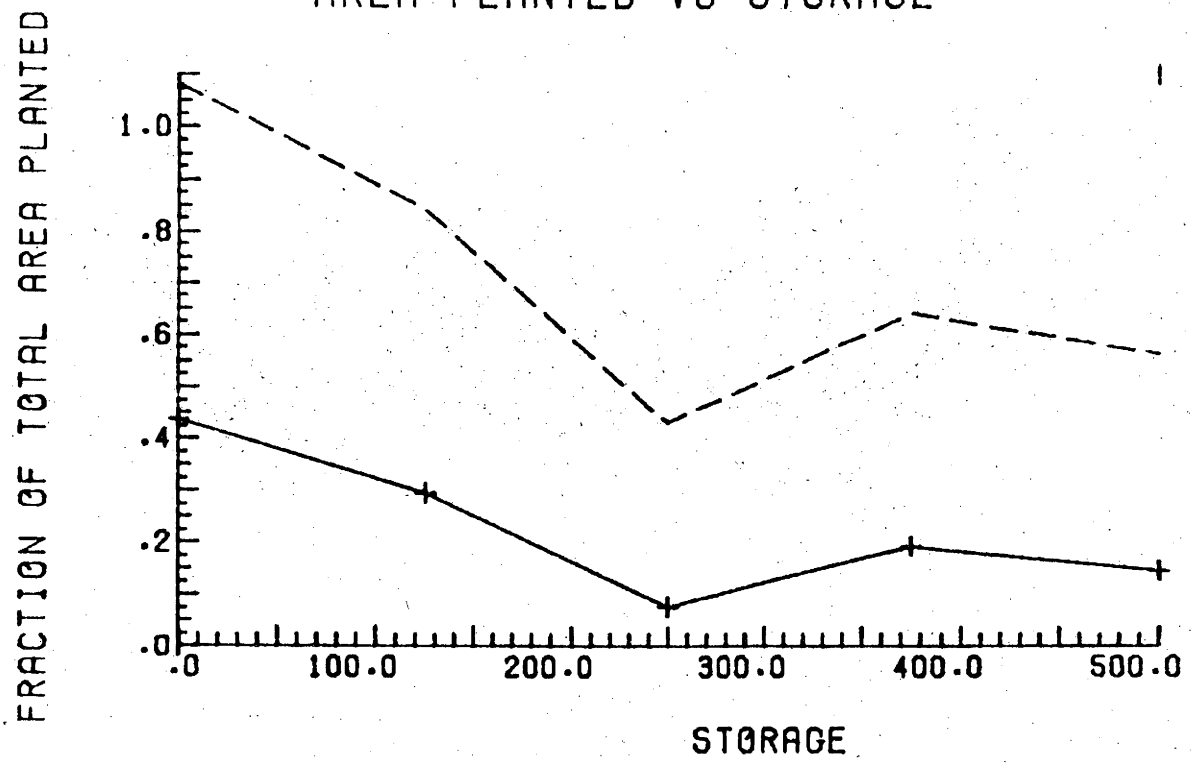

FIG 6.4.15B: AGRIBUSINESS BASELINE
AREA PLANTED VS SUPERRACE

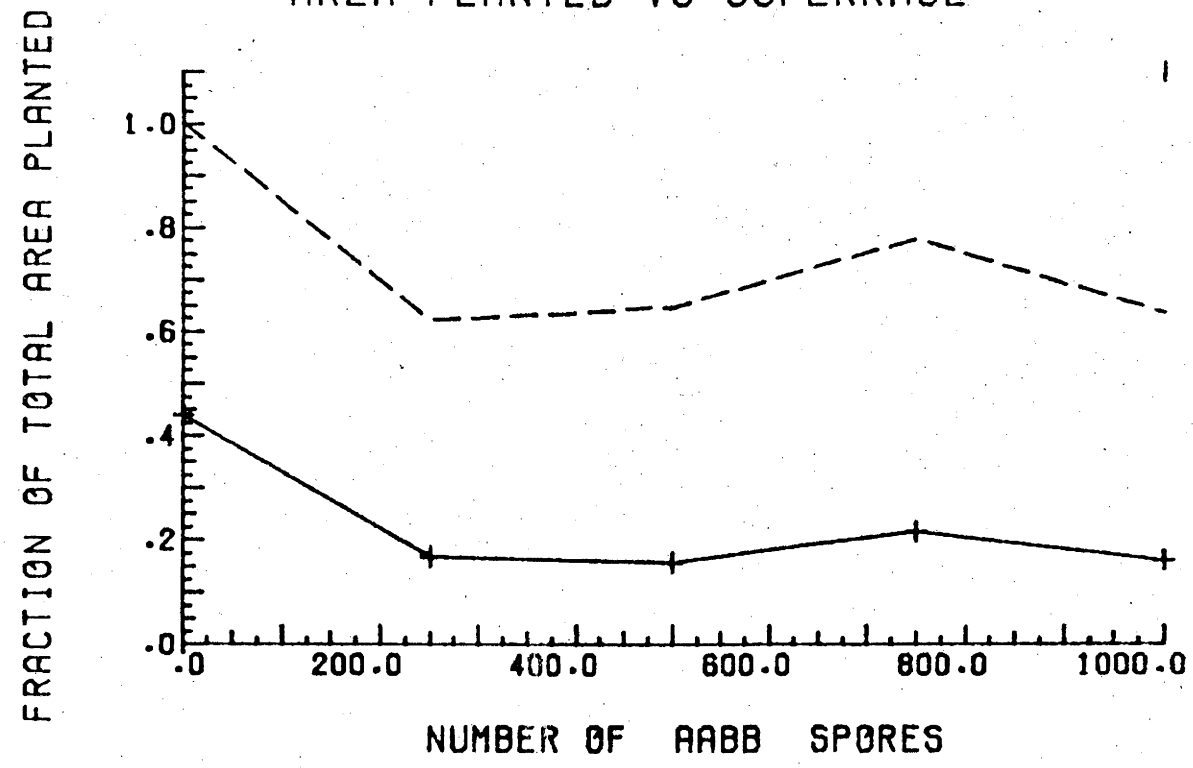

FIG 6.4.16A: AGRIBUSINESS BASELINE
E(RETURN) VS STORAGE

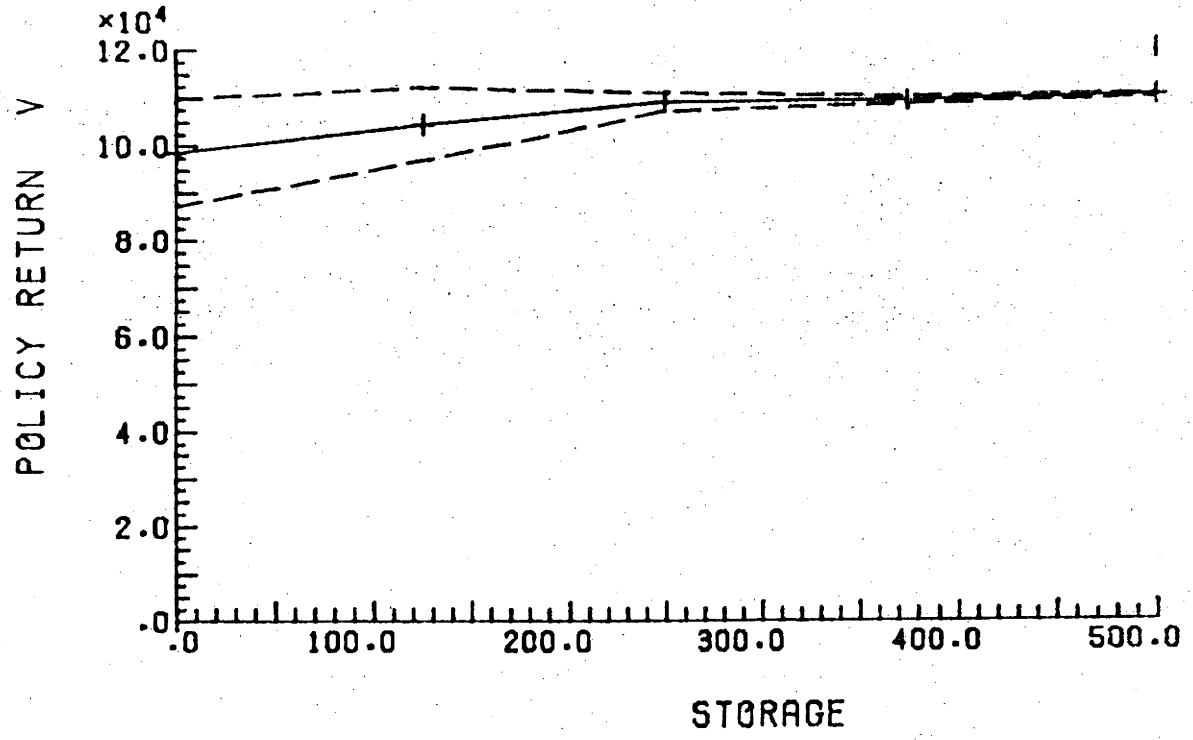

FIG 6.4.16B: AGRIBUSINESS BASELINE
E(RETURN) VS SUPERRACE

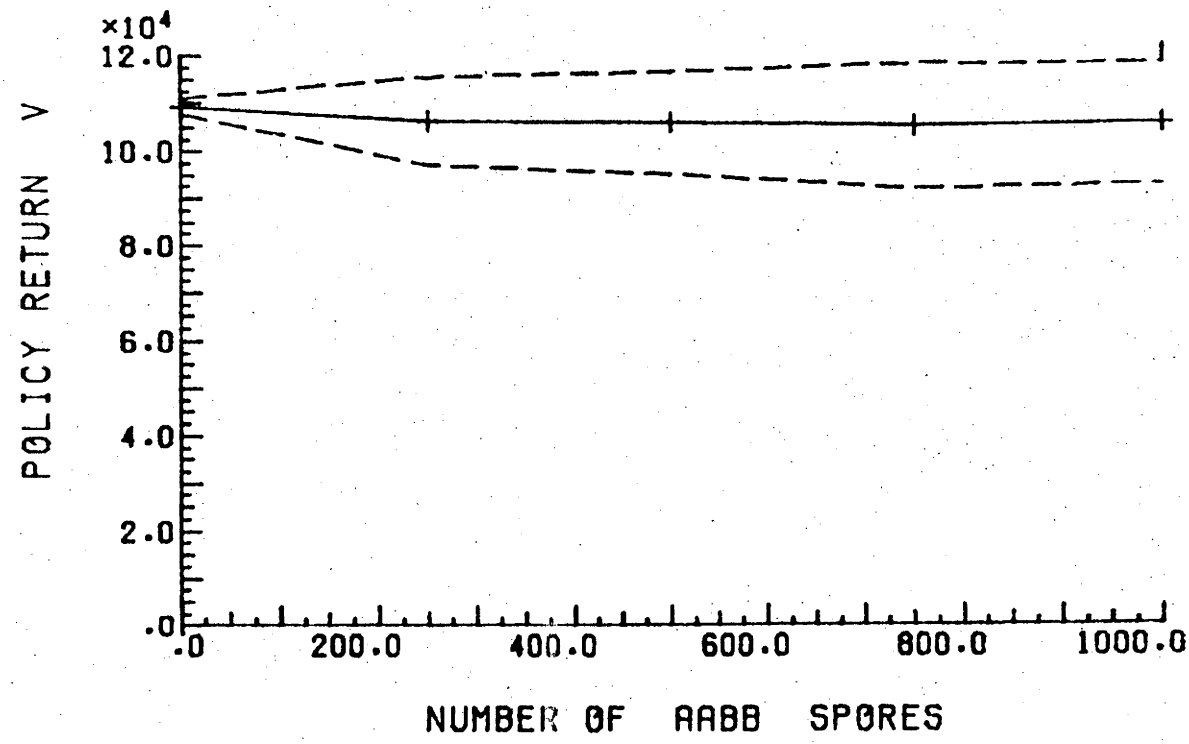

FIG 6.4.17A: AGRIBUSINESS BASELINE
STORAGE LEFT VS STORAGE

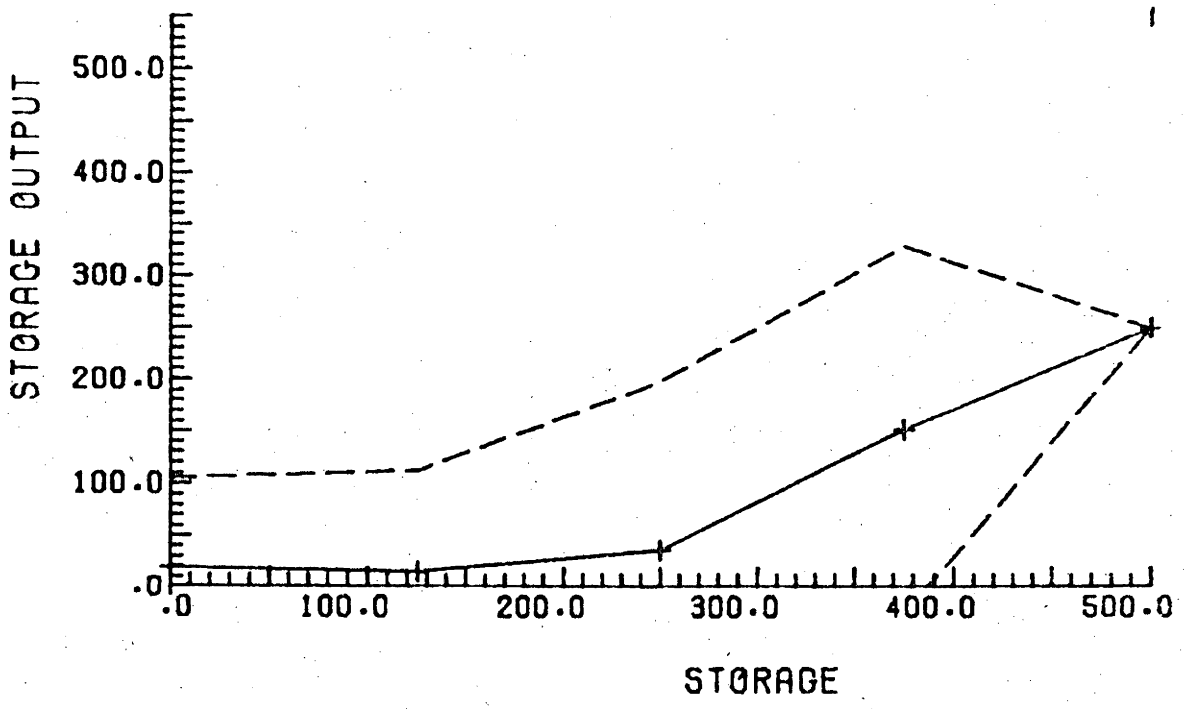

FIG 6.4.17B: AGRIBUSINESS BASELINE
STORAGE LEFT VS SUPERRACE

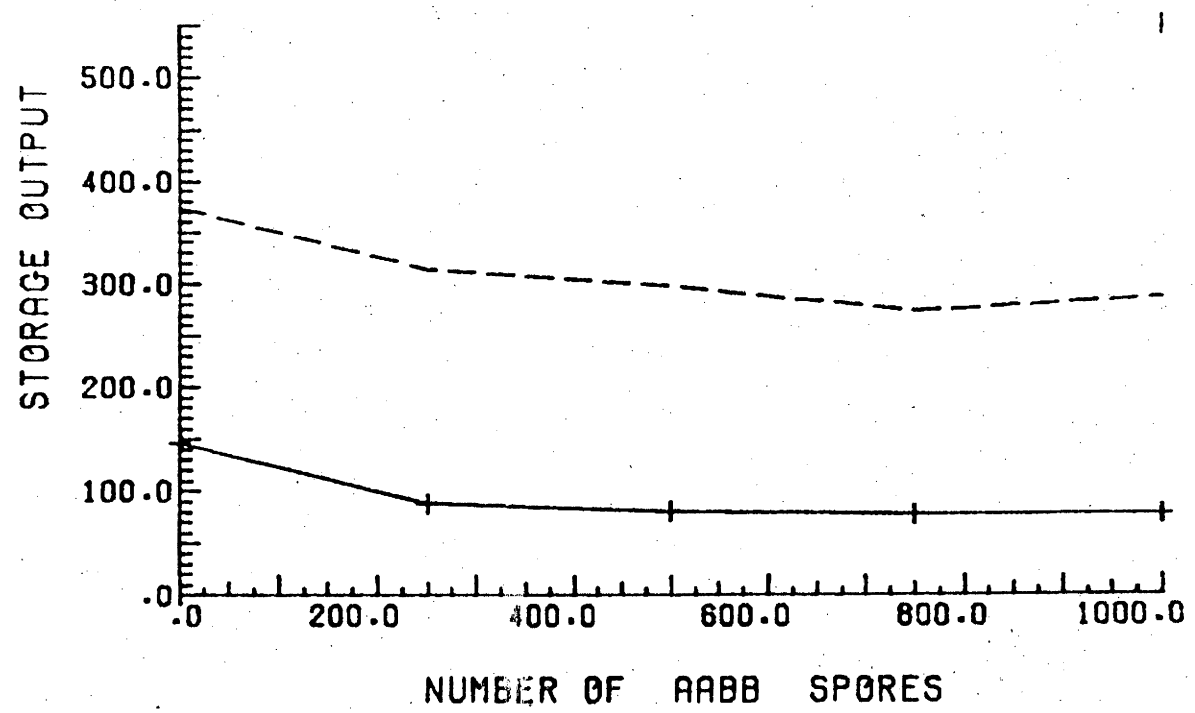

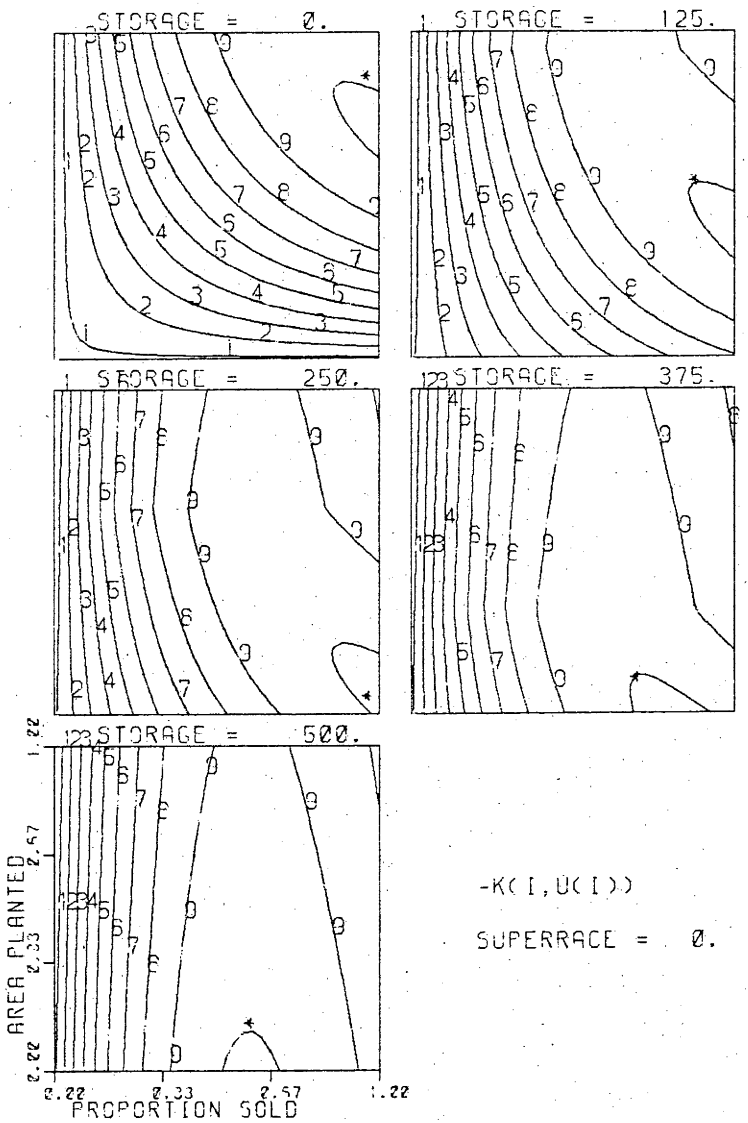

FIG. 5.4.19a: AGRIBUSINESS BASELINE RETURN

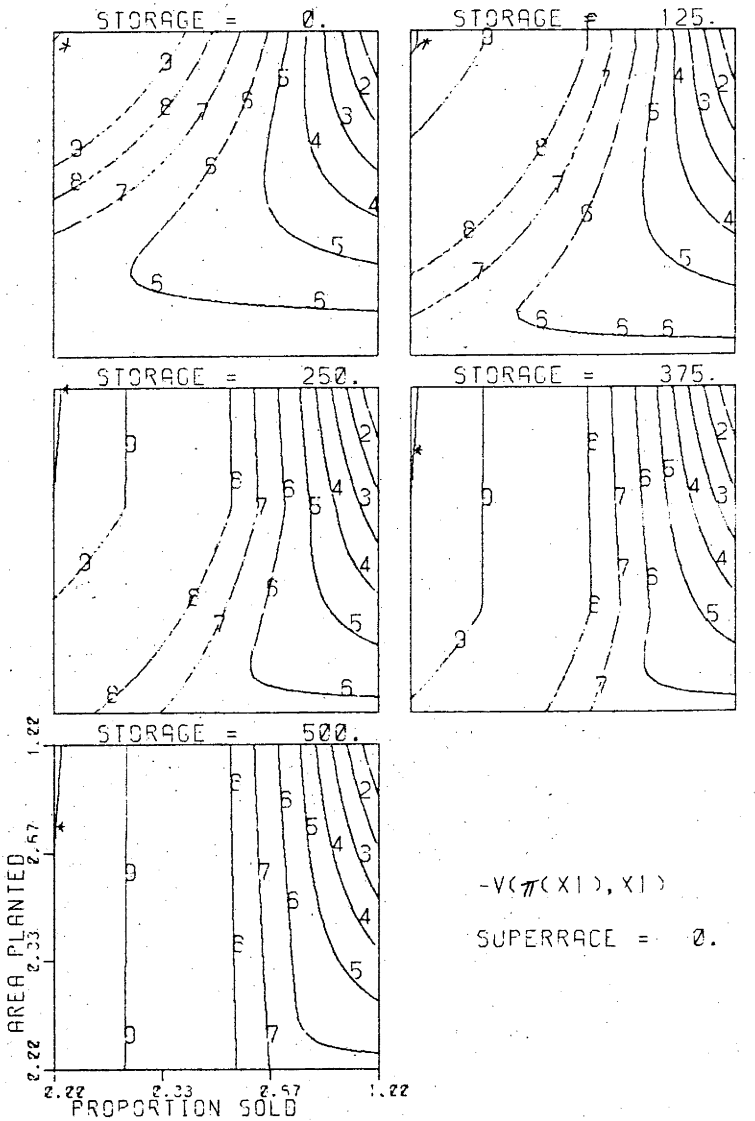

FIG. 5.4.19b: AGRIBUSINESS BASELINE RETURN

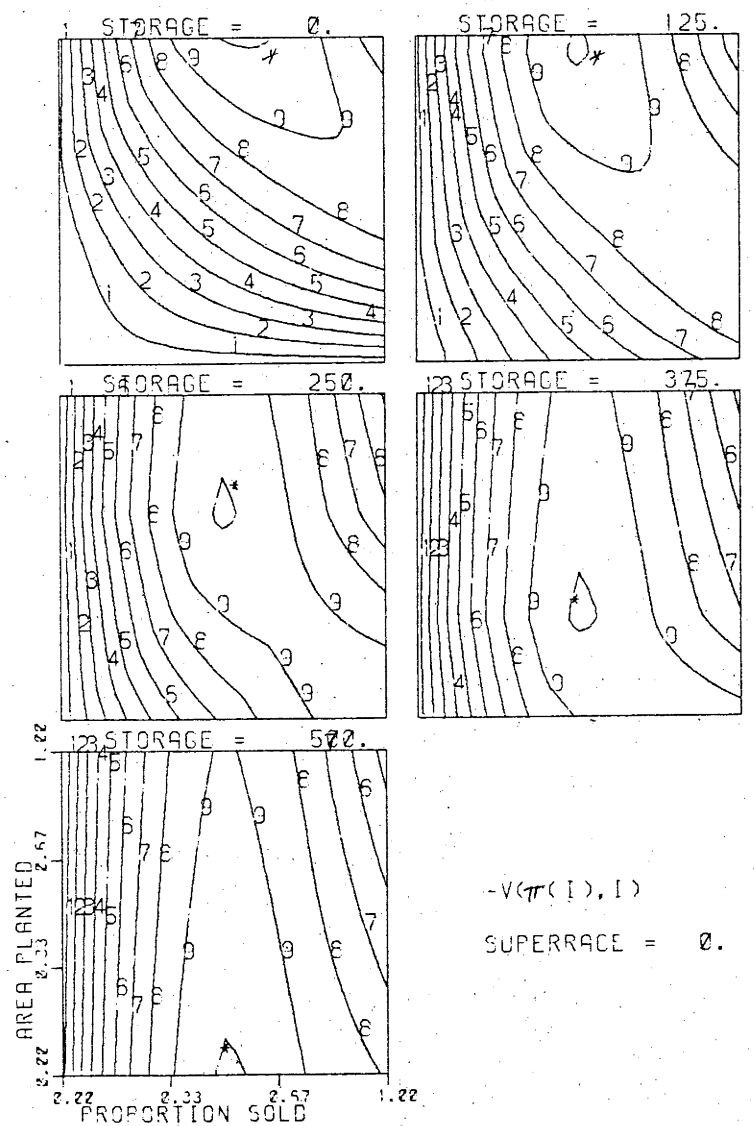

FIG. 6.4.19C: AGRIBUSINESS BASELINE RETURN

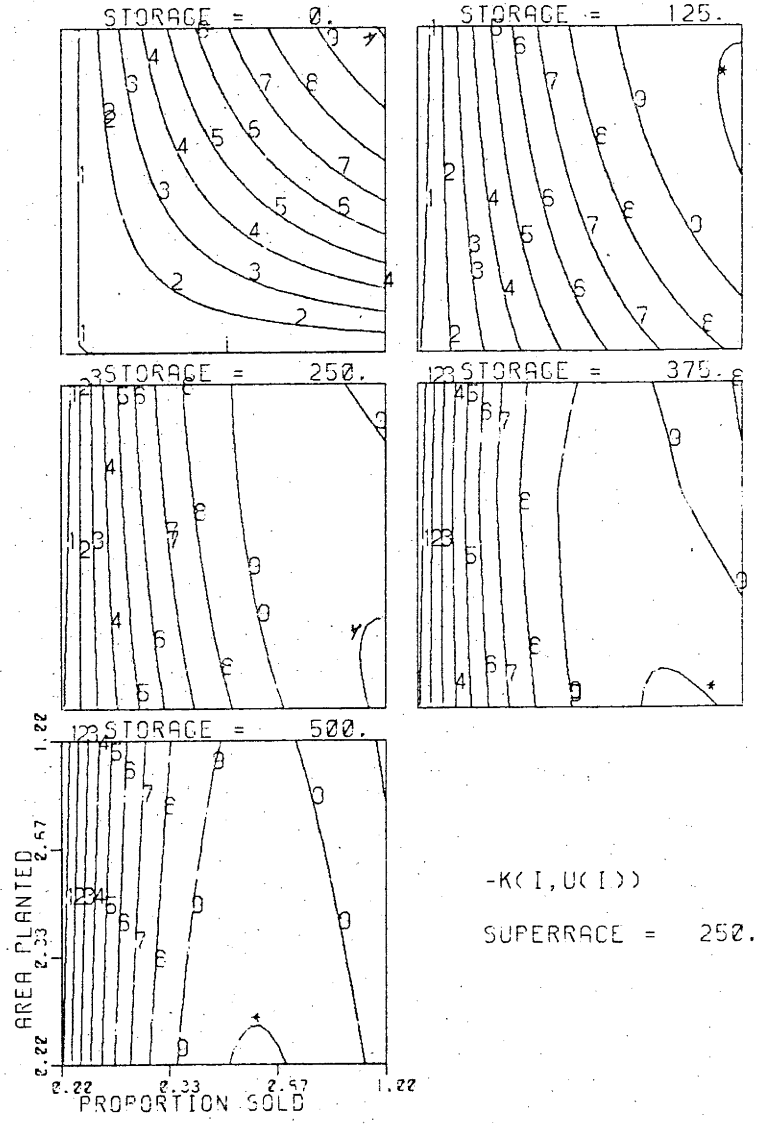

FIG. 6.4.20A: AGRIBUSINESS BASELINE RETURN

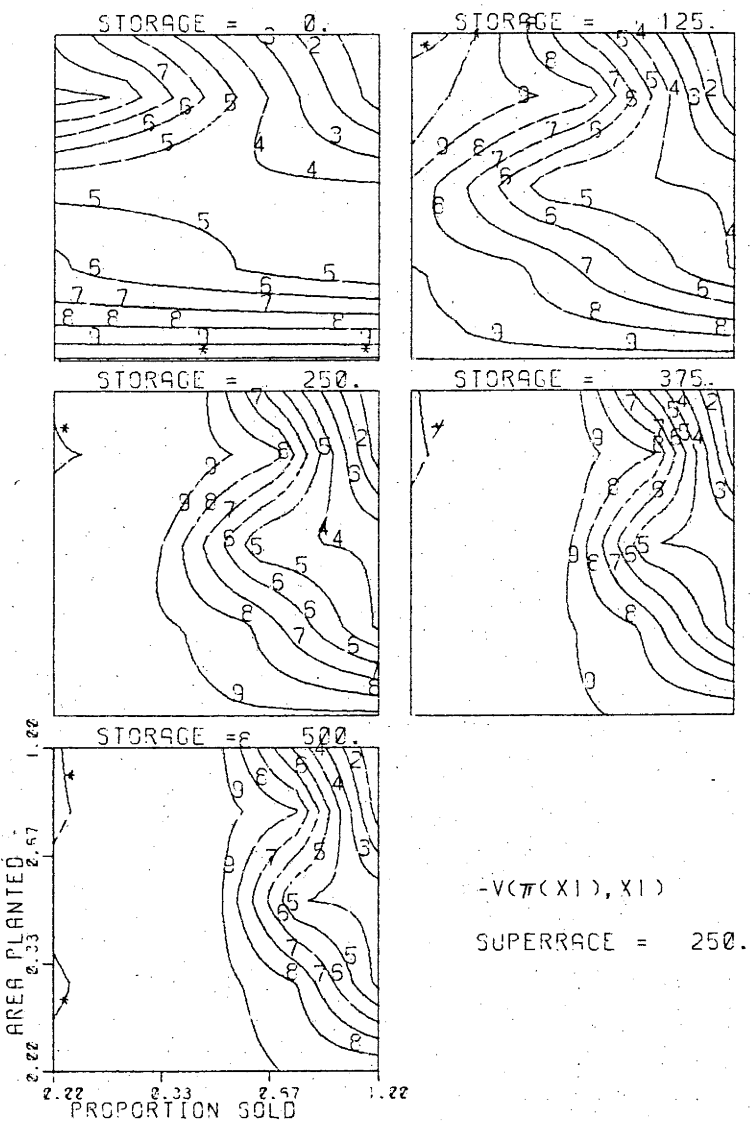

FIG. 6 4.20B: AGRIBUSINESS BASELINE RETURN

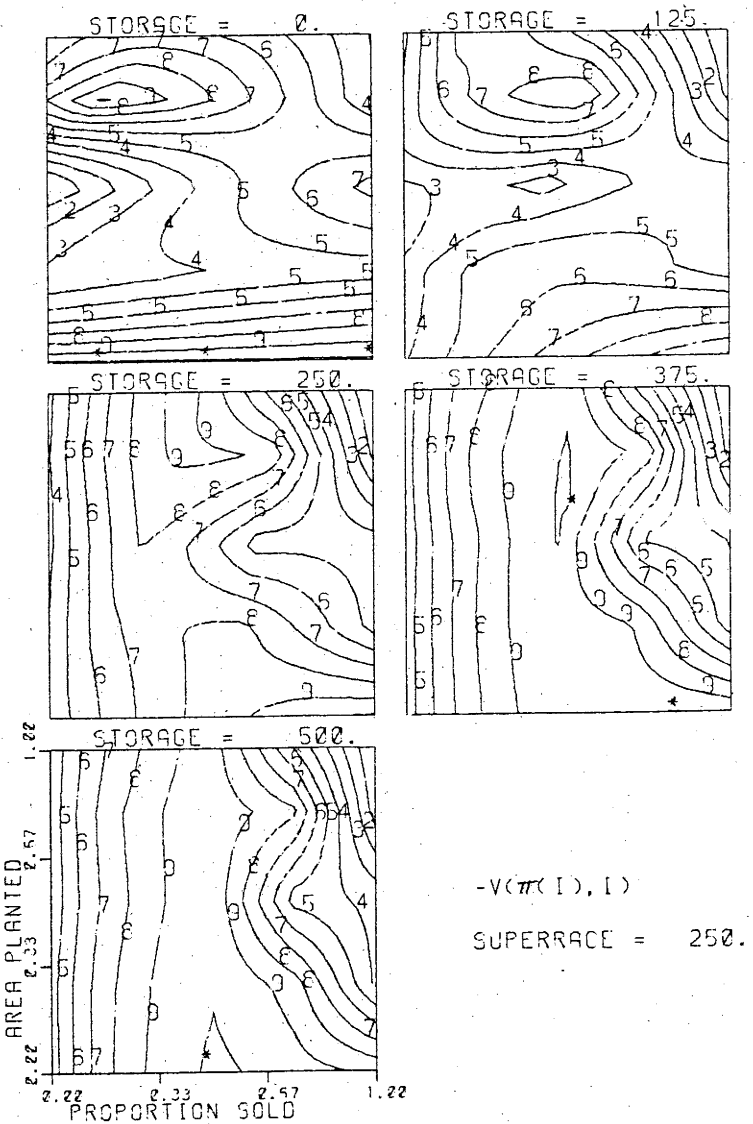

FIG. 6 4.20C: AGRIBUSINESS BASELINE RETURN

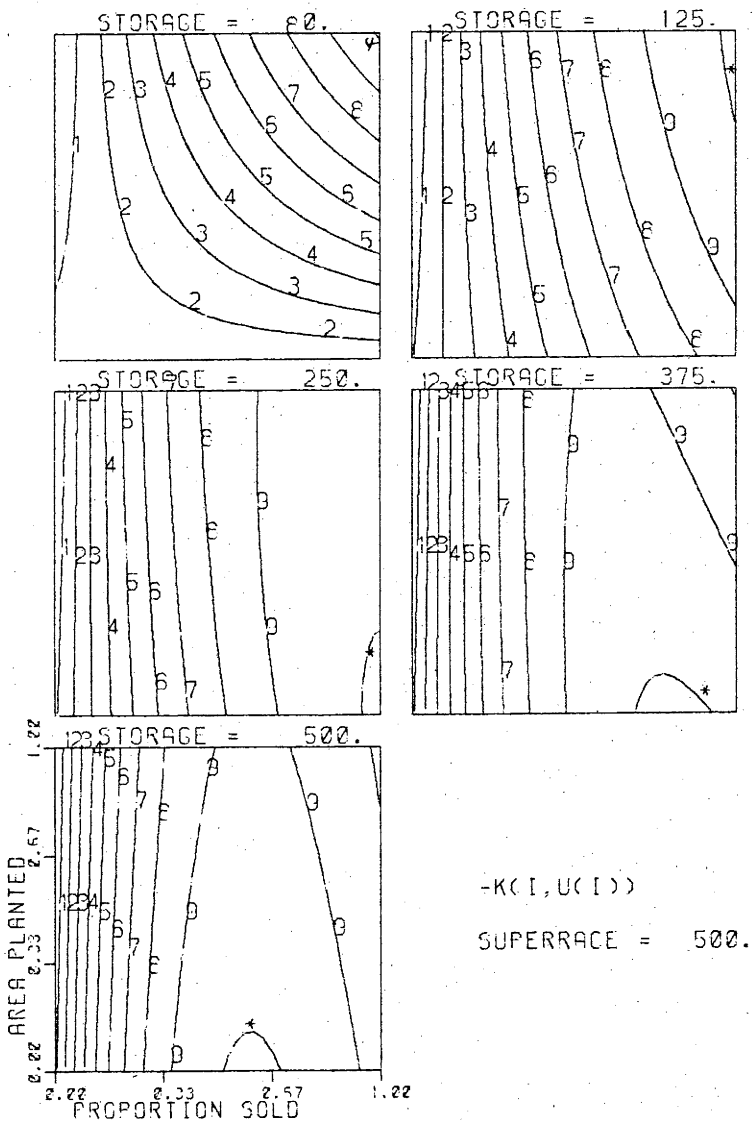

FIG. 5.4.21A: AGRIBUSINESS BASELINE RETURN

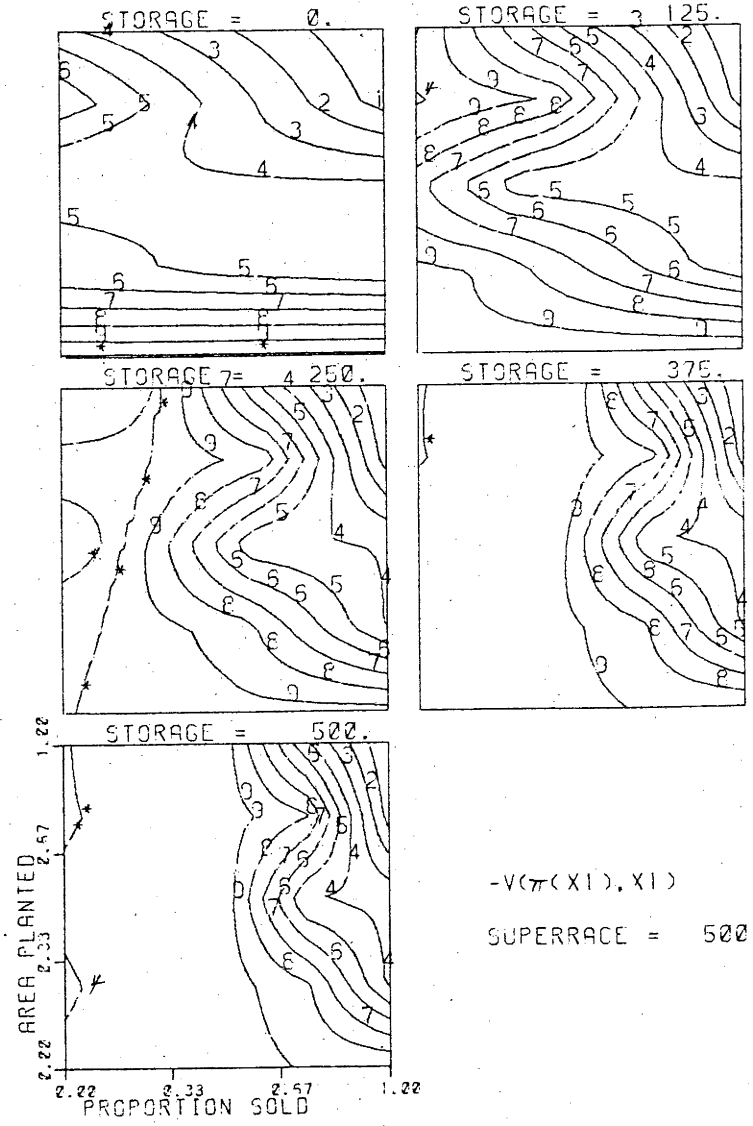

FIG. 5.4.21B: AGRIBUSINESS BASELINE RETURN

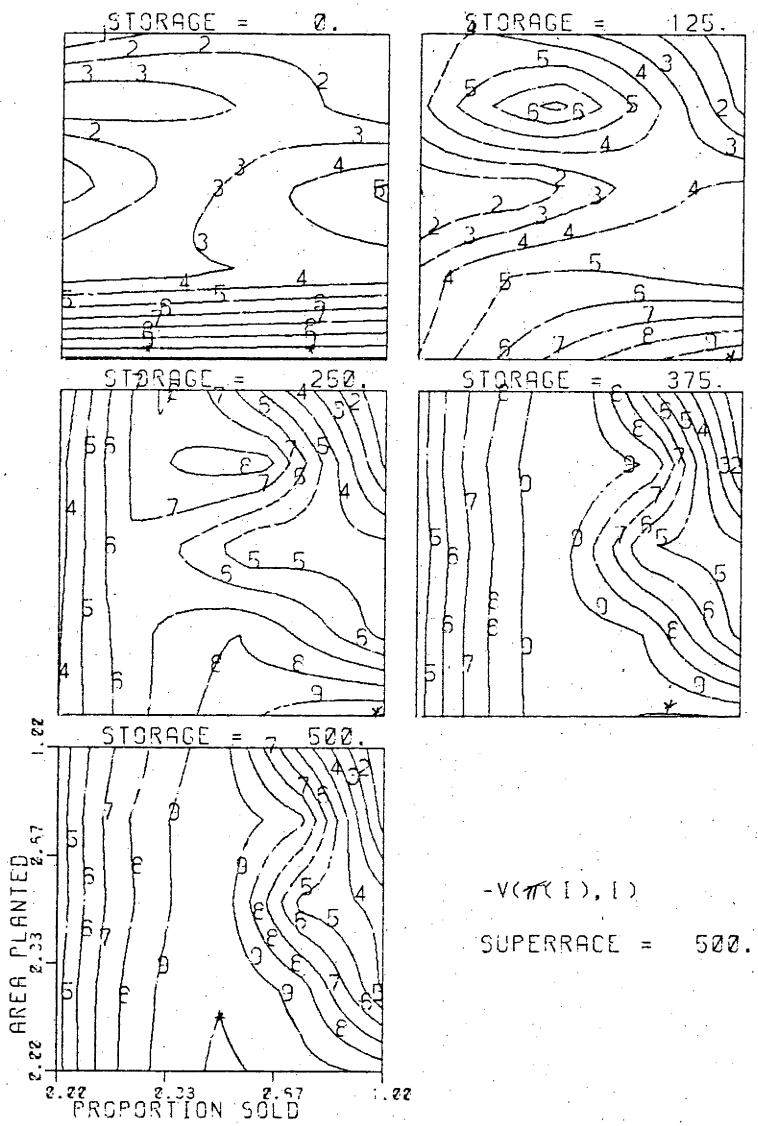

FIG. 6.4.21C: AGRIBUSINESS BASELINE RETURN

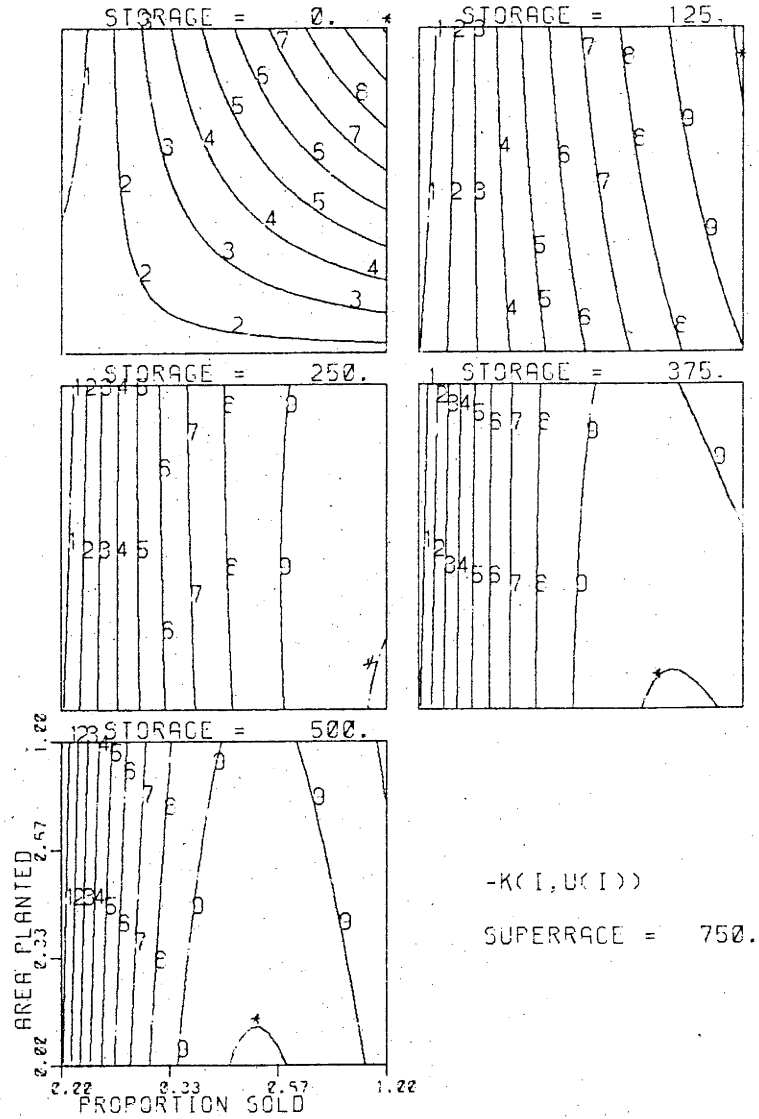

FIG. 6.4.22A: AGRIBUSINESS BASELINE RETURN

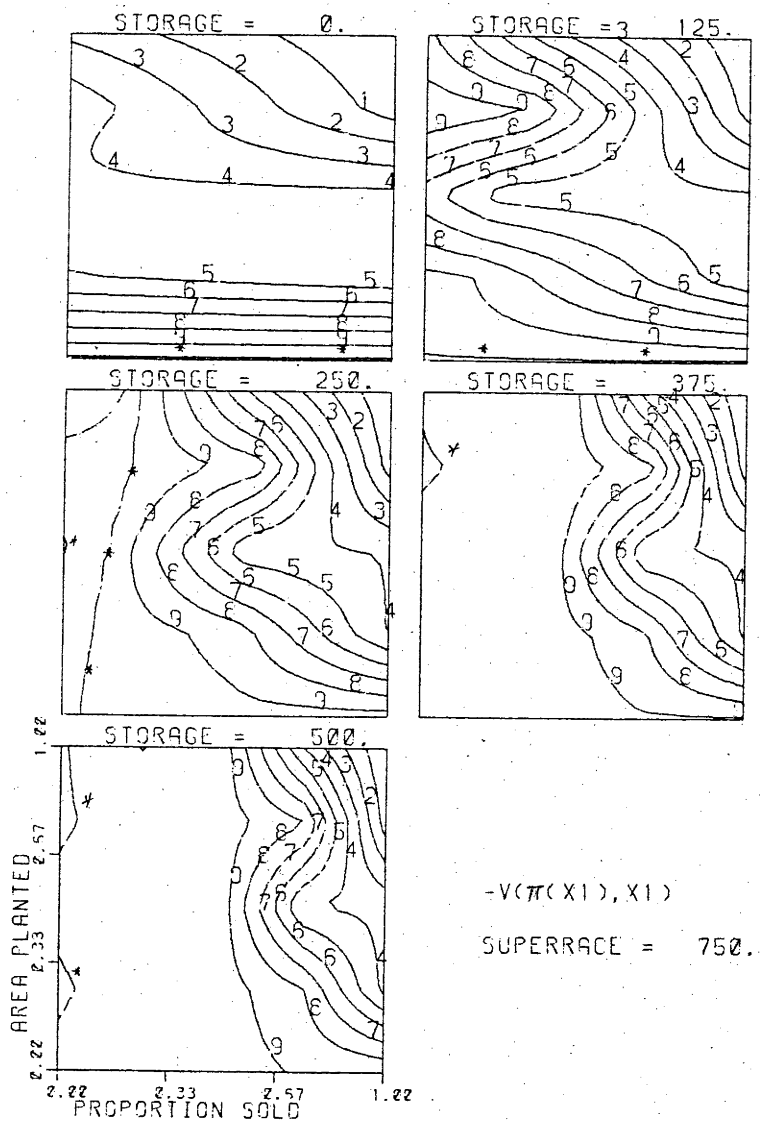

FIG. 6.4.22B: AGRIBUSINESS BASELINE RETURN

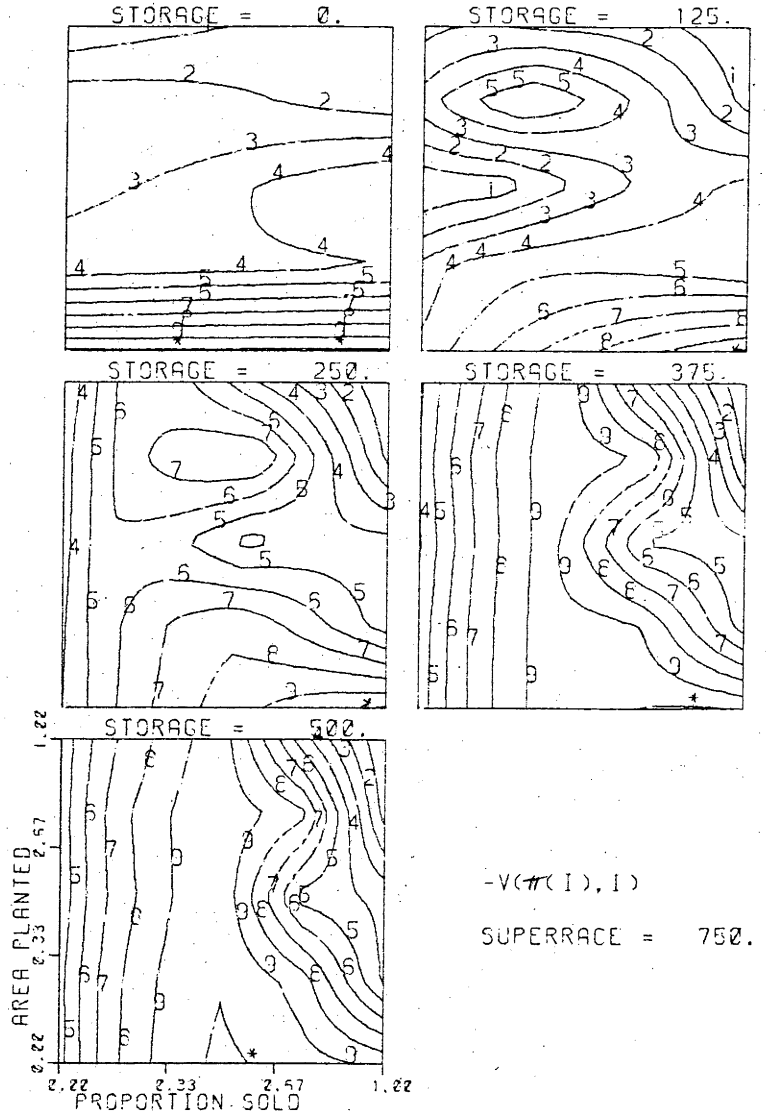

FIG. 6.4.22C: AGRIBUSINESS BASELINE RETURN

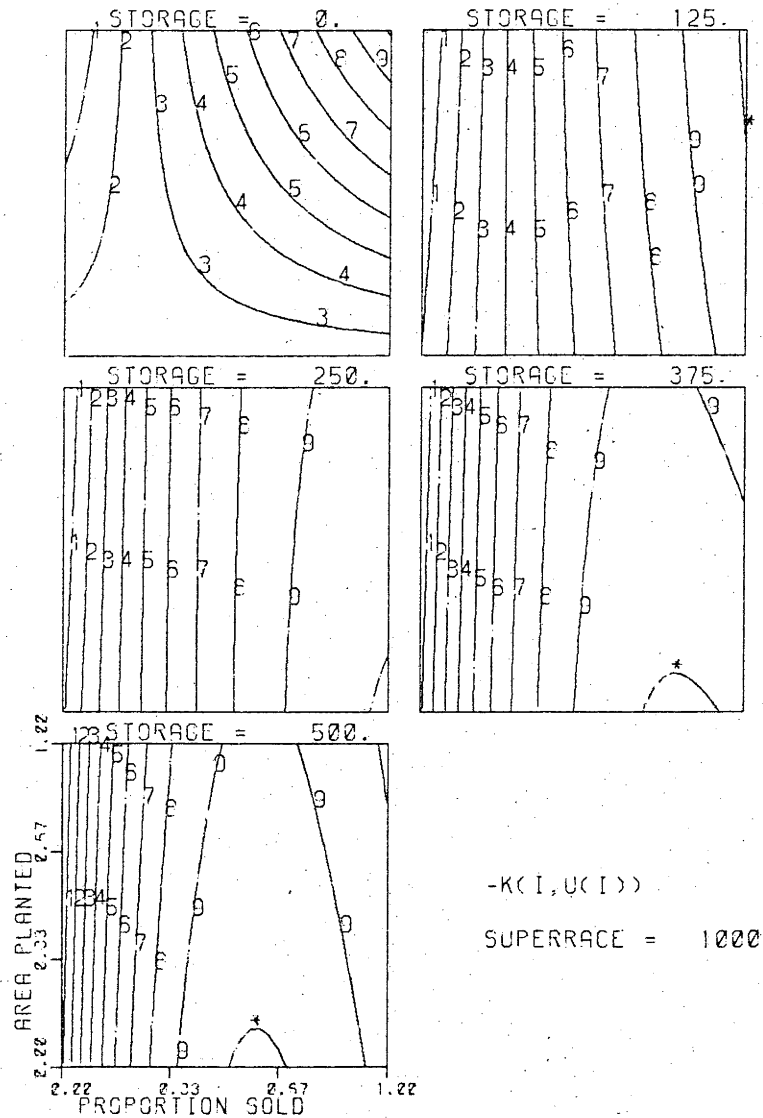

FIG. 6.4.23A: AGRIBUSINESS BASELINE RETURN

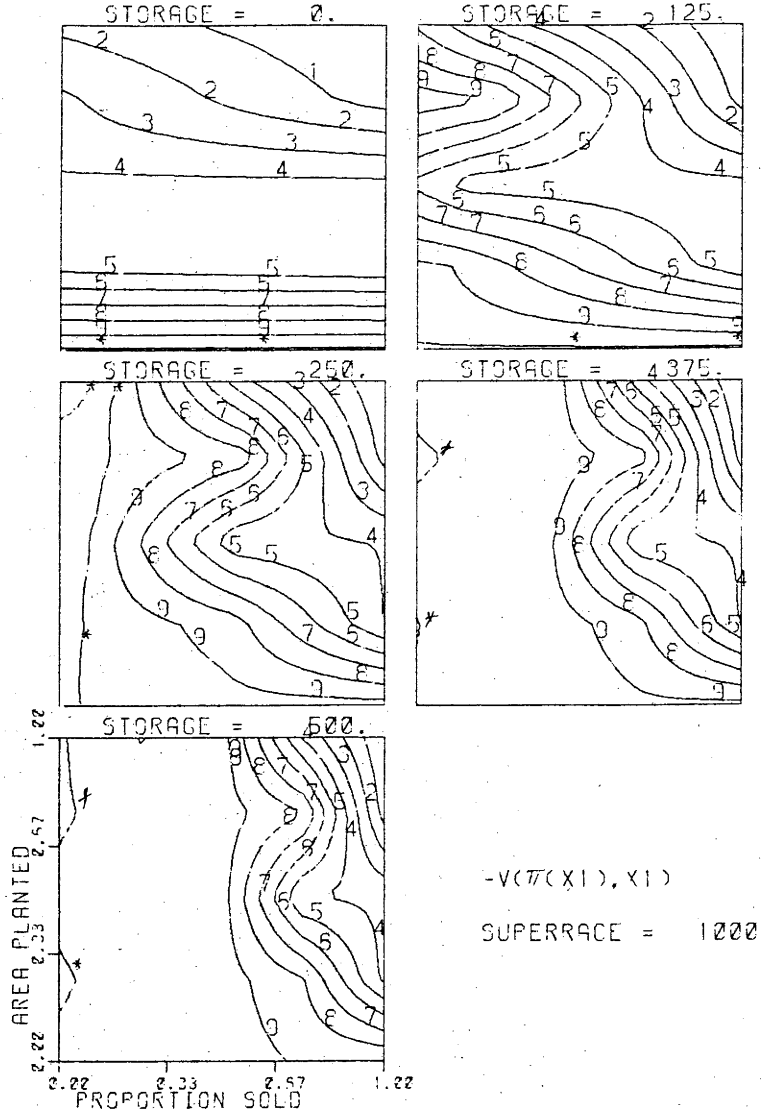

FIG. 6.4.23B: AGRIBUSINESS BASELINE RETURN

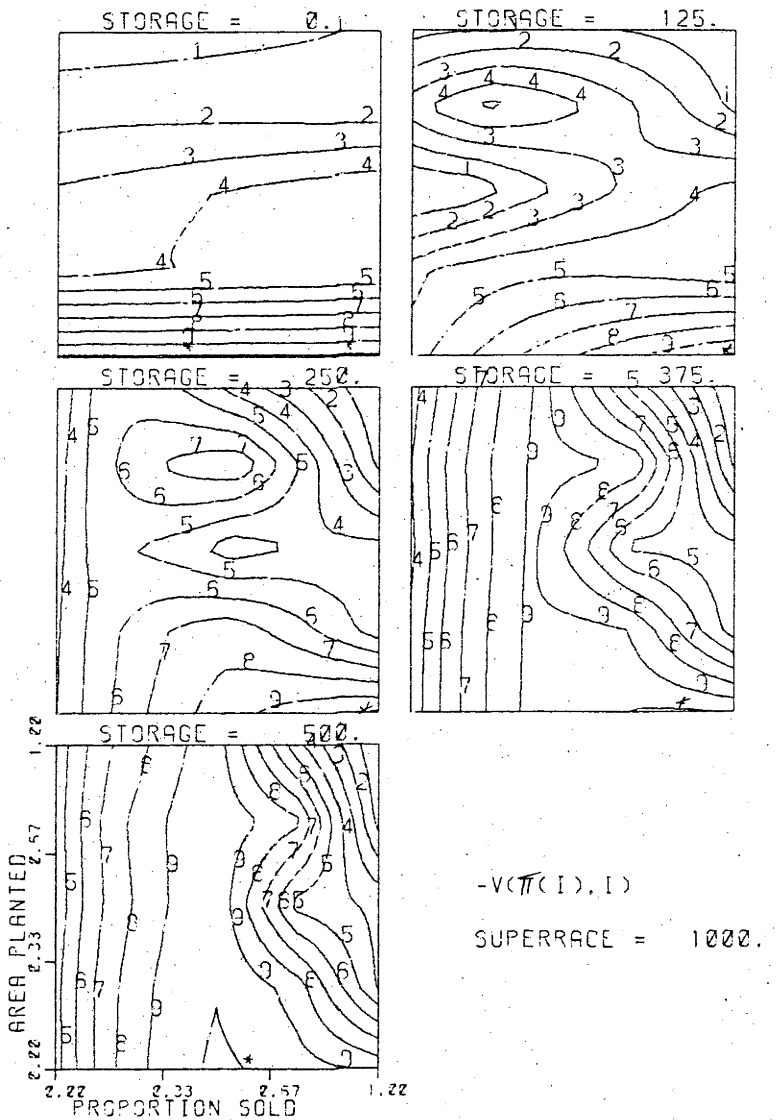

FIG. 6 4.23C: AGRIBUSINESS BASELINE RETURN

FIG 6.4.24: AGRIBUSINESS BASELINE
FRACTION SOLD THROUGH TIME

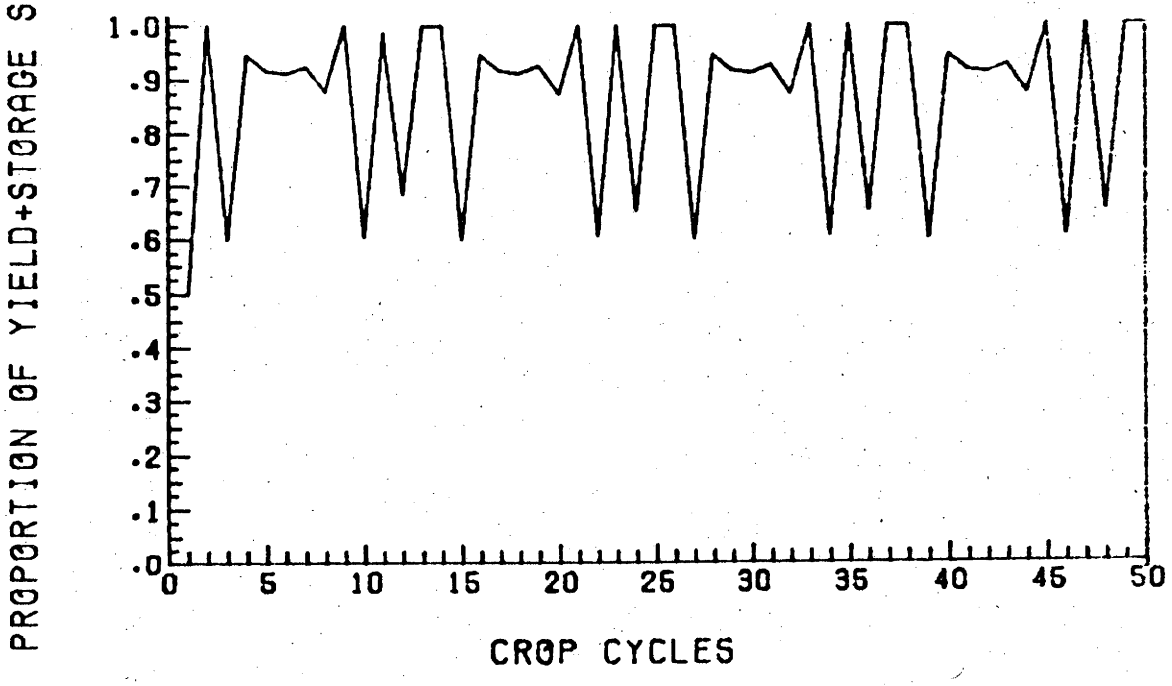

FIG 6.4.25: AGRIBUSINESS BASELINE
AREA PLANTED THROUGH TIME

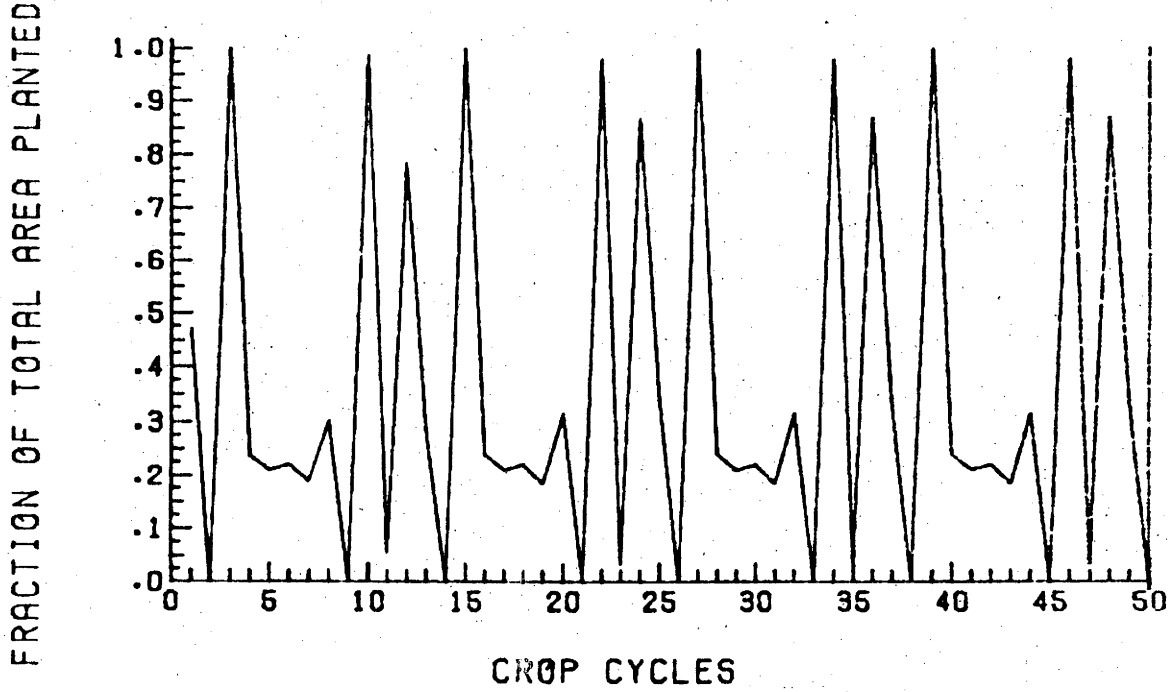

FIG 6.4.26: AGRIBUSINESS BASELINE
RETURN THROUGH TIME

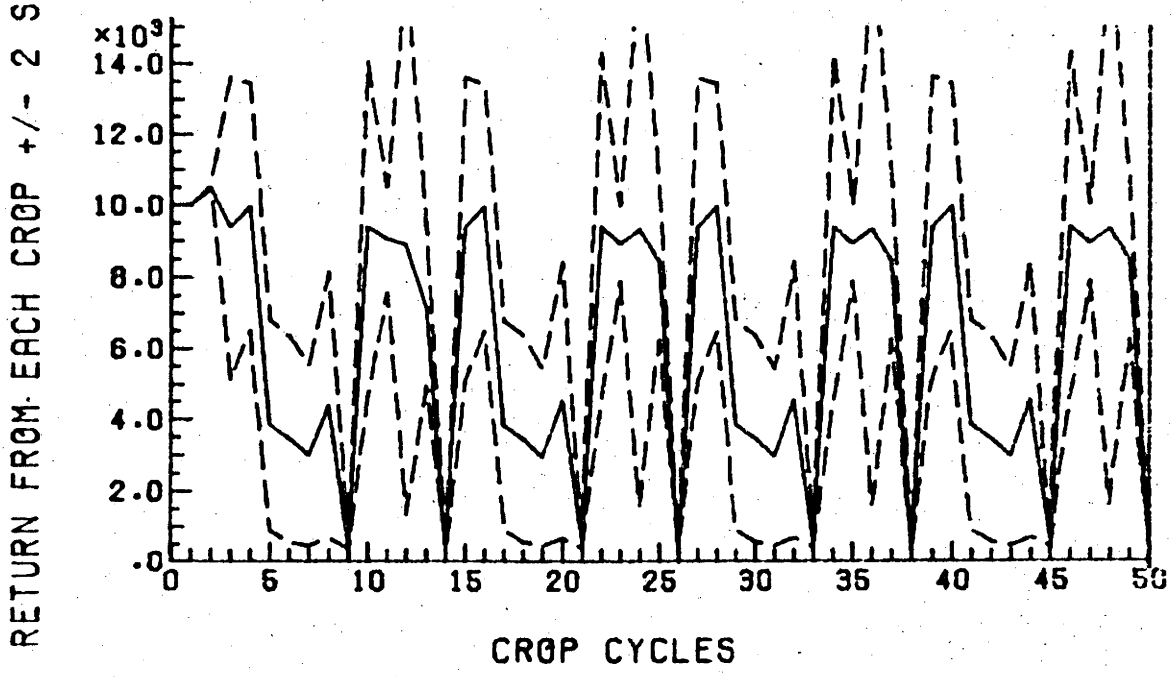

FIG 6.4.27: AGRIBUSINESS BASELINE
YIELD THROUGH TIME

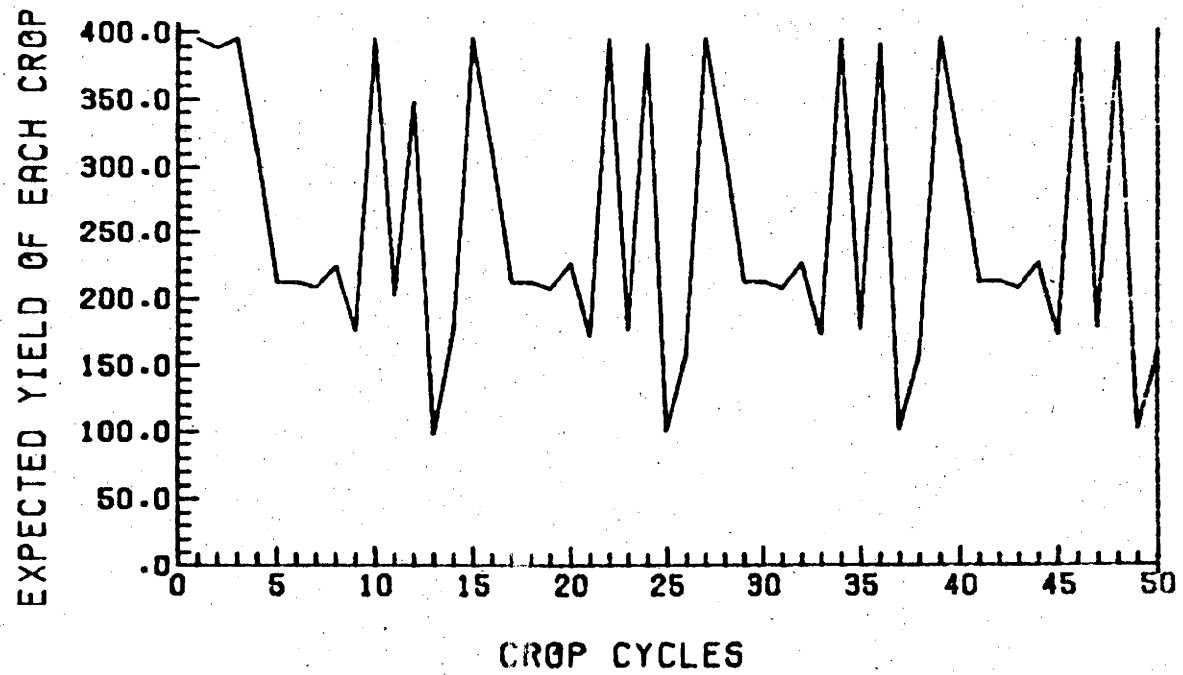

FIG 6.4.28: AGRIBUSINESS BASELINE
STORAGE THROUGH TIME

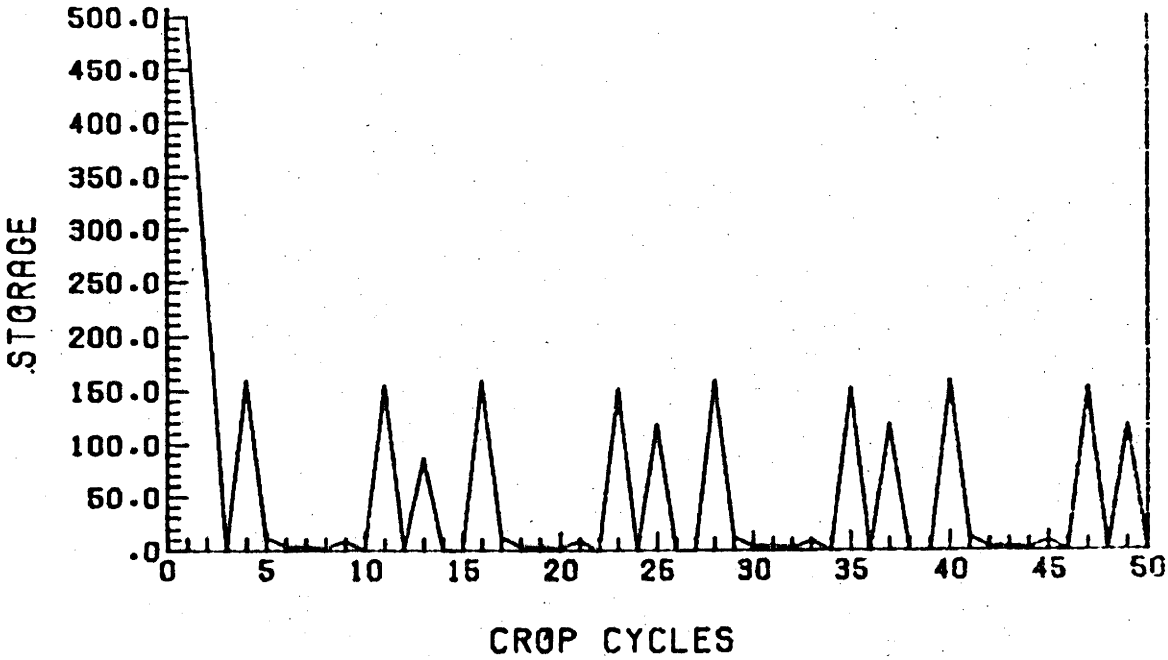

FIG 6.4.29: AGRIBUSINESS BASELINE
SUPERRACE THROUGH TIME

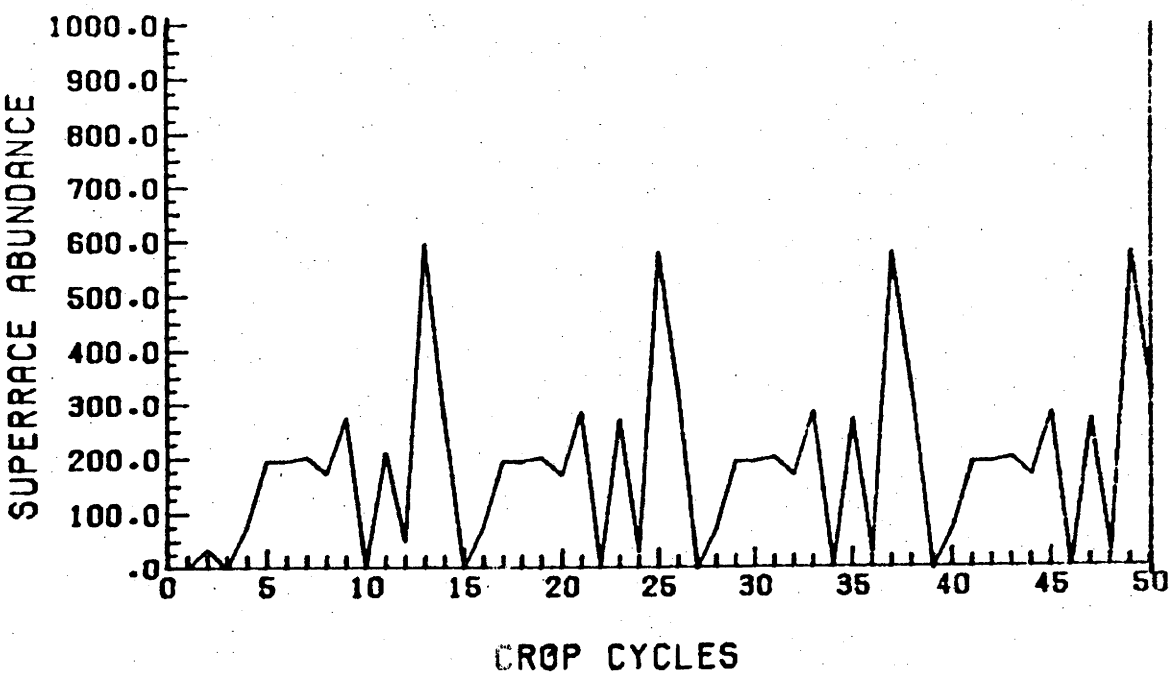

FIG 6.4.30A: AGRIBUSINESS BASELINE GAME
FRACTION SOLD VS STORAGE

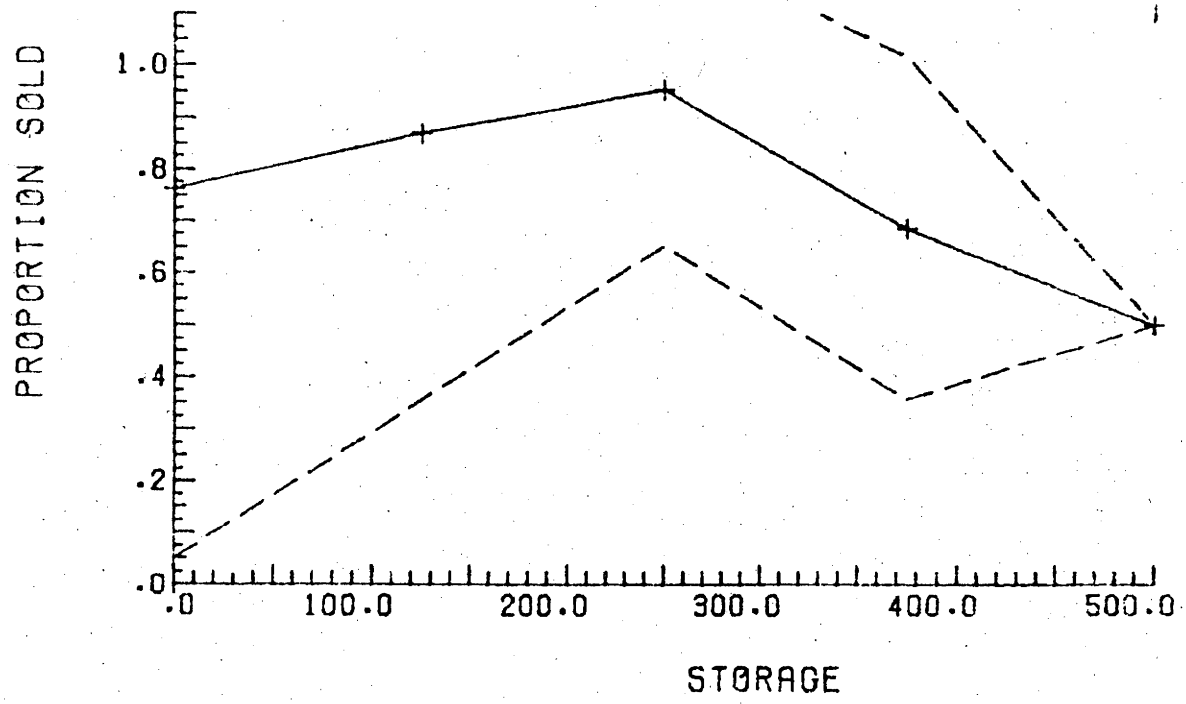

FIG 6.4.30B: AGRIBUSINESS BASELINE GAME
FRACTION SOLD VS SUPERRACE

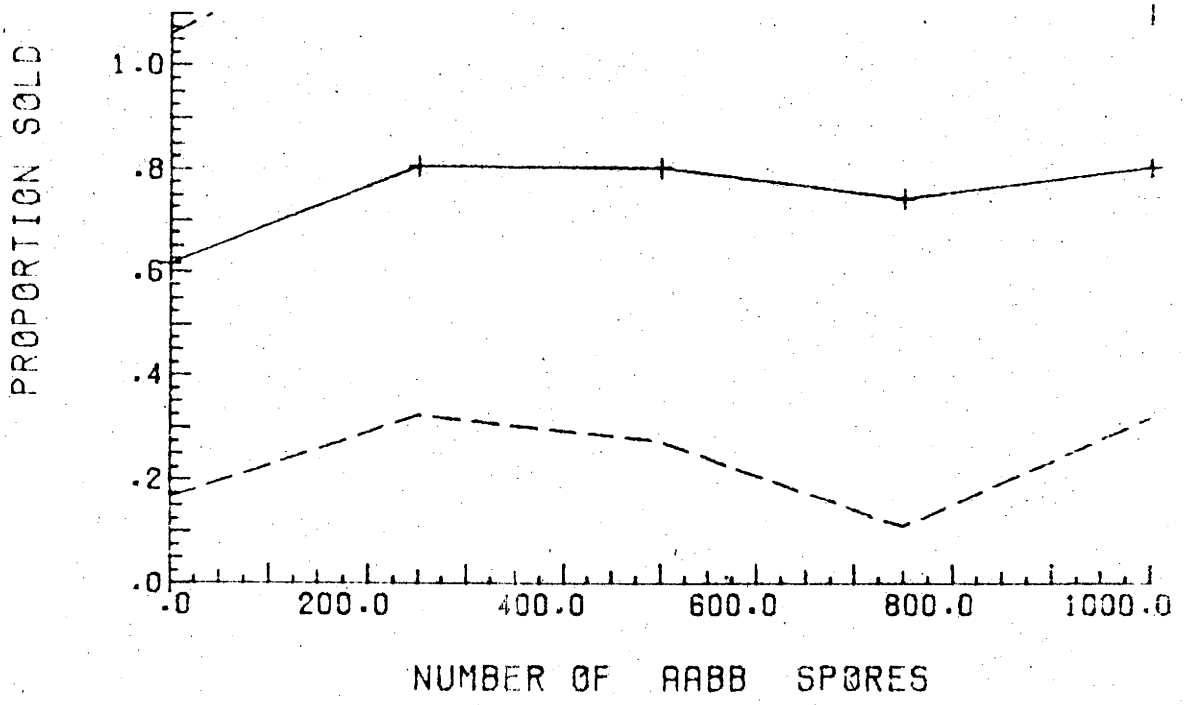

FIG 6.4.30C: AGRIBUSINESS BASELINE GAME
AREA PLANTED VS STORAGE

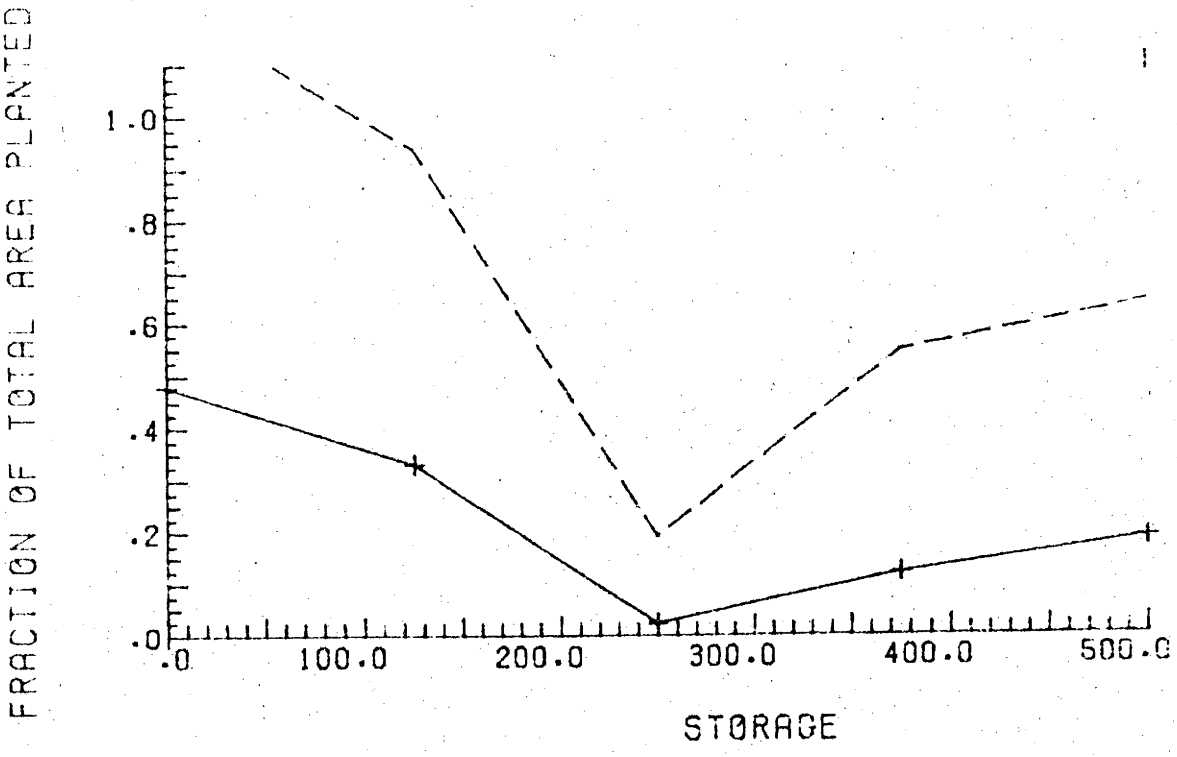

FIG 6.4.30D: AGRIBUSINESS BASELINE GAME
AREA PLANTED VS SUPERRACE

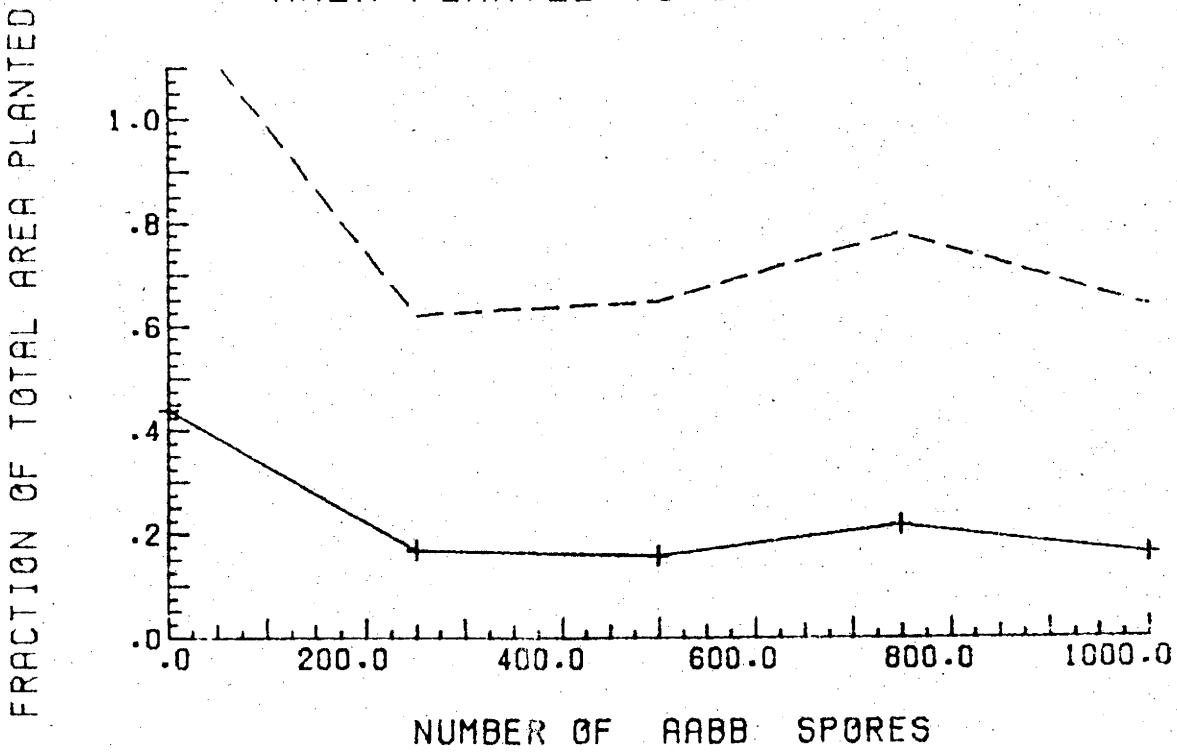

FIG 6.4.31A: AGRIBUSINESS BASELINE GAME
(1 - P(EXT)) VS STORAGE

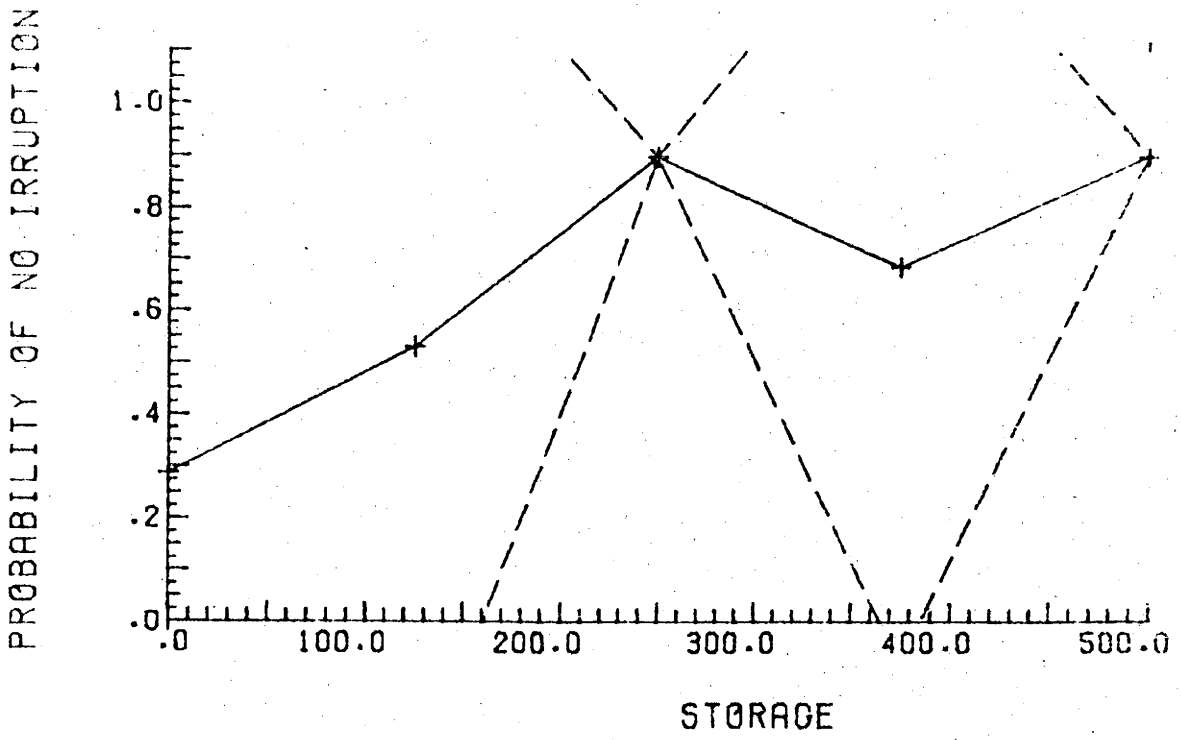

FIG 6.4.31B: AGRIBUSINESS BASELINE GAME
(1 - P(EXT)) VS SUPERRACE

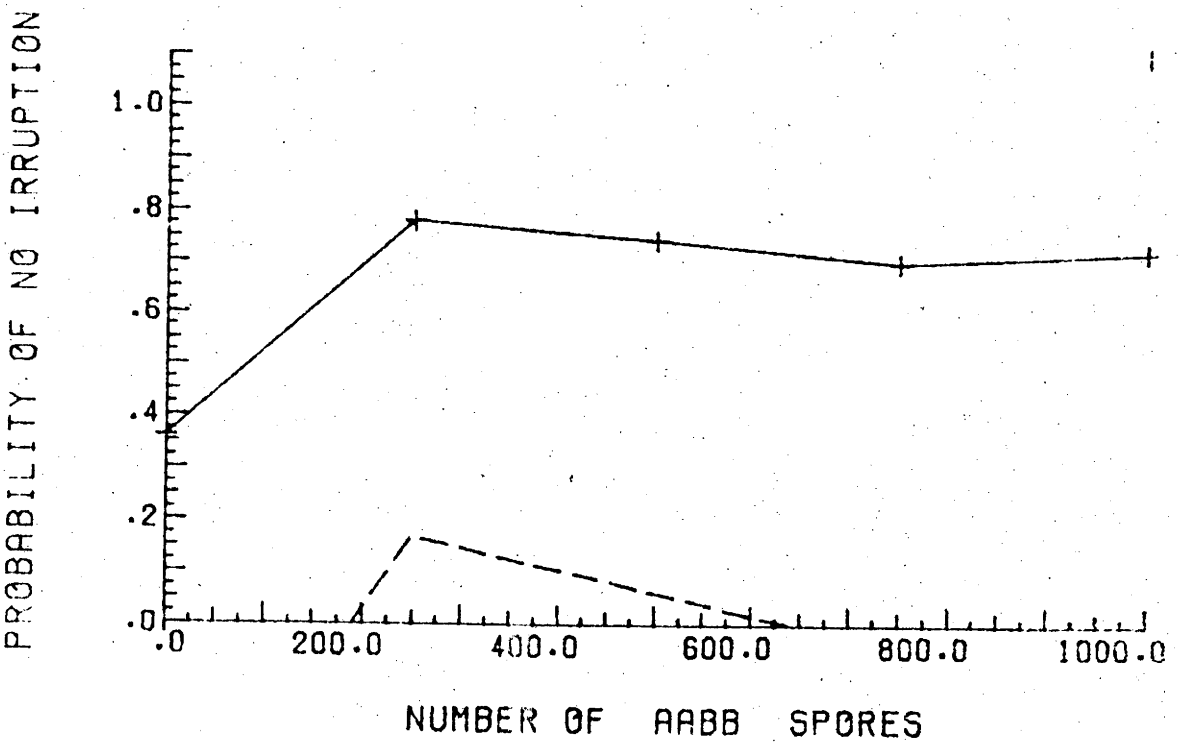

FIG 6.4.32A: AGRIBUSINESS BASELINE GAME
MUTATION RATE VS STORAGE

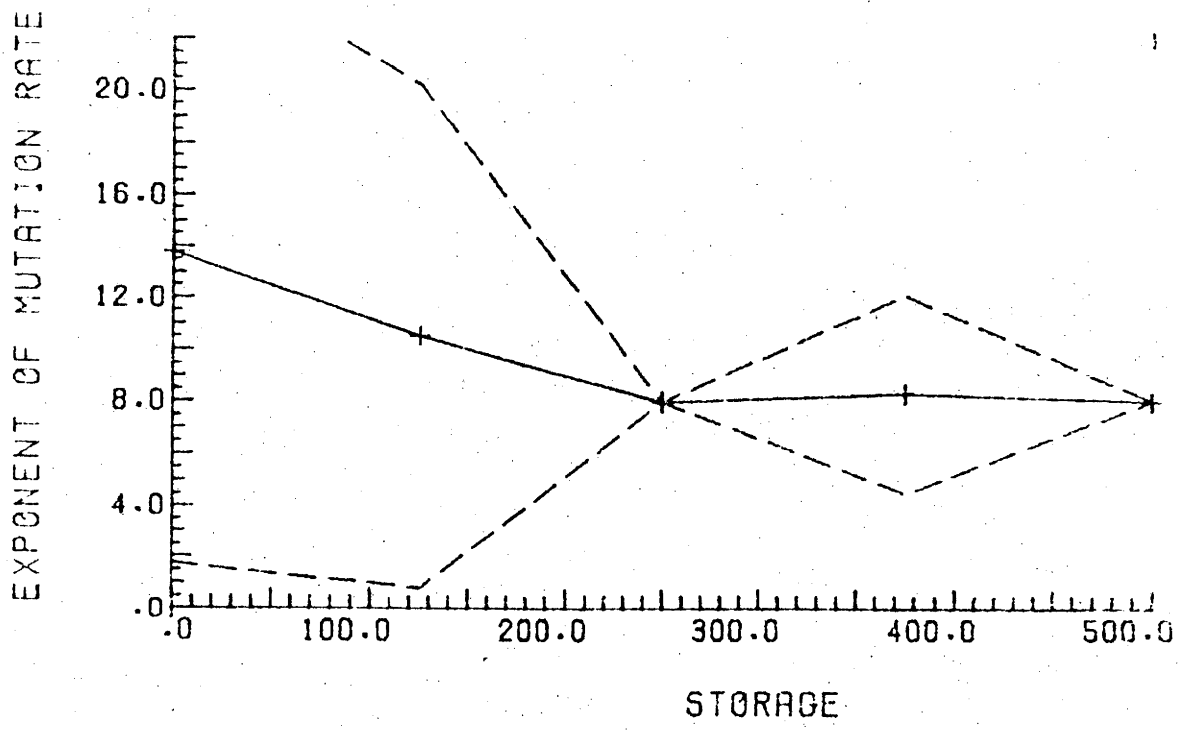

FIG 6.4.32B: AGRIBUSINESS BASELINE GAME
MUTATION RATE VS SUPERRACE

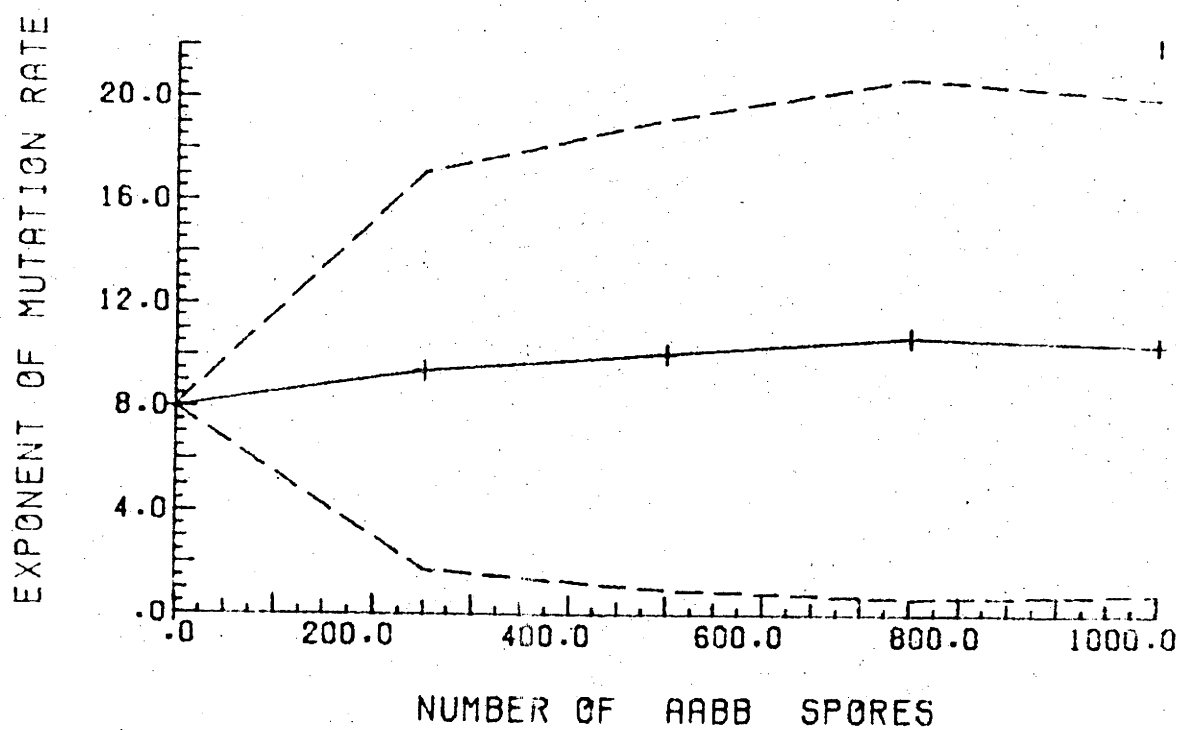

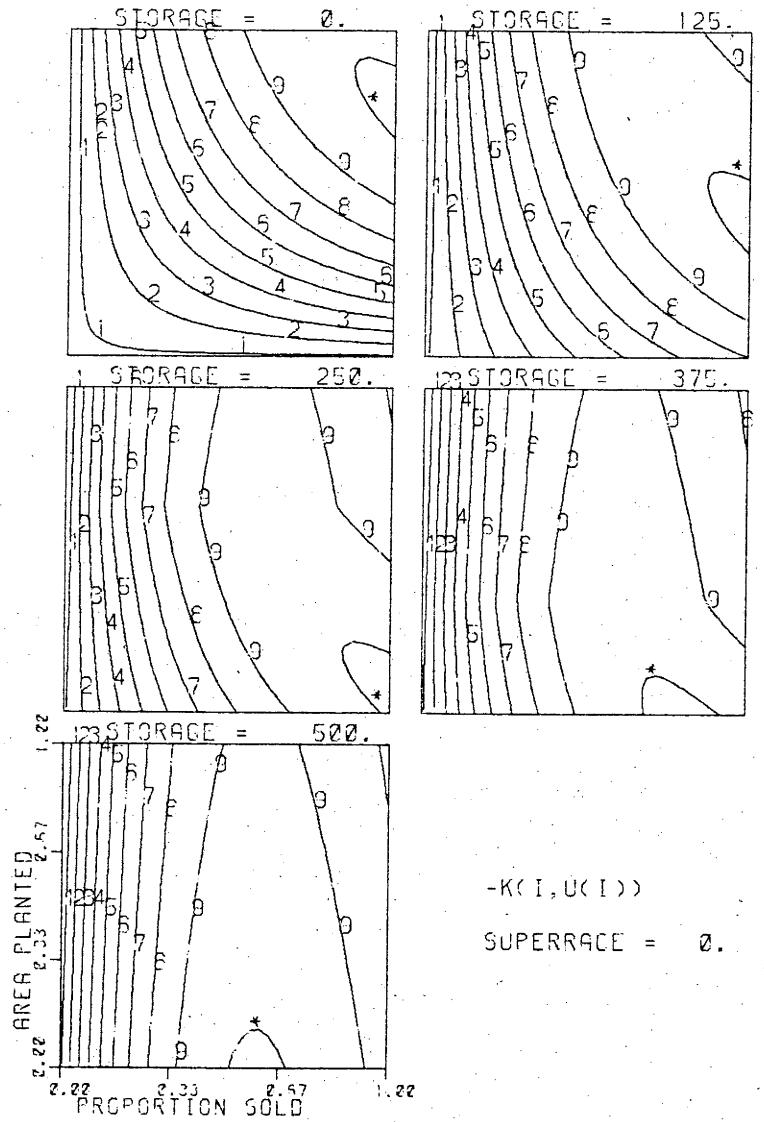

FIG. 6.4.33A: AGRIBUSINESS BASELINE GAME RETURN

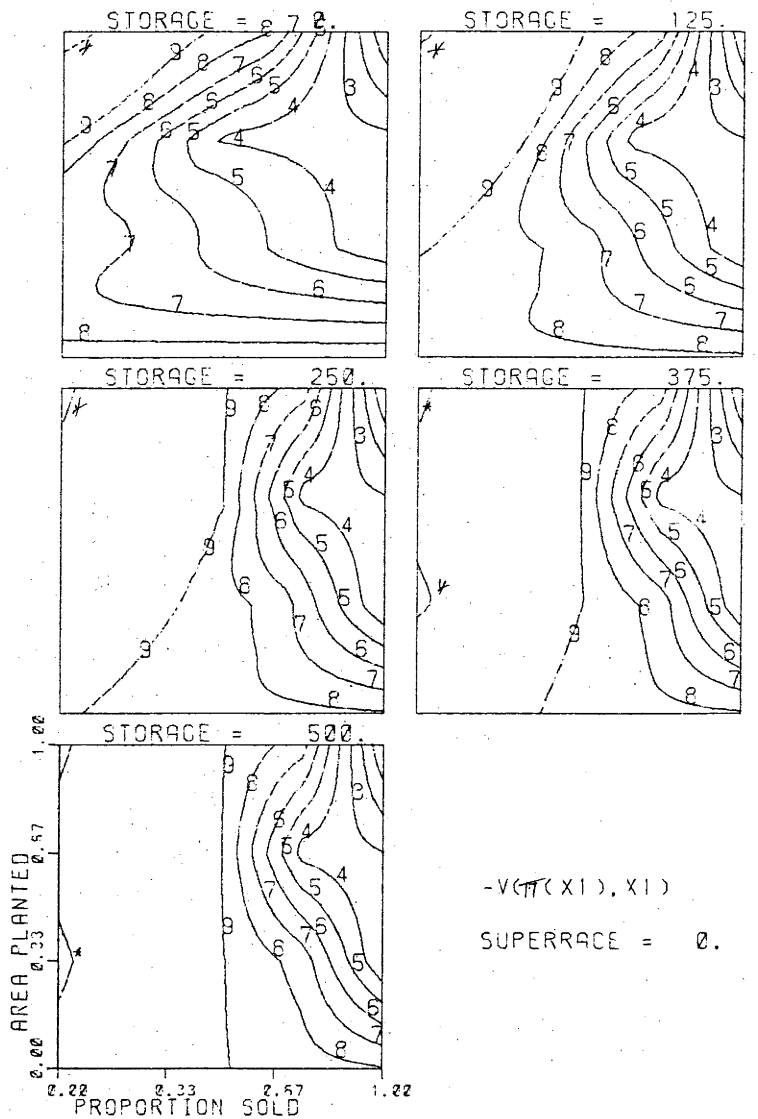

FIG. 6.4.33B: AGRIBUSINESS BASELINE GAME RETURN

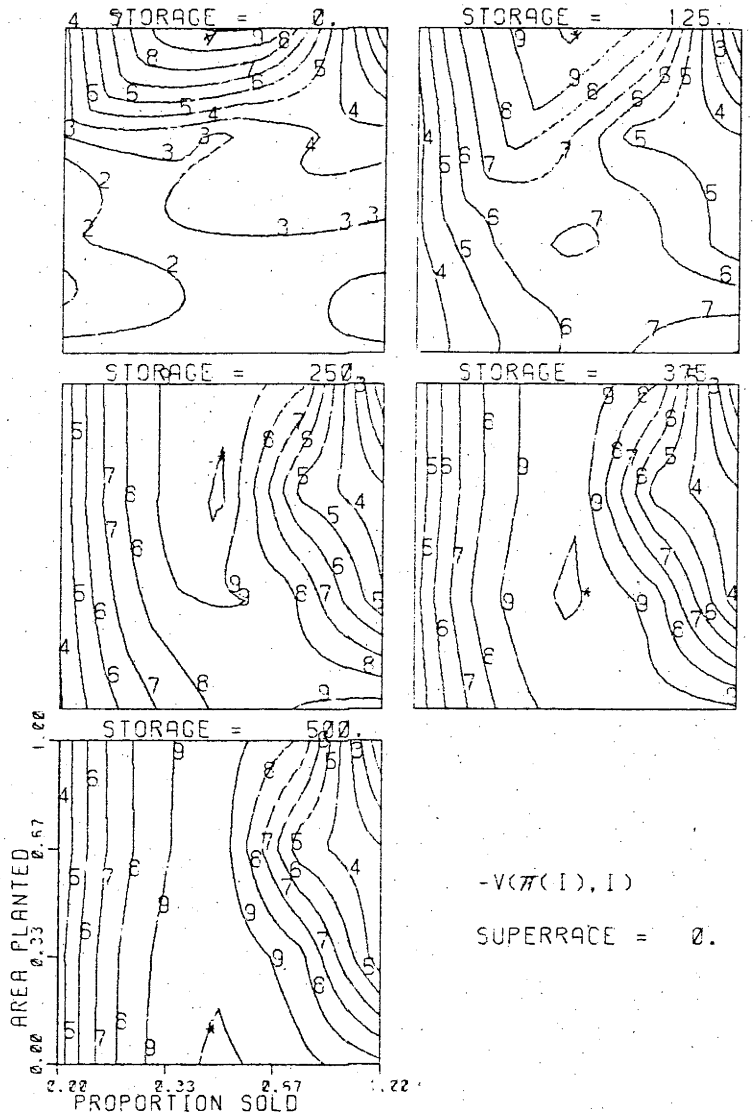

FIG. 5.4.33C: AGRIBUSINESS BASELINE GAME RETURN

FIG 6.4.34: AGRIBUSINESS BASELINE GAME
FRACTION SOLD THROUGH TIME

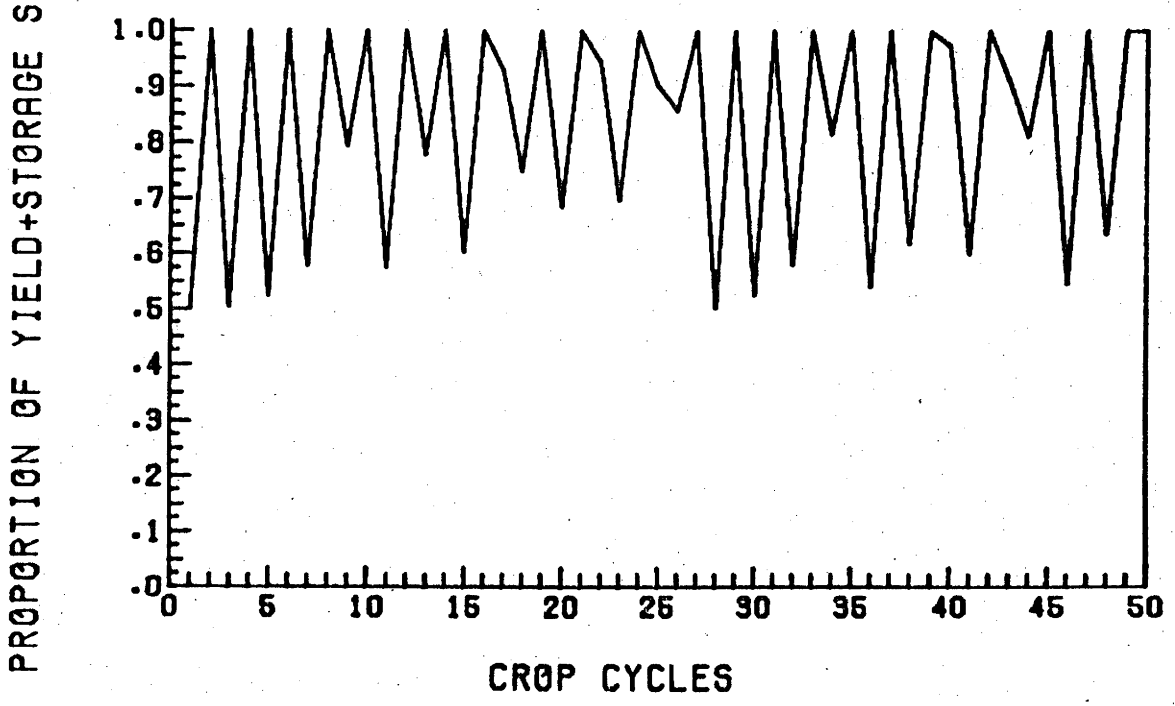

FIG 6.4.35: AGRIBUSINESS BASELINE GAME
AREA PLANTED THROUGH TIME

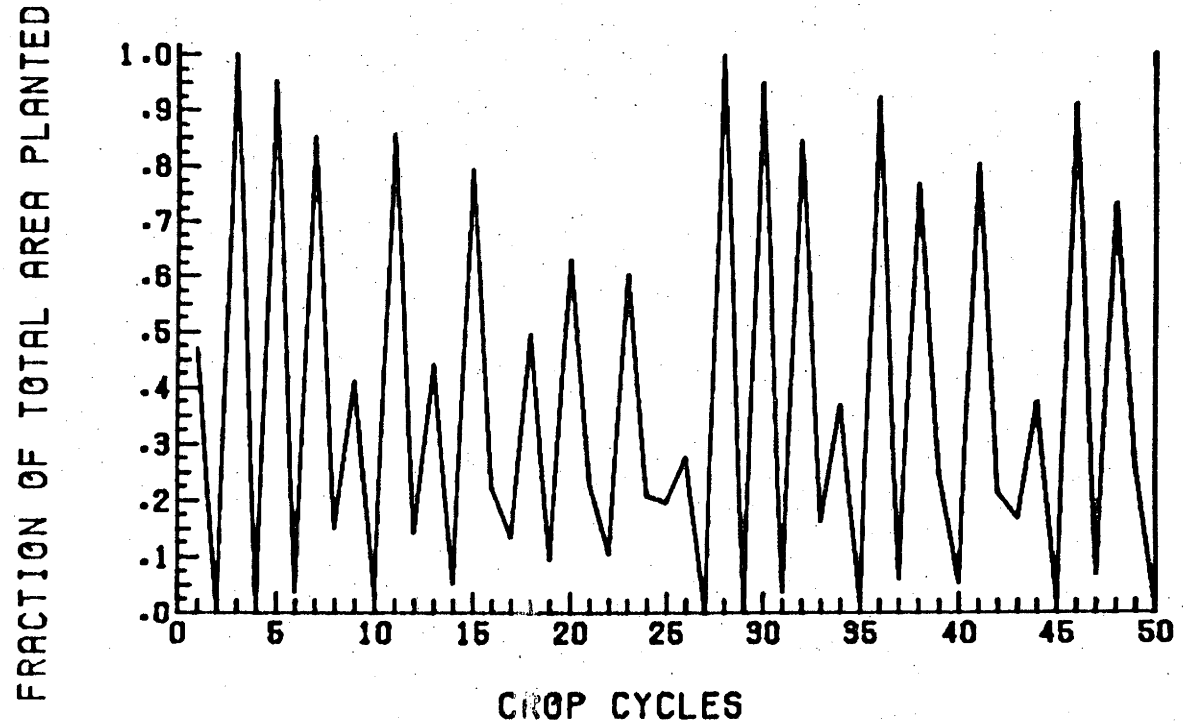

FIG 6.4.36: AGRIBUSINESS BASELINE GAME
RETURN THROUGH TIME

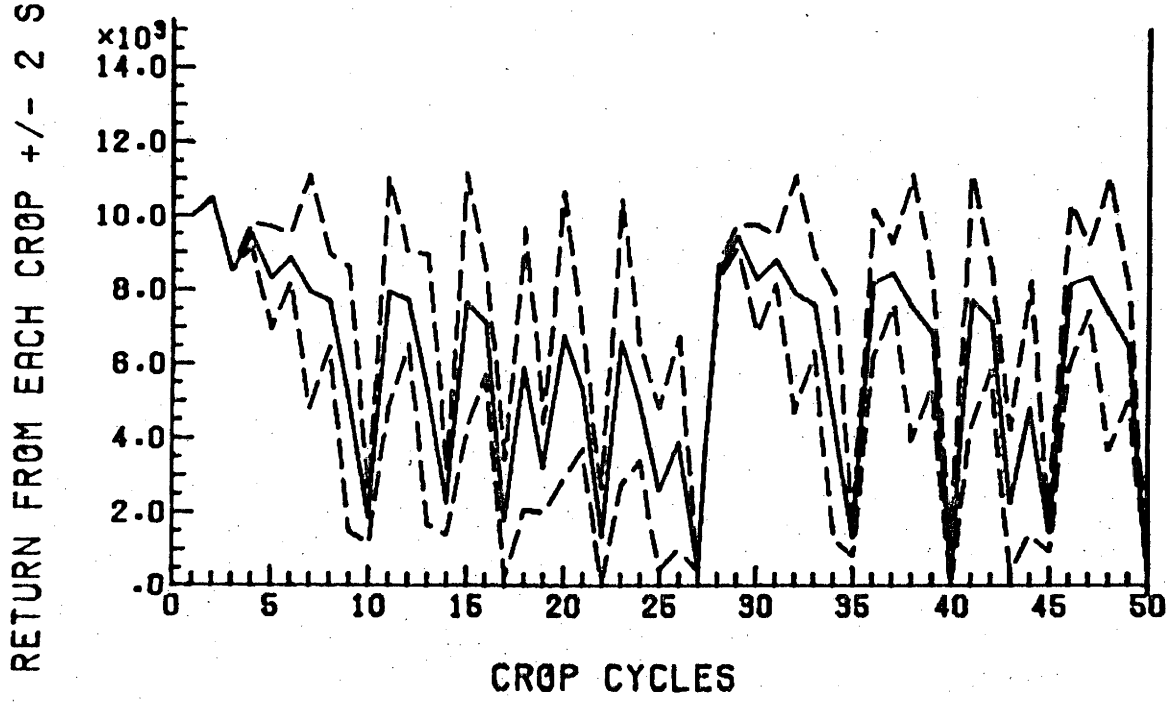

FIG 6.4.37: AGRIBUSINESS BASELINE GAME
SUPERRACE THROUGH TIME

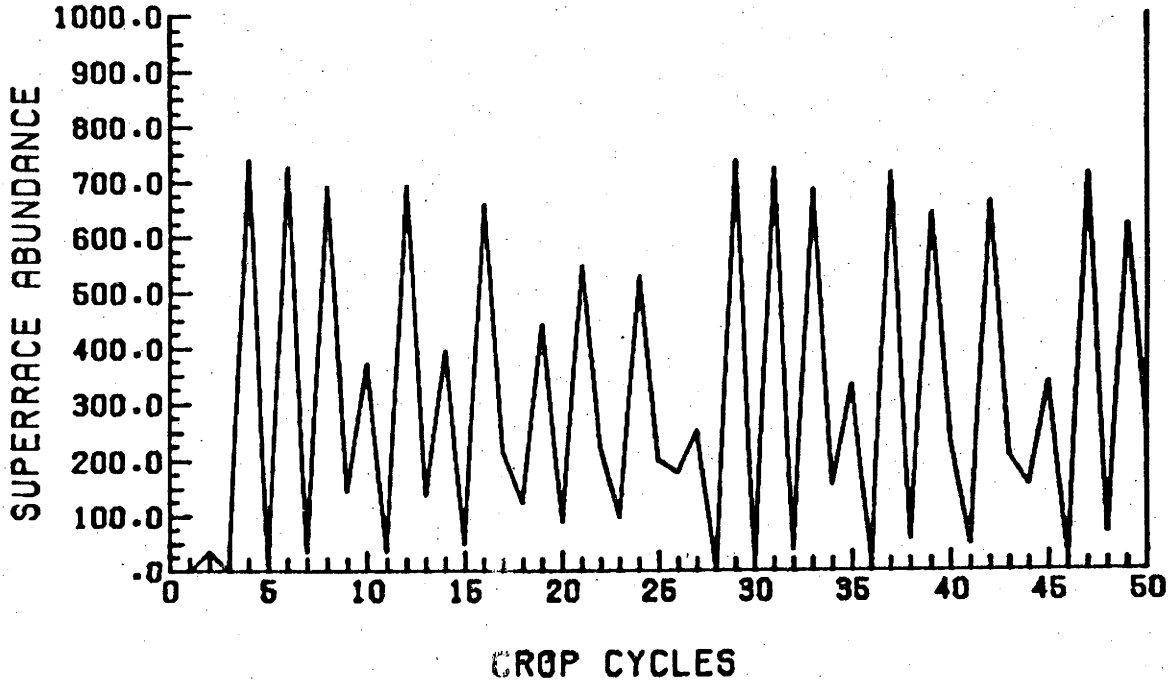

FIG 6.4.38: AGRIBUSINESS BASELINE GAME
(1-P(EXT)) THROUGH TIME

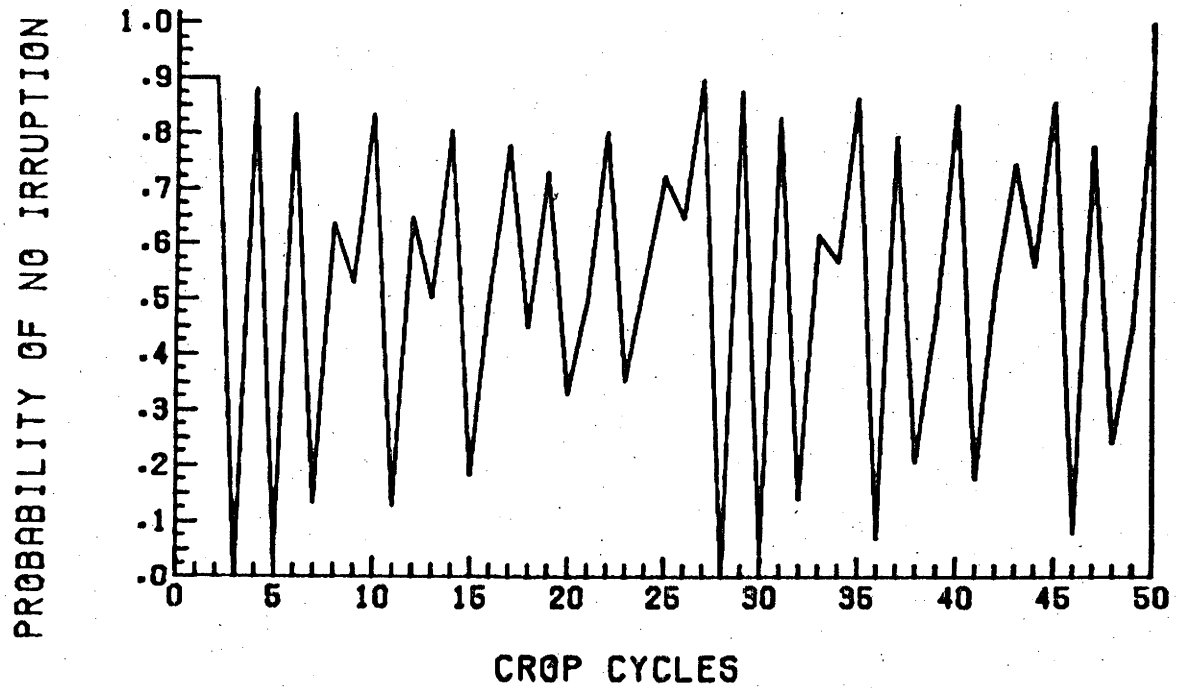

FIG 6.4.39: AGRIBUSINESS BASELINE GAME
MUTATION RATE THROUGH TIME

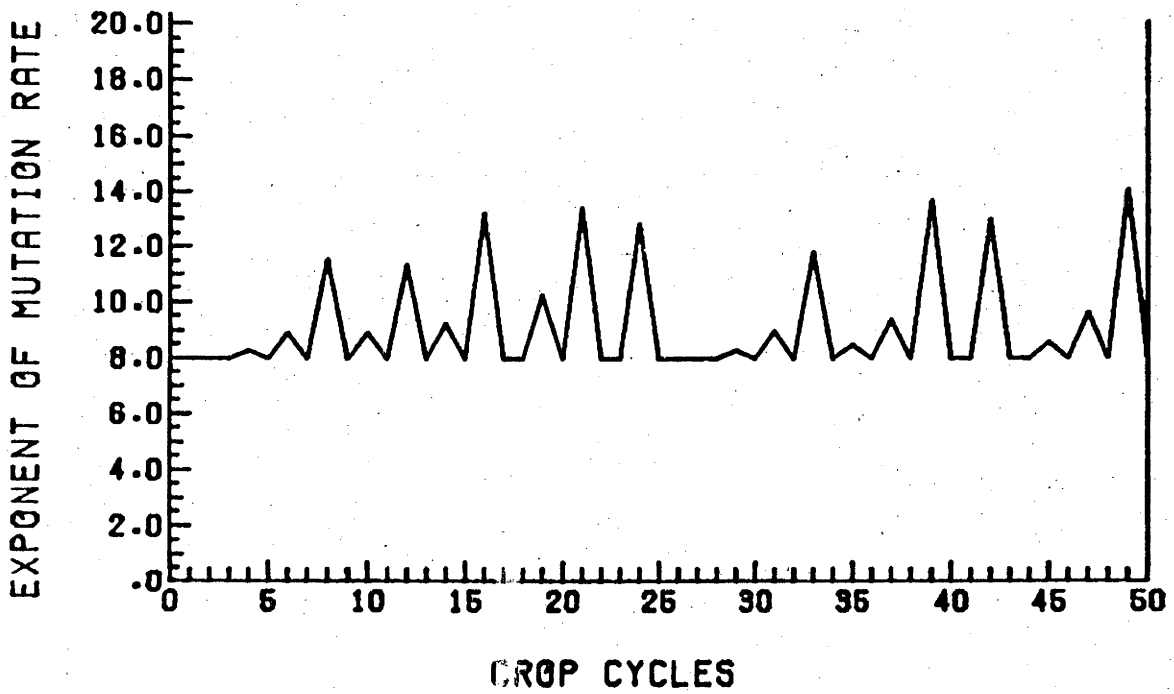

FIG. 6.4.40: PLANT WITH GENERAL RESISTANCE
DEPENDENCE OF YIELD ON PATHOGEN INPUT

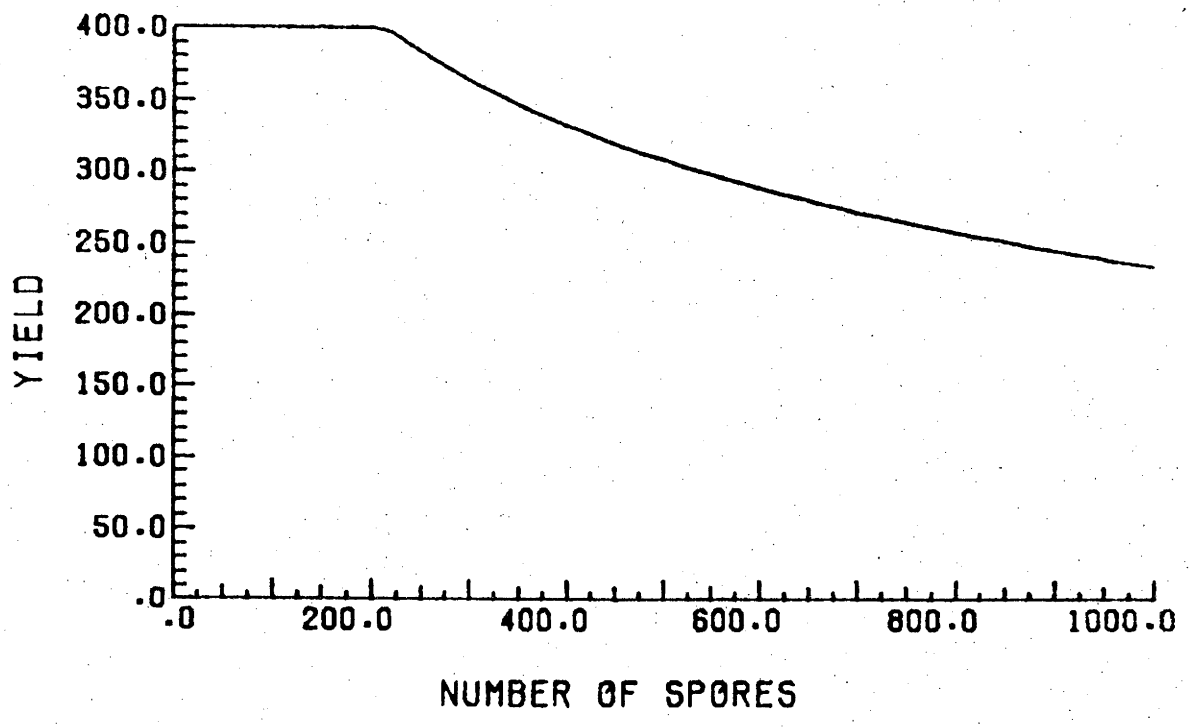

FIG. 6.4.41: NATURAL INCREASE OF PATHOGEN
ON PLANT WITH GENERAL RESISTANCE

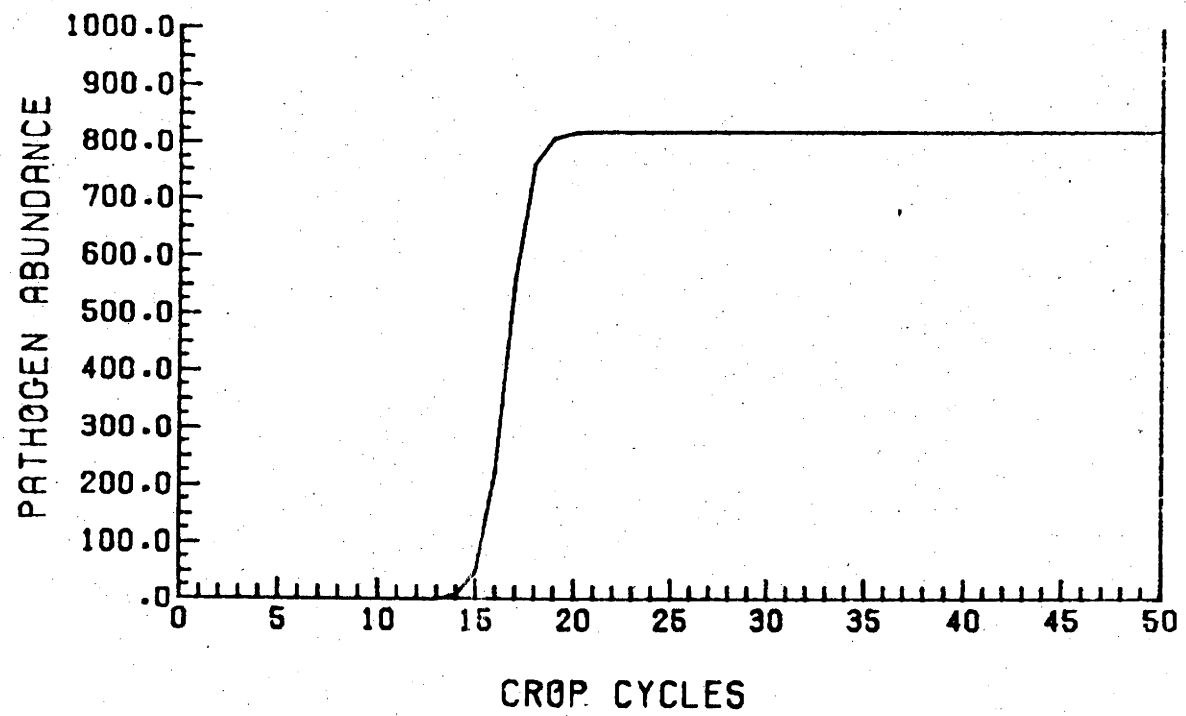

FIG 6.4.42A: AGRIBUSINESS + GENERAL RESISTANCE
FRACTION SOLD VS STORAGE

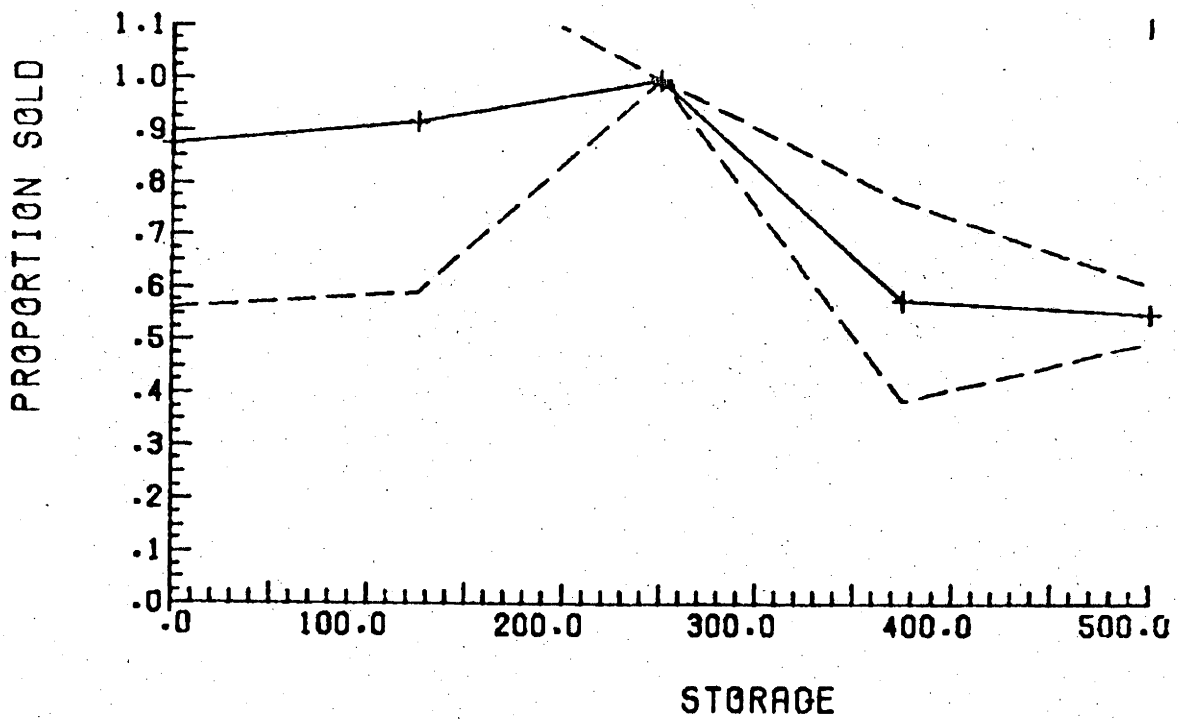

FIG 6.4.42B: AGRIBUSINESS + GENERAL RESISTANCE
FRACTION SOLD VS SUPERRACE

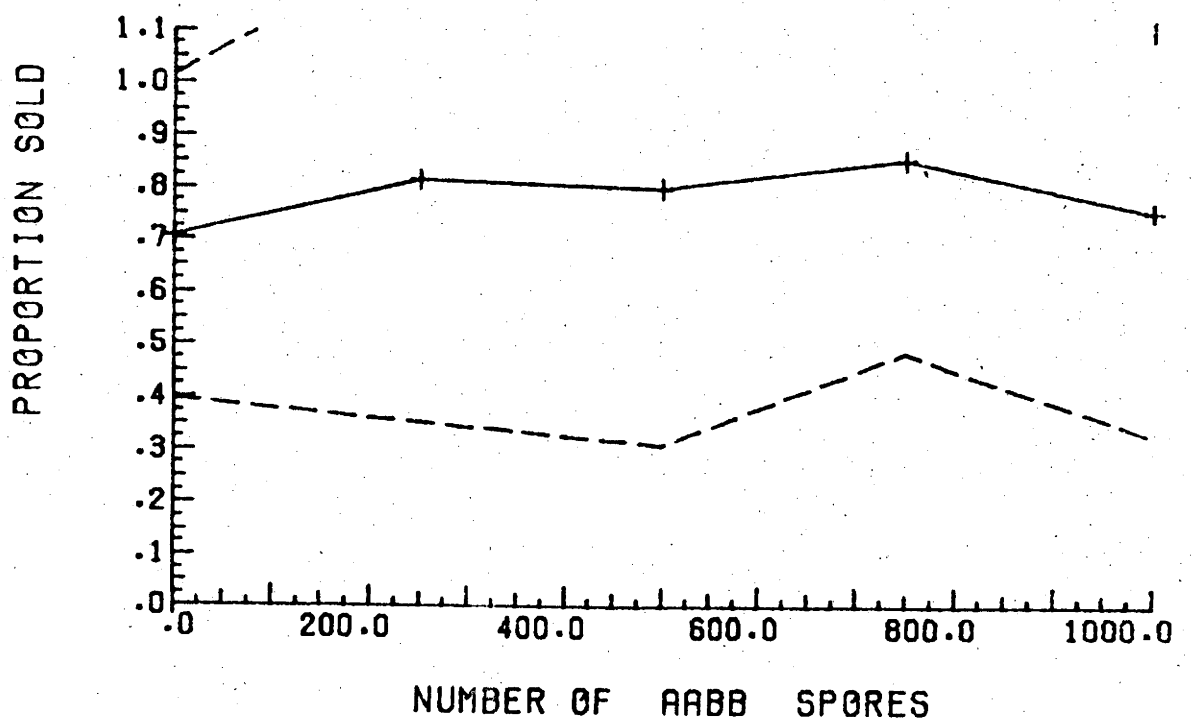

FIG 6.4.43A: AGRIBUSINESS + GENERAL RESISTANCE
AREA PLANTED VS STORAGE

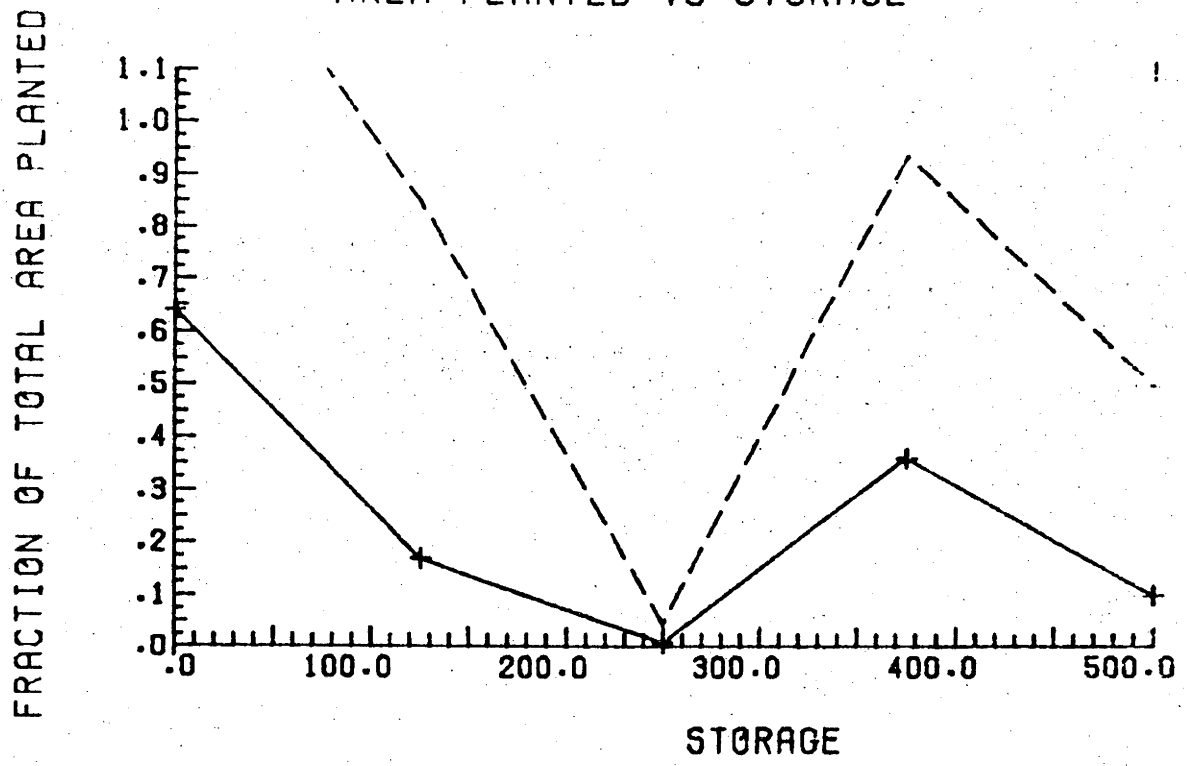

FIG 6.4.43B: AGRIBUSINESS + GENERAL RESISTANCE
AREA PLANTED VS SUPERRACE

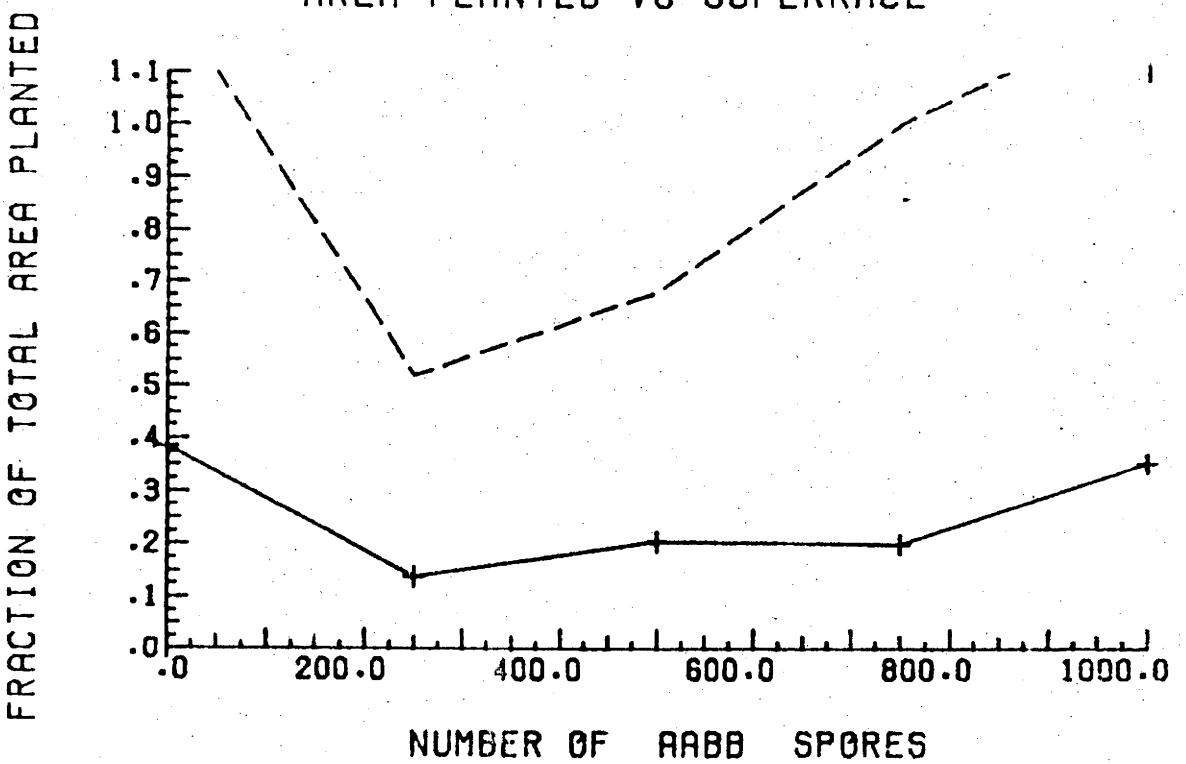

FIG 6.4.44A: AGRIBUSINESS + GENERAL RESISTANCE
E(RETURN) VS STORAGE

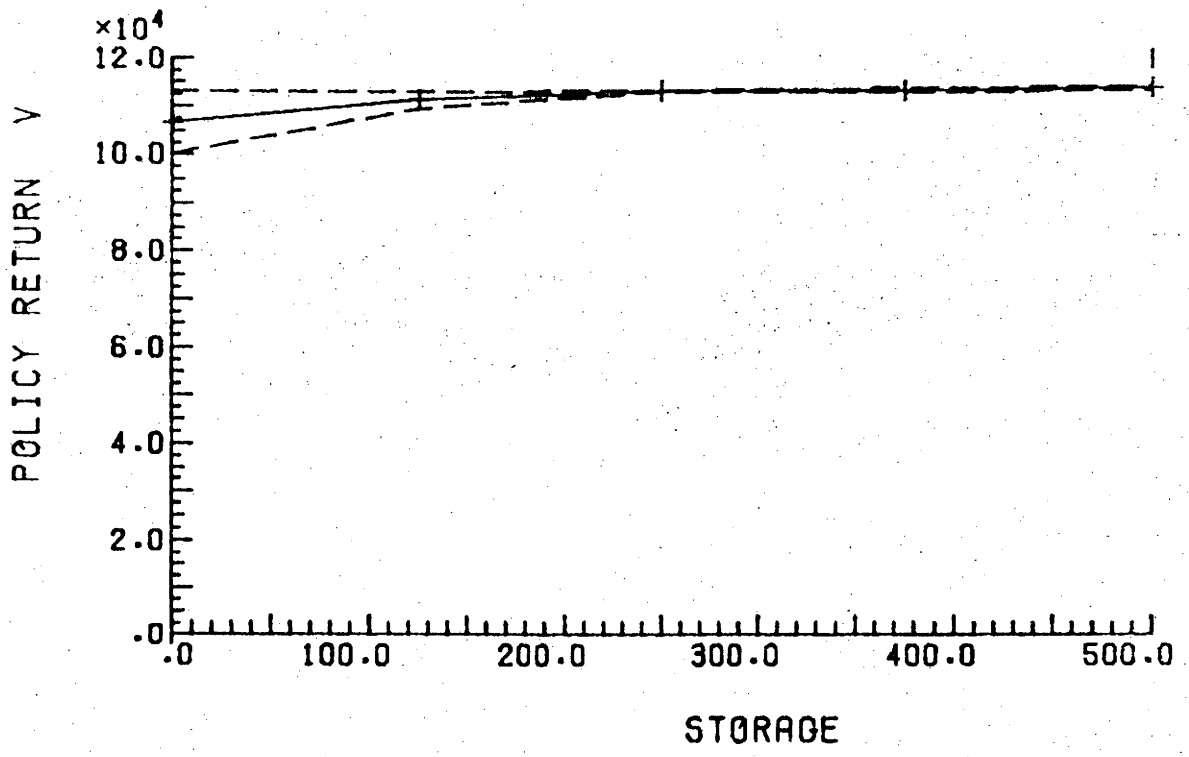

FIG 6.4.44B: AGRIBUSINESS + GENERAL RESISTANCE
E(RETURN) VS SUPERRACE

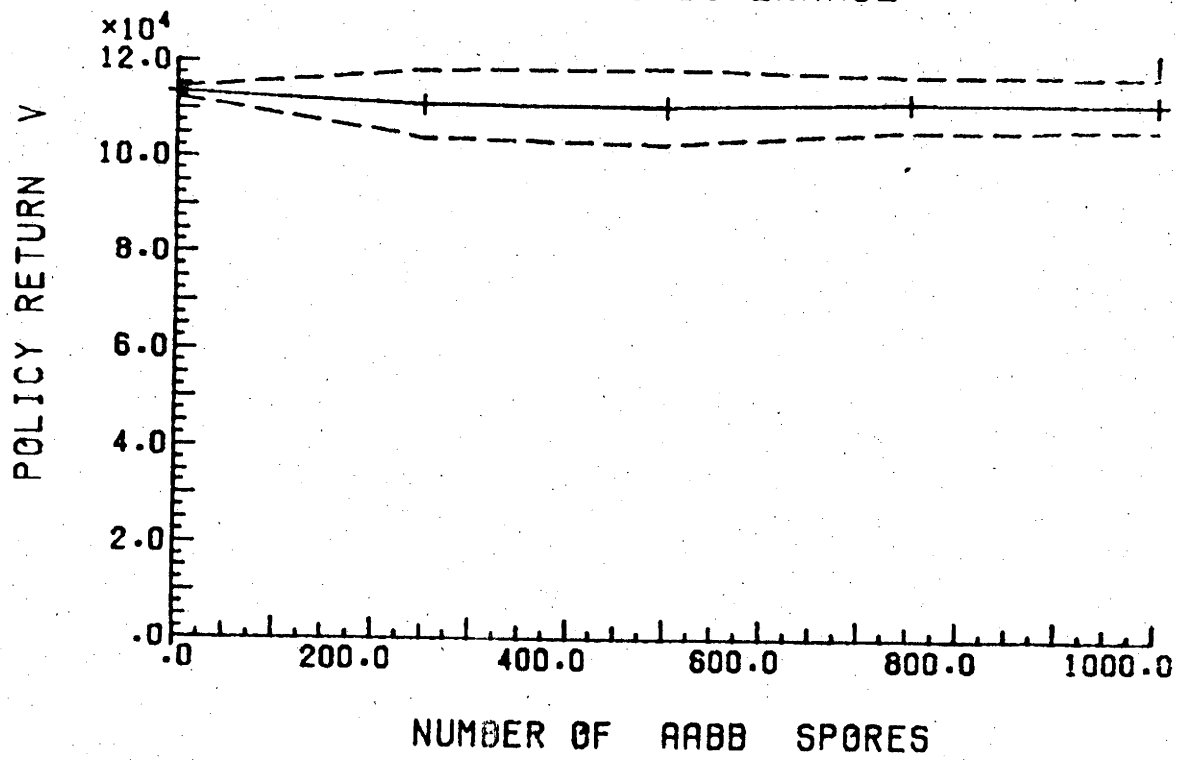

FIG 6.4.45A: AGRIBUSINESS + GENERAL RESISTANCE
STORAGE LEFT VS STORAGE

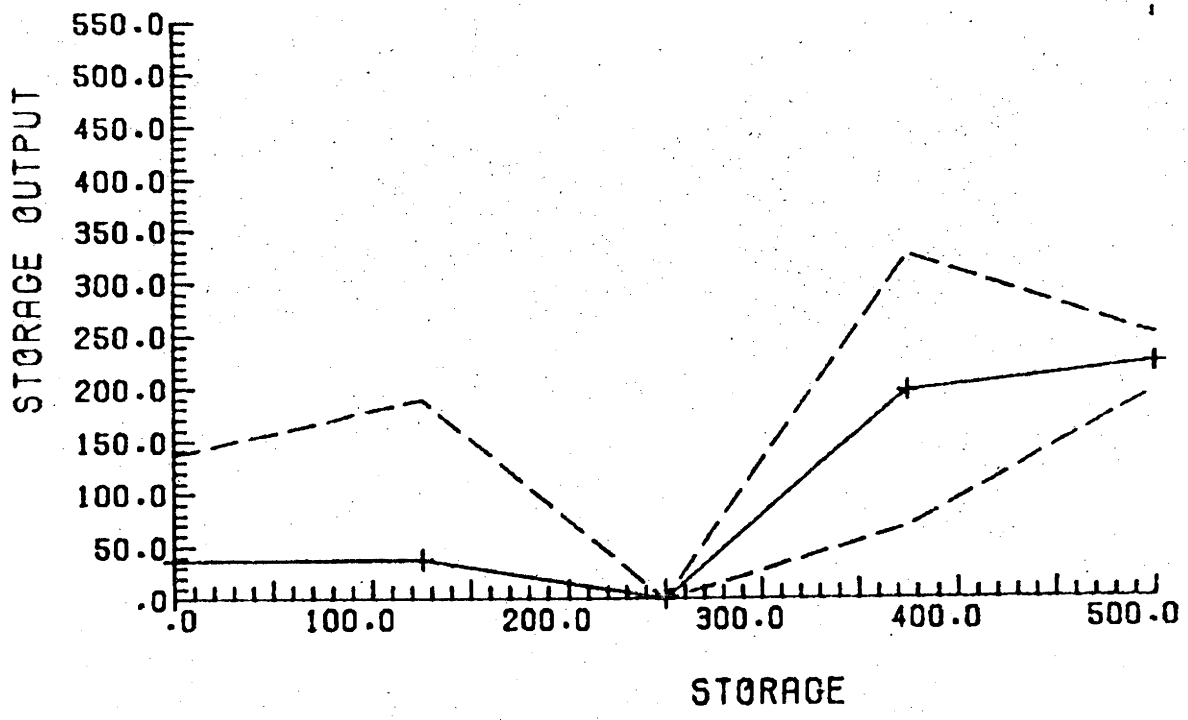

FIG 6.4.45B: AGRIBUSINESS + GENERAL RESISTANCE
STORAGE LEFT VS SUPERRACE

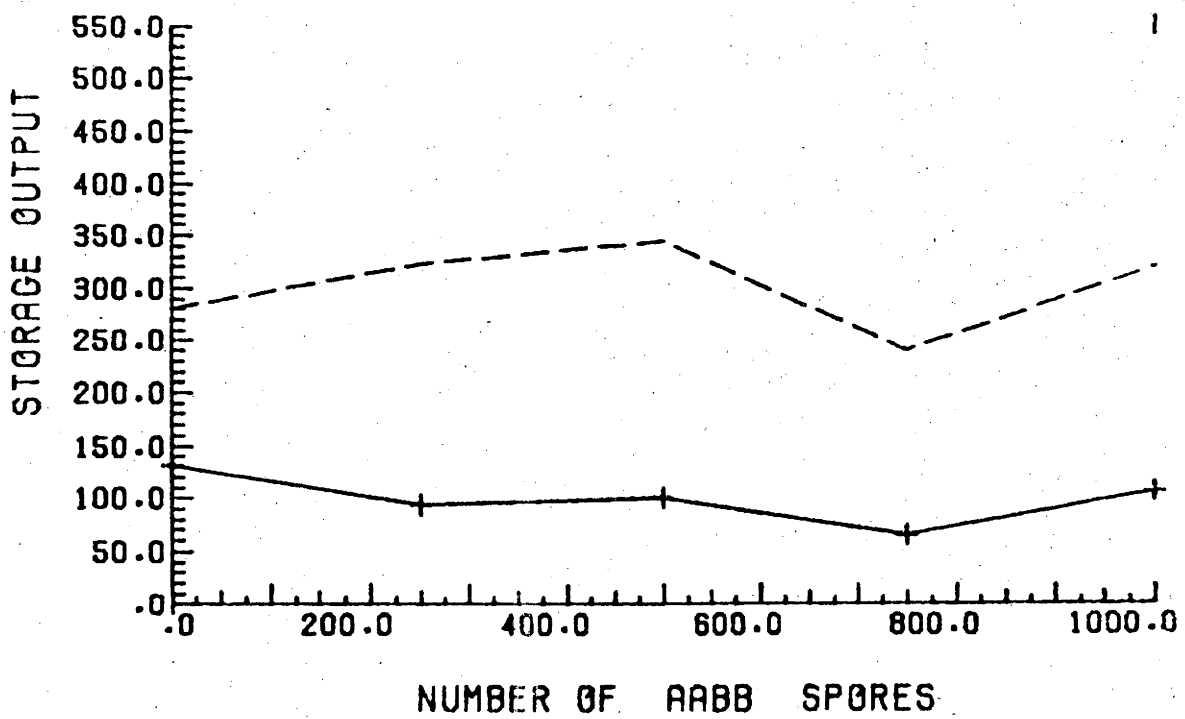

FIG 6.4.46A: AGRIBUSINESS + GENERAL RESISTANCE
AABB OUTPUT VS STORAGE

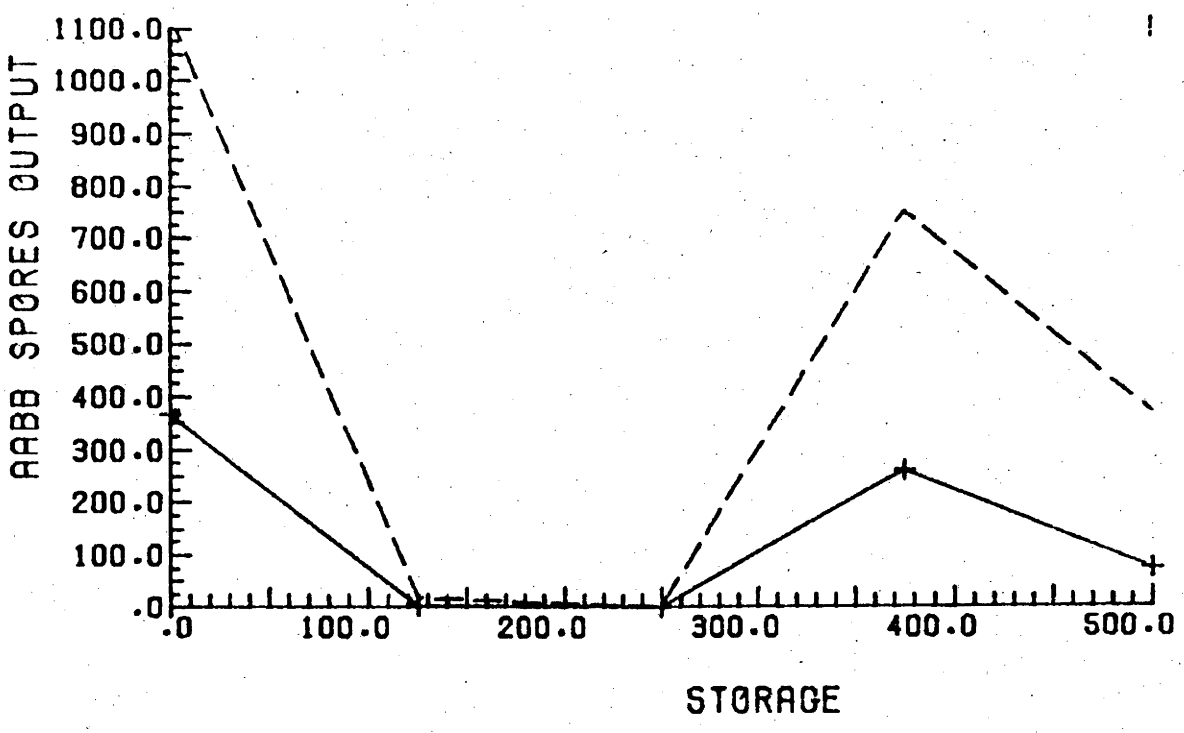

FIG 6.4.46B: AGRIBUSINESS + GENERAL RESISTANCE
AABB OUTPUT VS SUPERRACE

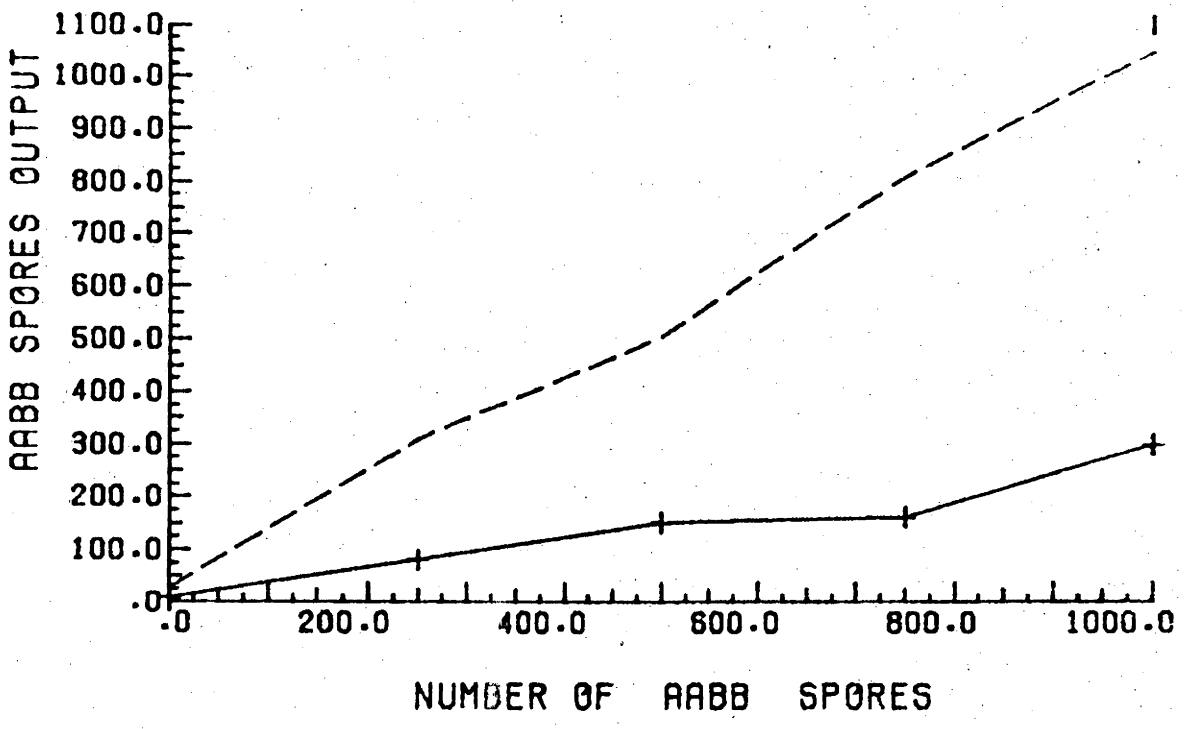

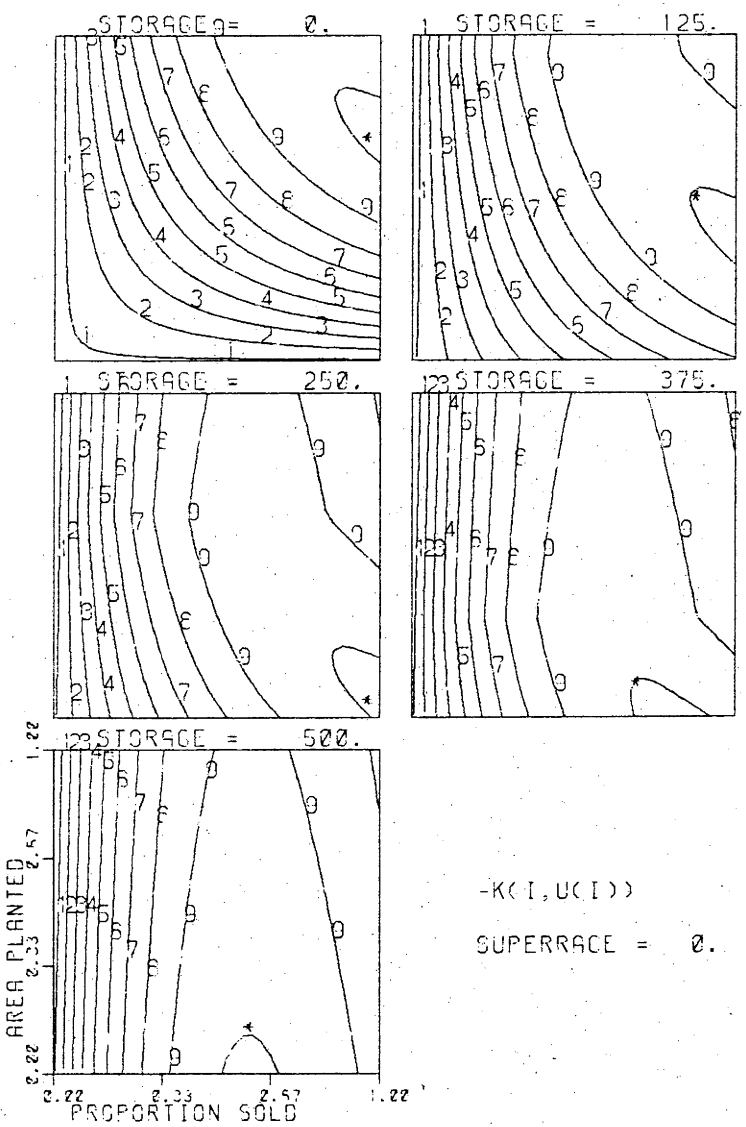

FIG. 5 4.47A: RETURN WITH GENERAL RESISTANCE

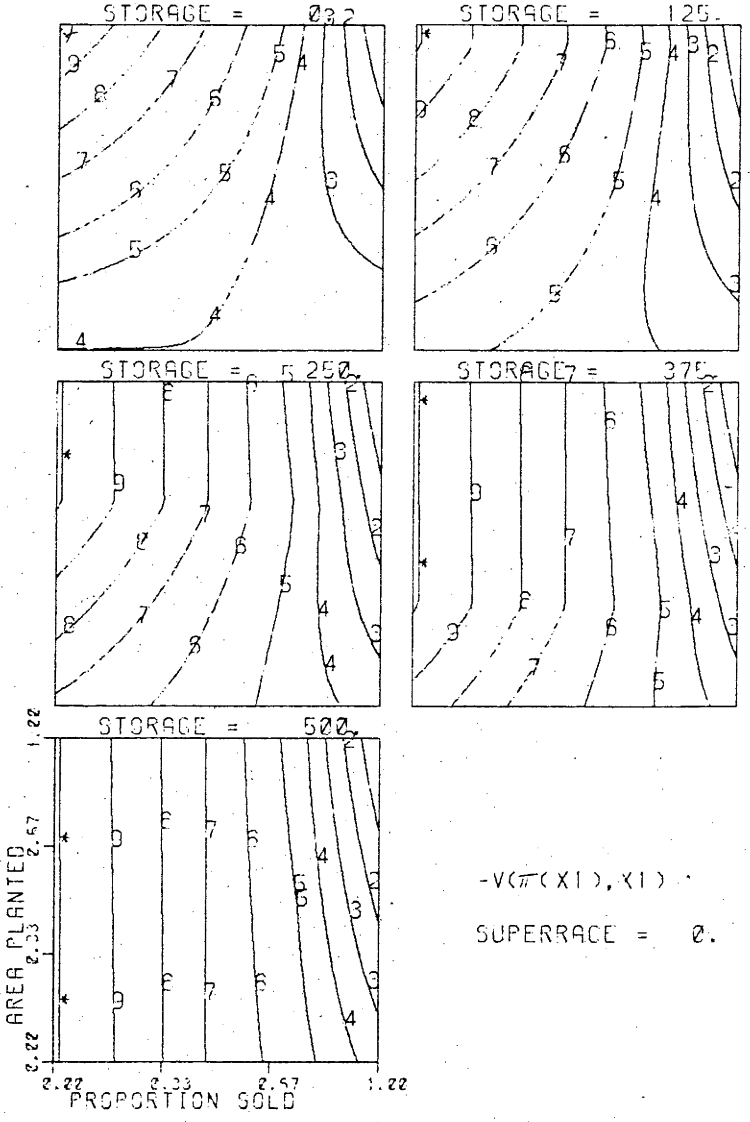

FIG. 5 4.47B: RETURN WITH GENERAL RESISTANCE

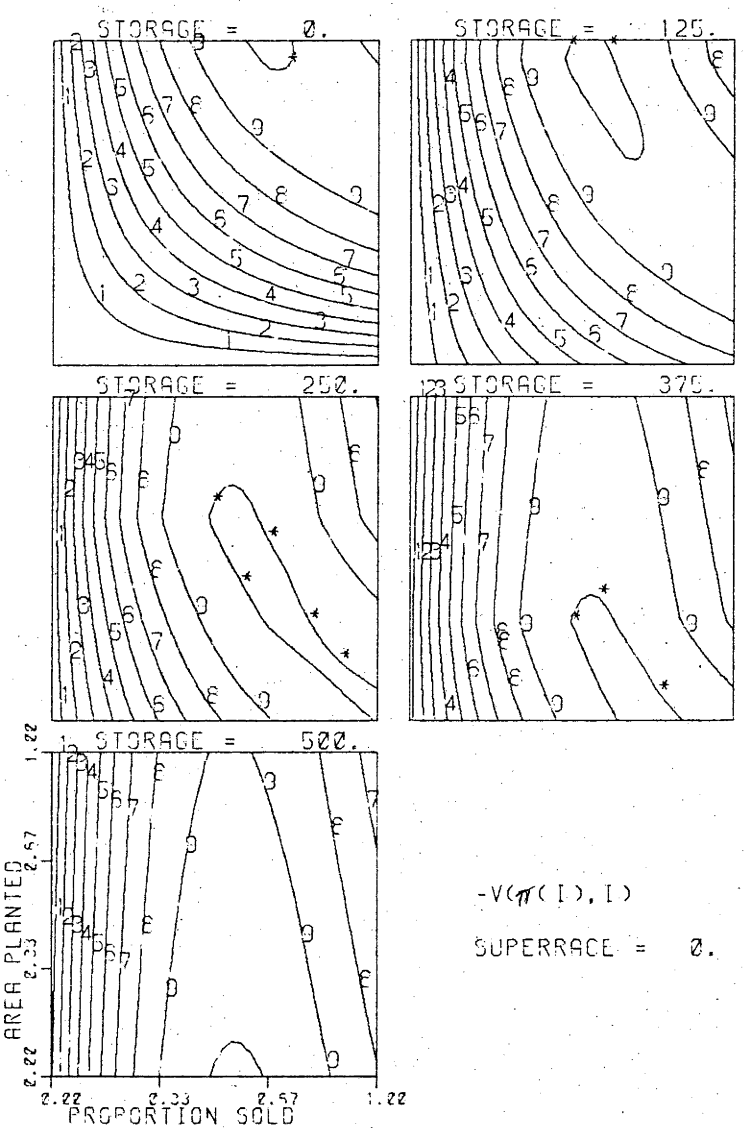

FIG. 6 4.47C: RETURN WITH GENERAL RESISTANCE

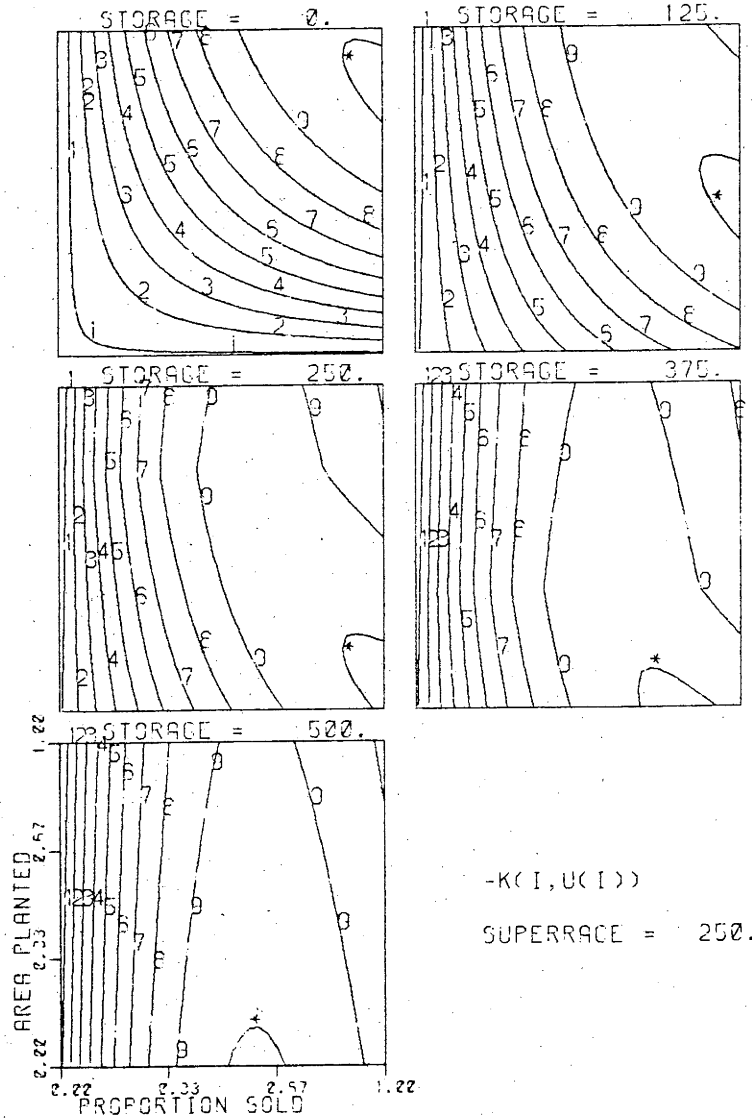

FIG. 6 4.48A· RETURN WITH GENERAL RESISTANCE

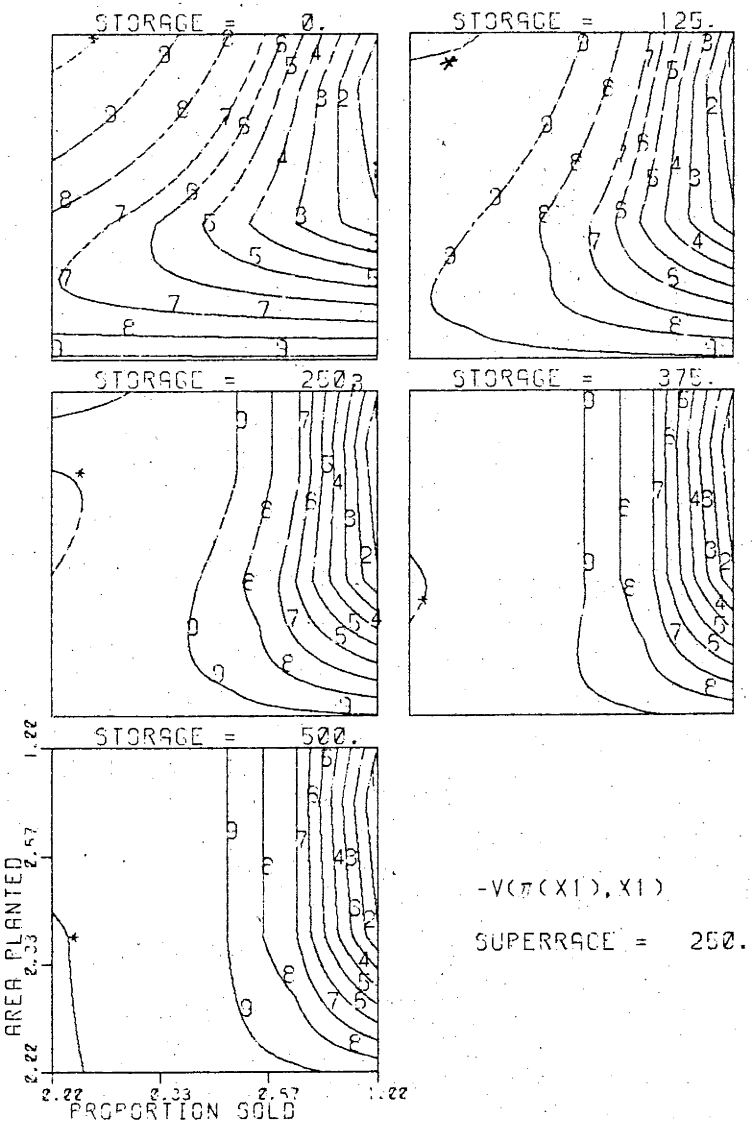

FIG. 5.4.48B- RETURN WITH GENERAL RESISTANCE

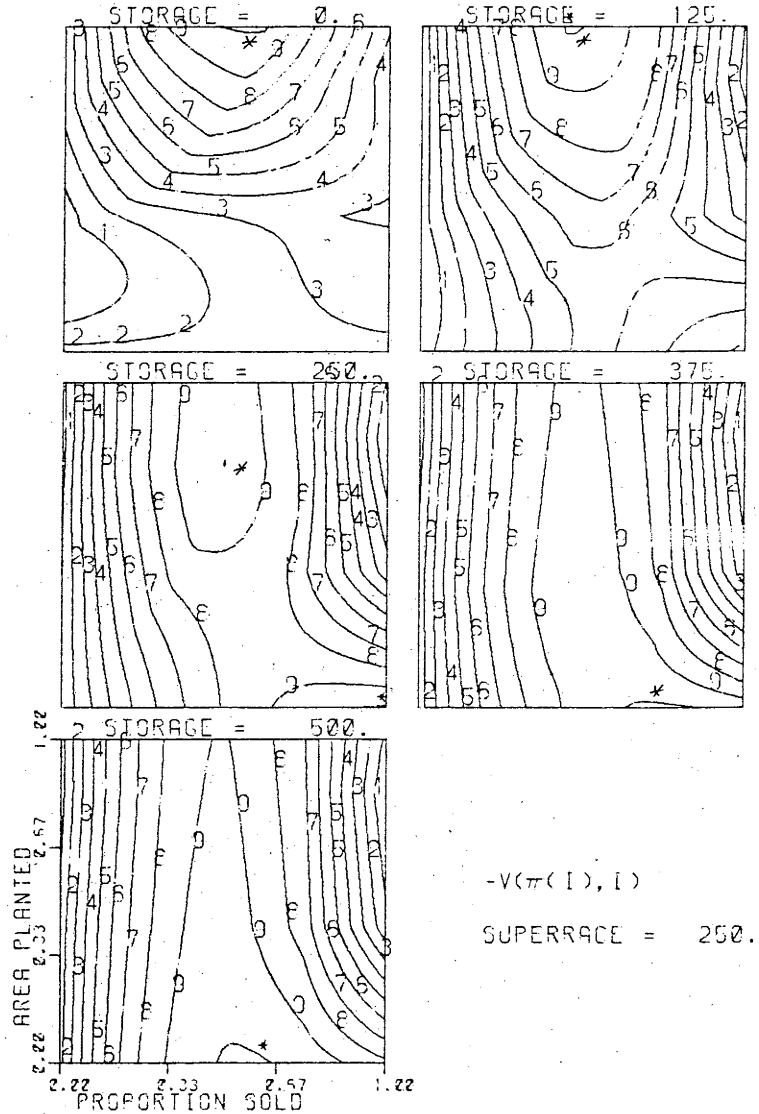

FIG. 5.4.48C- RETURN WITH GENERAL RESISTANCE

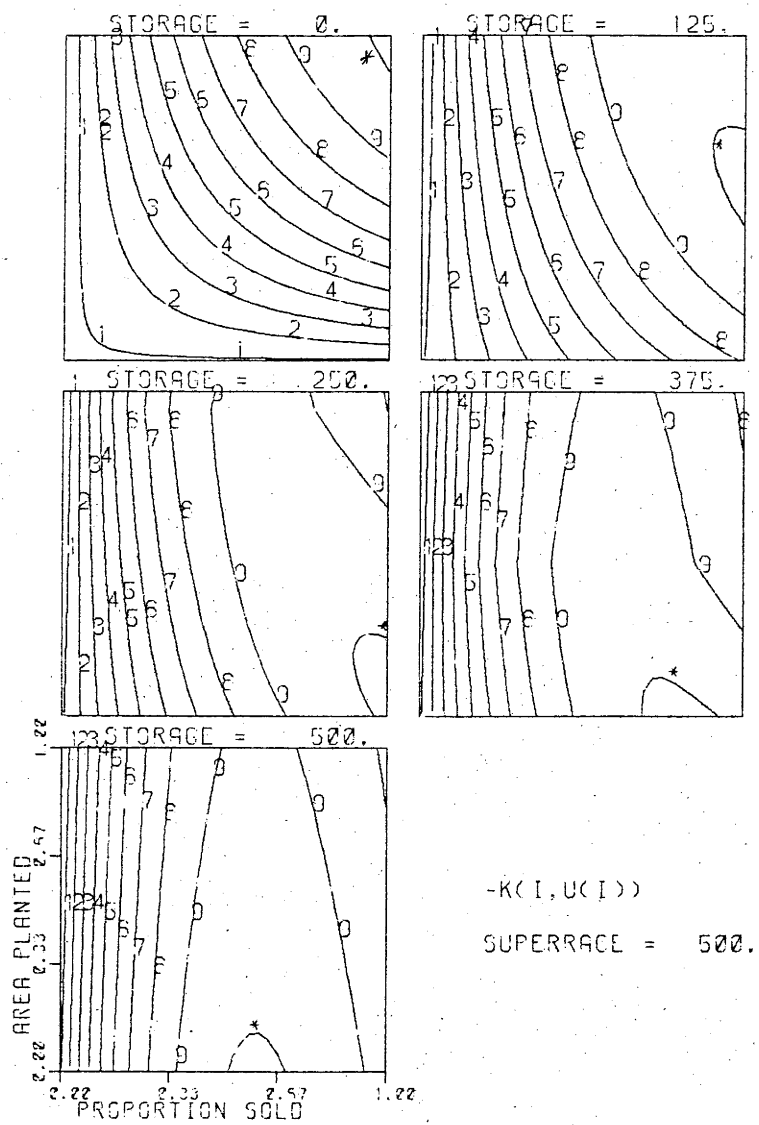

FIG. 6.4.49A: RETURN WITH GENERAL RESISTANCE

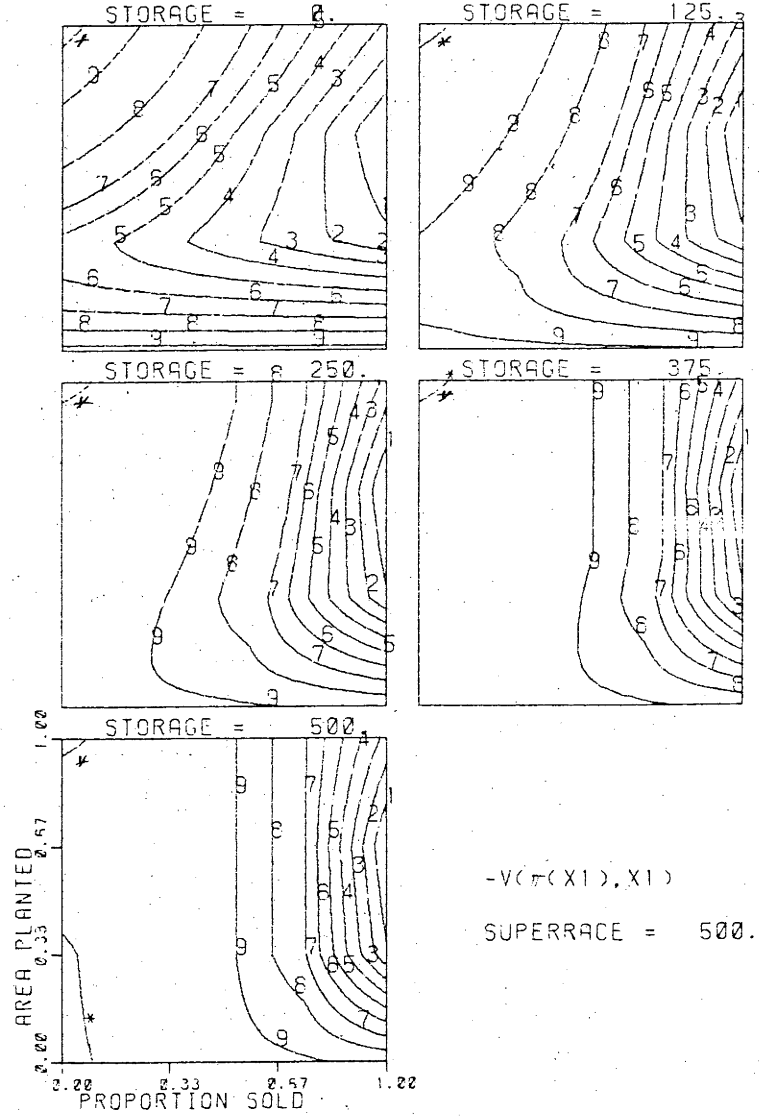

FIG. 6.4.49B: RETURN WITH GENERAL RESISTANCE

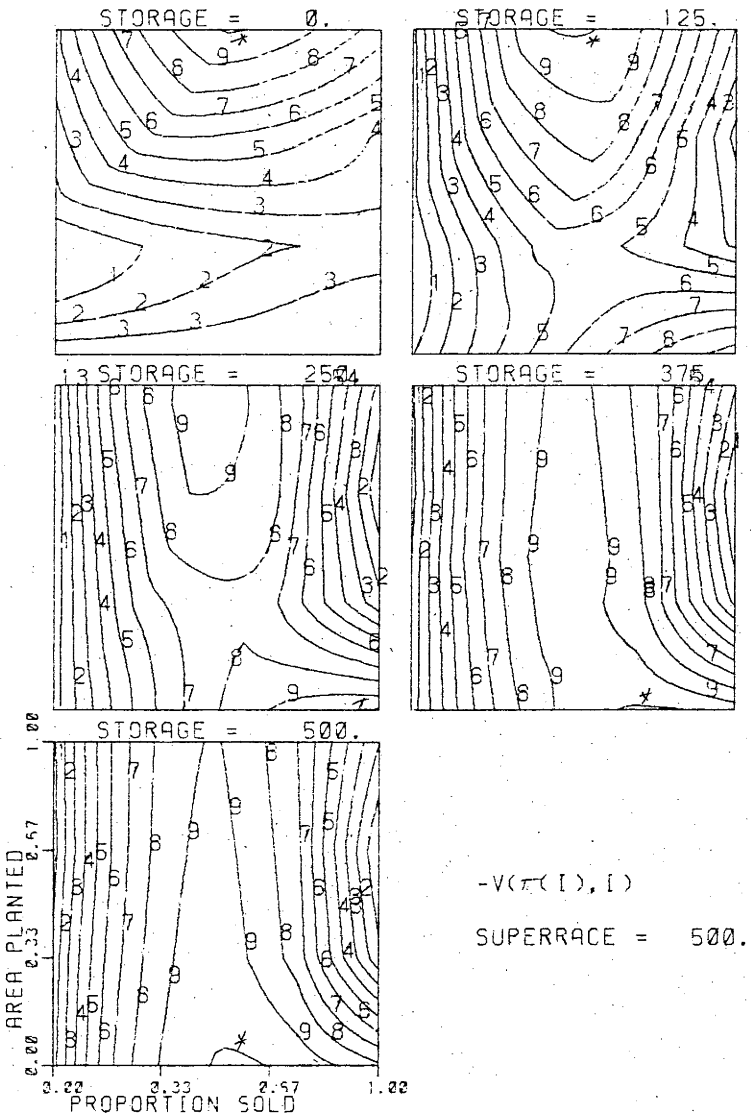

FIG. 6.4.49C: RETURN WITH GENERAL RESISTANCE

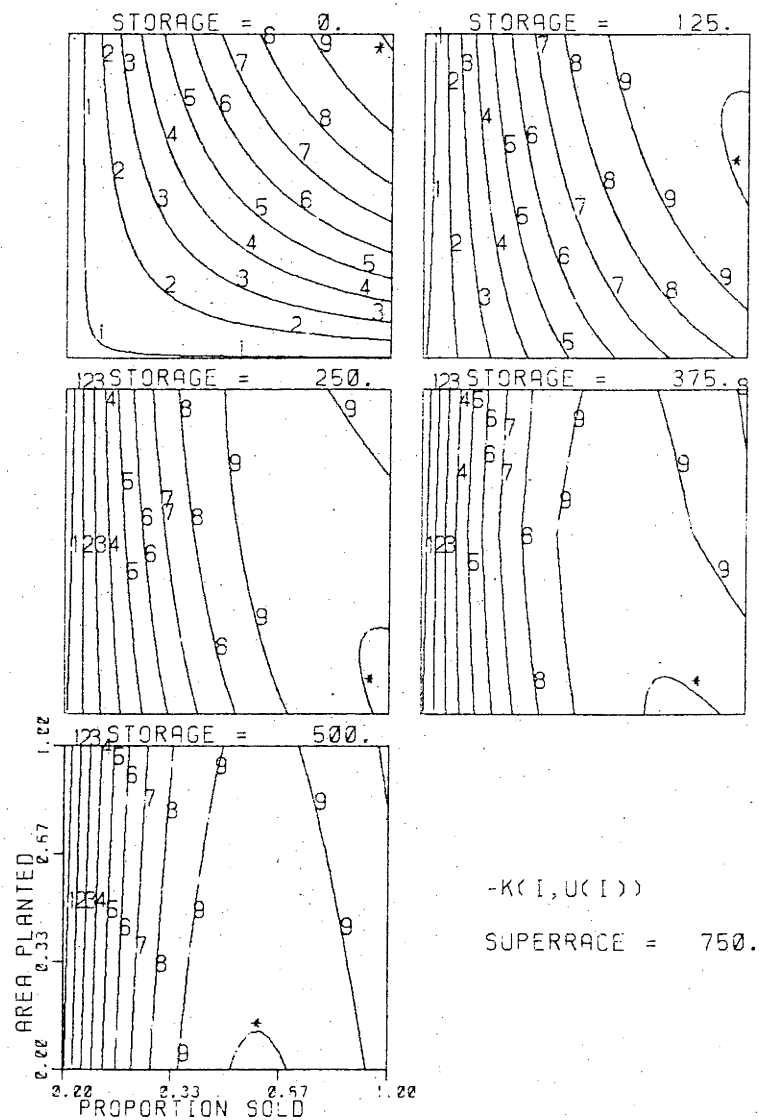

FIG. 6.4.50A: RETURN WITH GENERAL RESISTANCE

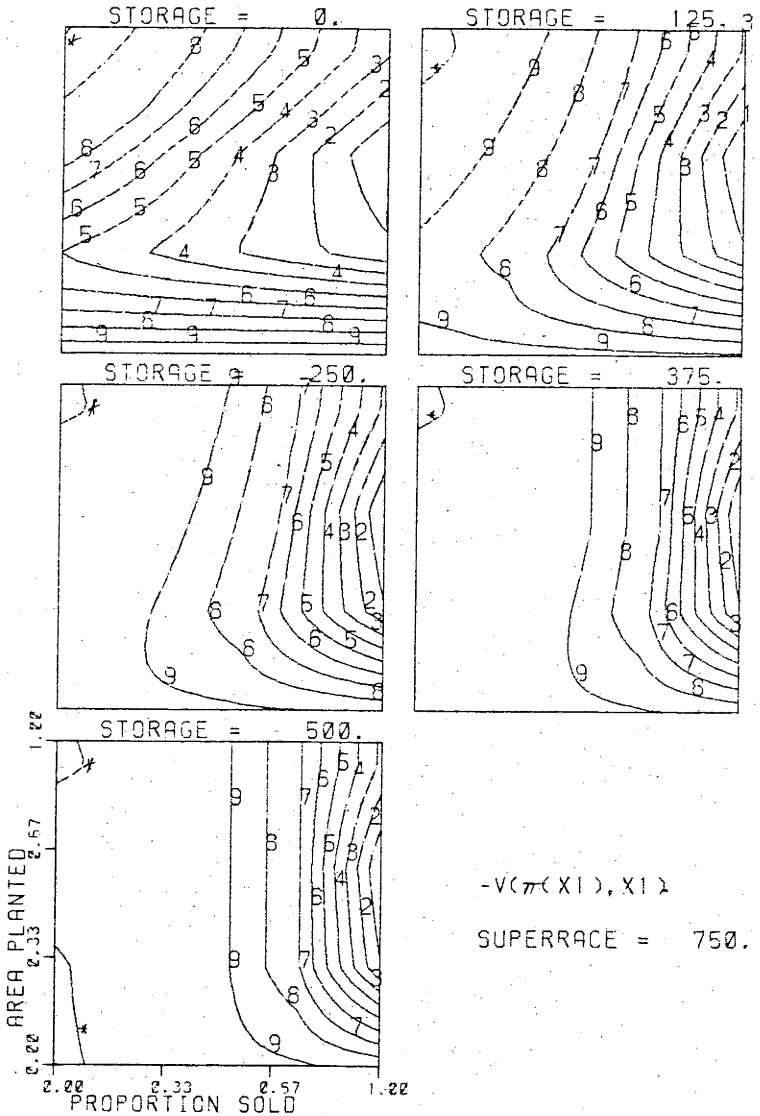

FIG. 6.4.50B· RETURN WITH GENERAL RESISTANCE

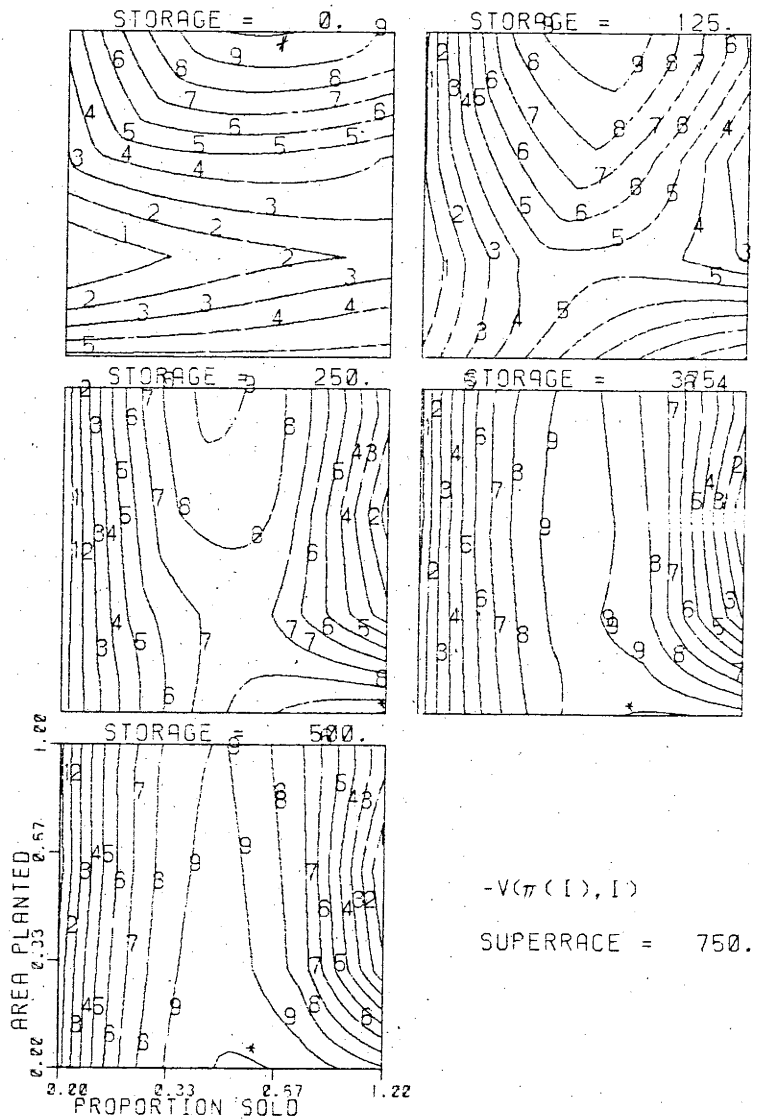

FIG. 6.4.50C·· RETURN WITH GENERAL RESISTANCE

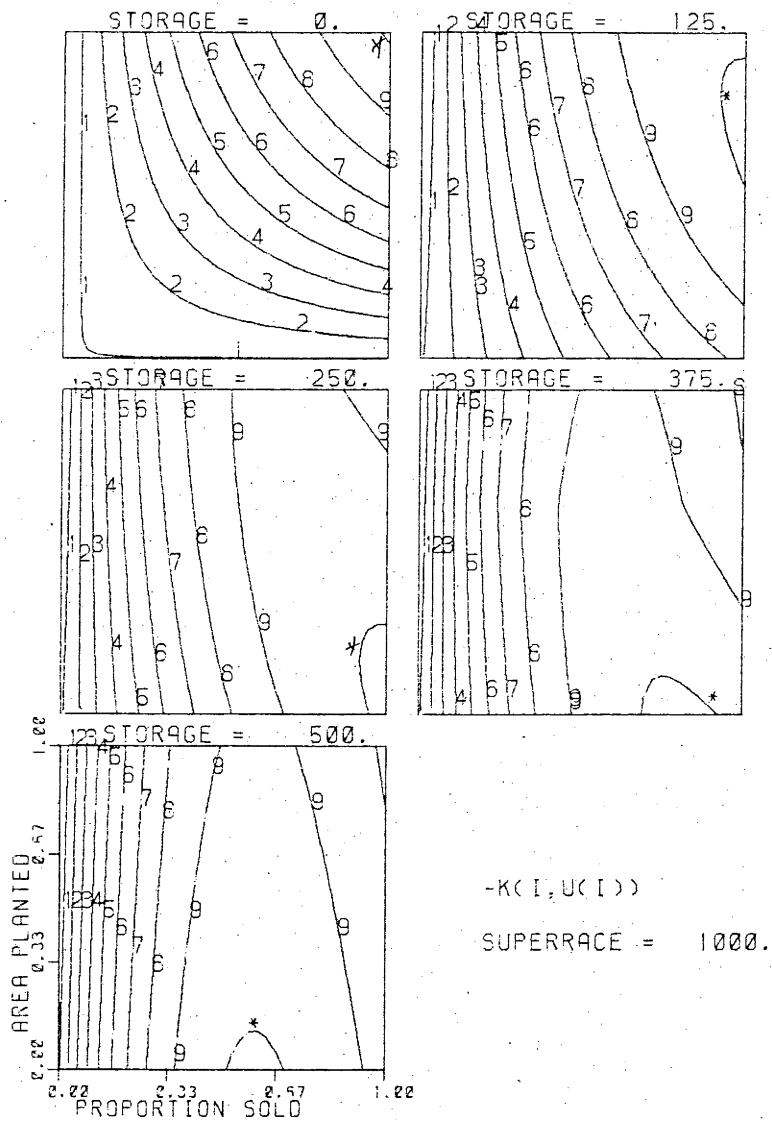

FIG. 6.4.51A: RETURN WITH GENERAL RESISTANCE

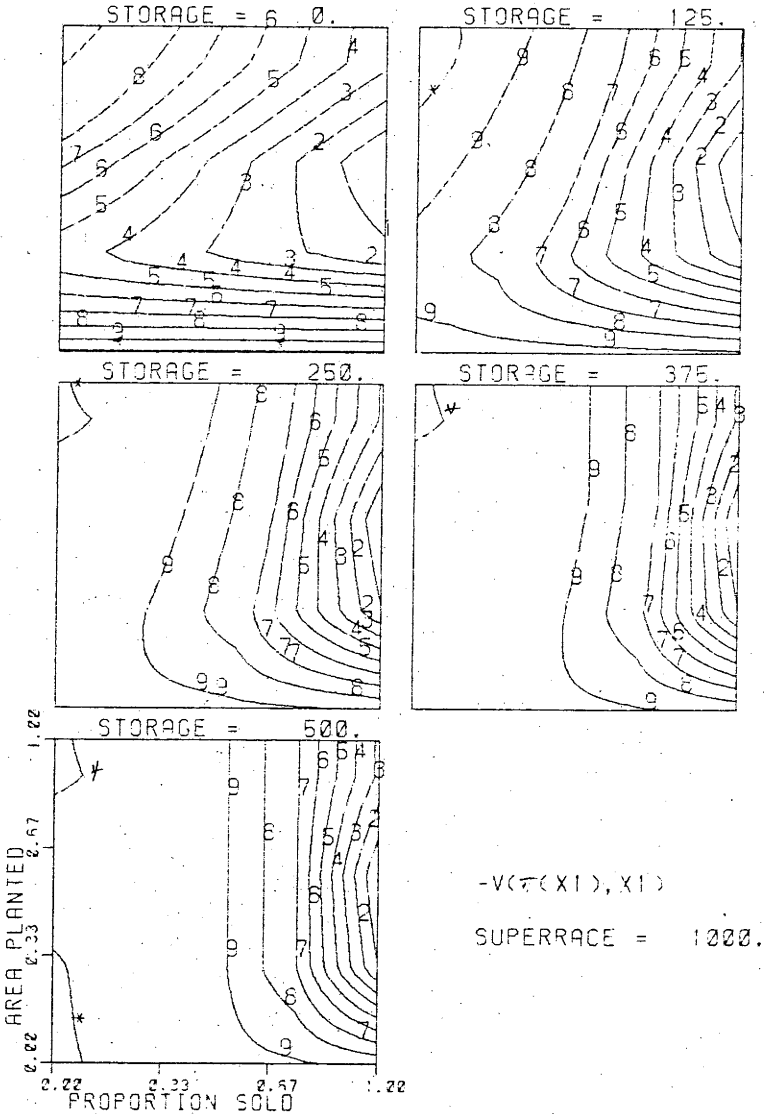

FIG. 6.4.51B: RETURN WITH GENERAL RESISTANCE

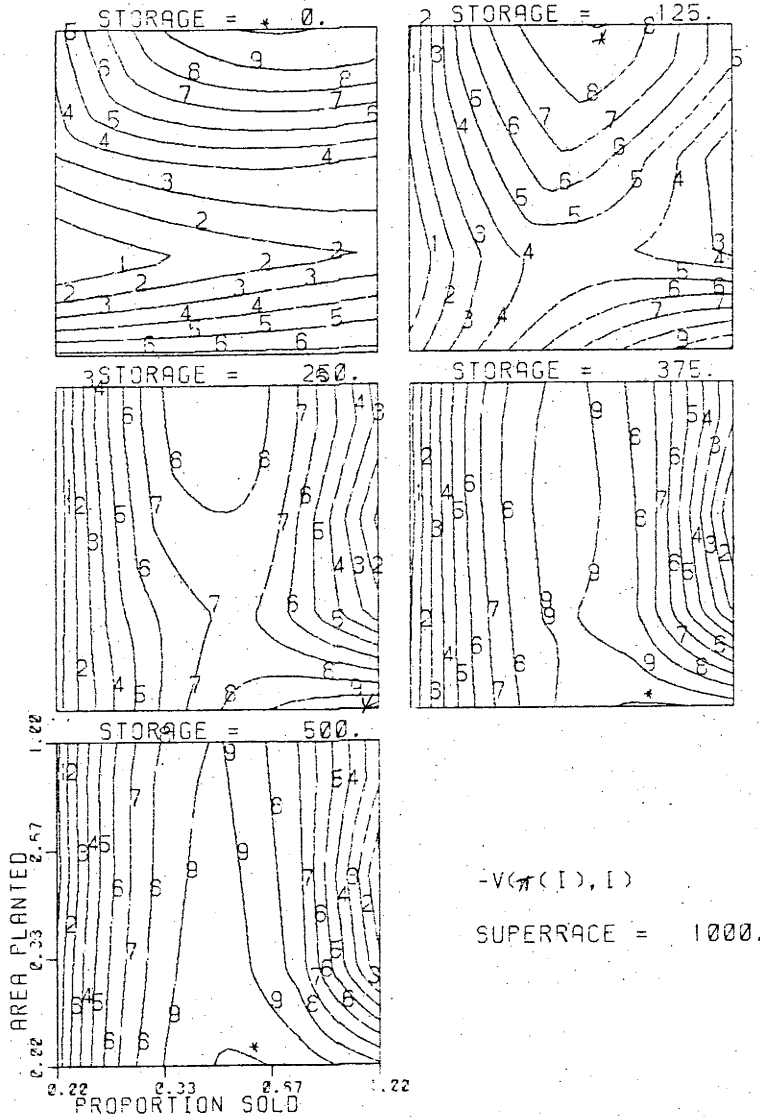

FIG. 6.4.51C. RETURN WITH GENERAL RESISTANCE

FIG 6.4.52: AGRIBUSINESS + GENERAL RESISTANCE
FRACTION SOLD THROUGH TIME

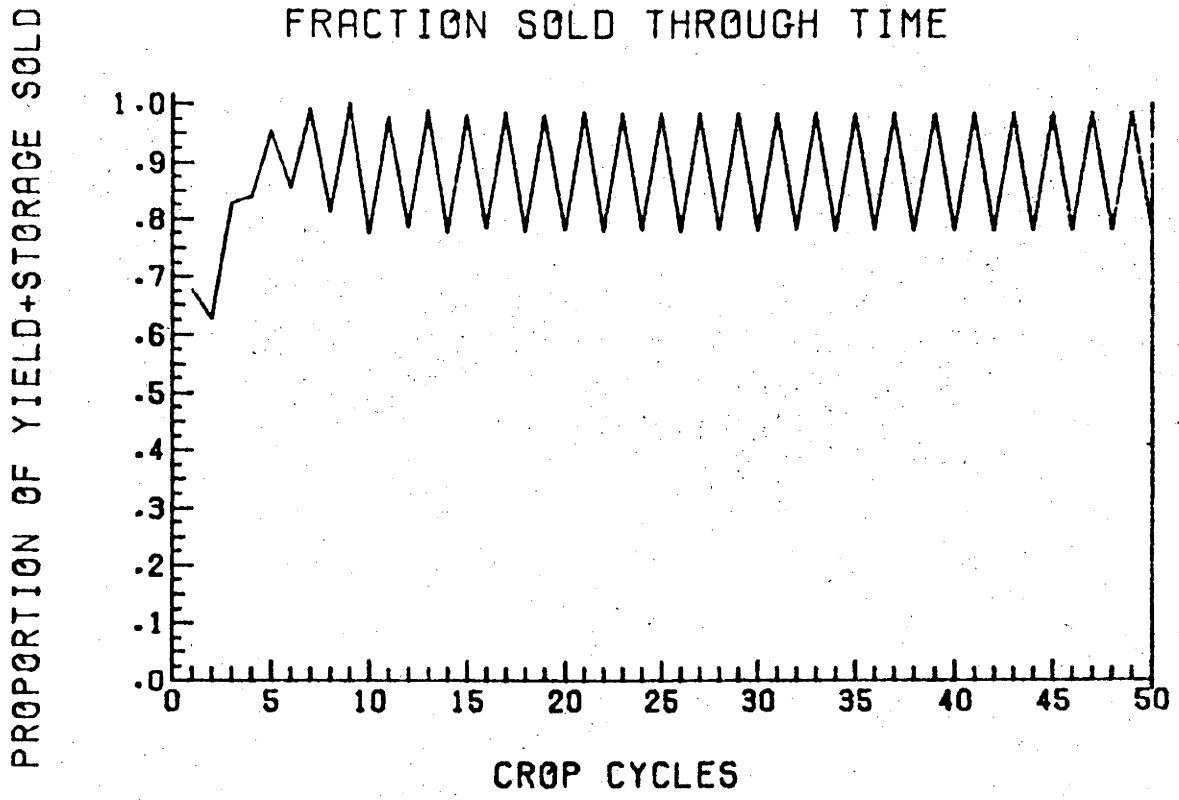

FIG 6.4.53: AGRIBUSINESS + GENERAL RESISTANCE
AREA PLANTED THROUGH TIME

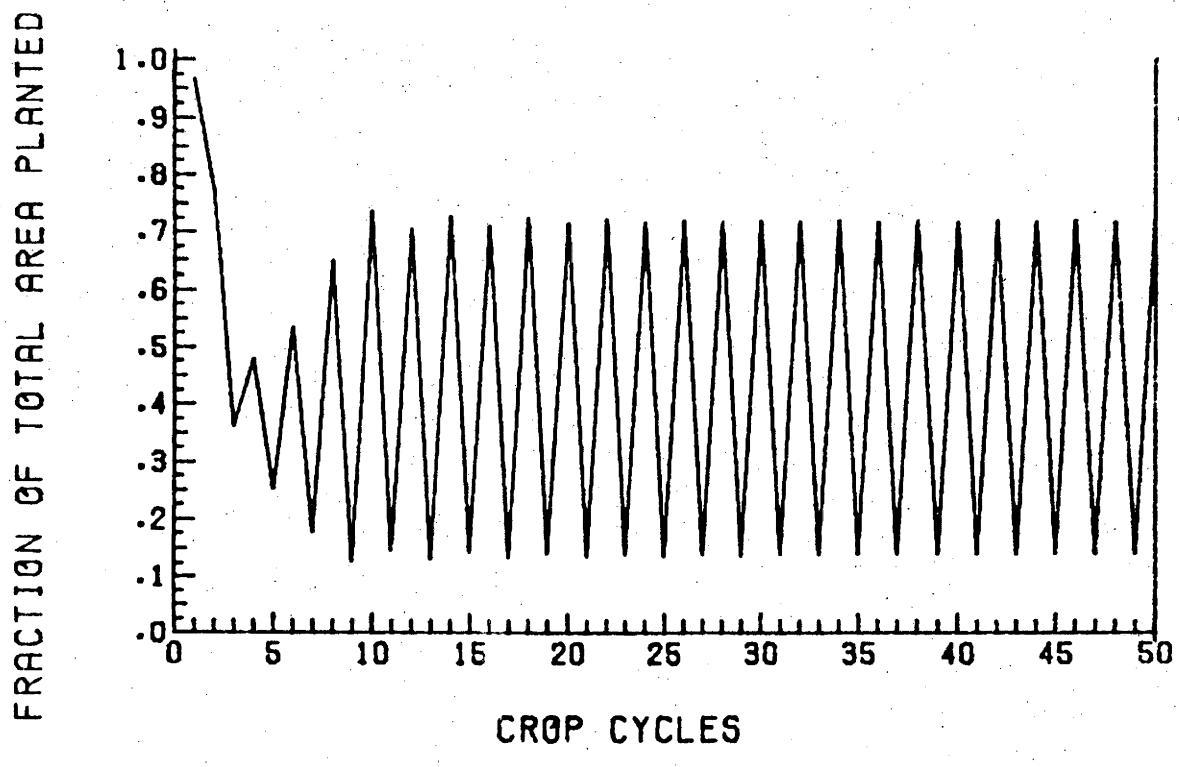

FIG 6.4.54: AGRIBUSINESS + GENERAL RESISTANCE
RETURN THROUGH TIME

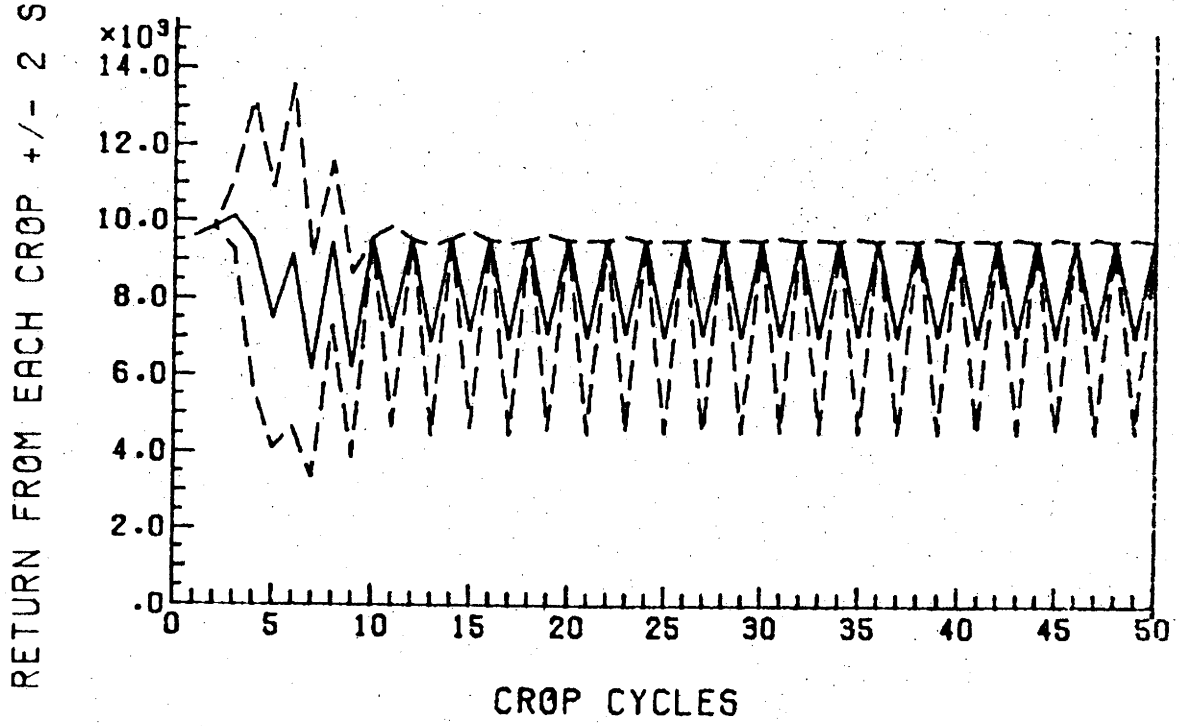

FIG 6.4.55: AGRIBUSINESS + GENERAL RESISTANCE
YIELD THROUGH TIME

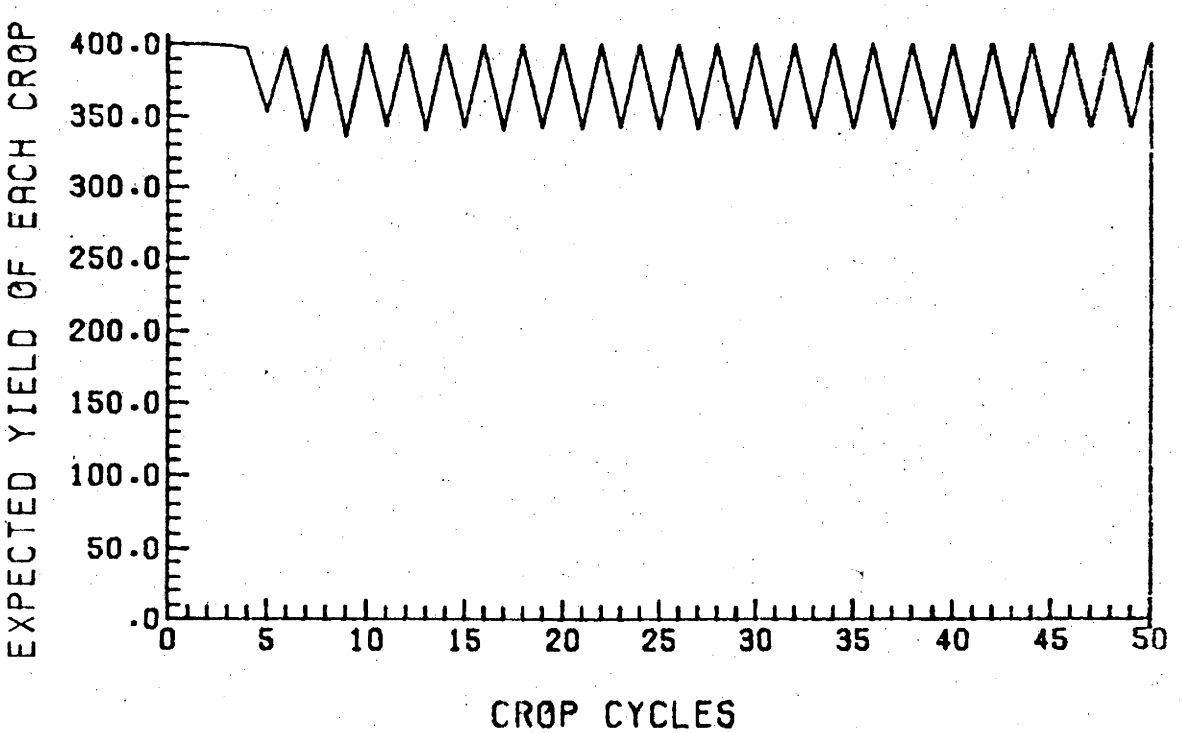

FIG 6.4.56: AGRIBUSINESS + GENERAL RESISTANCE STORAGE THROUGH TIME

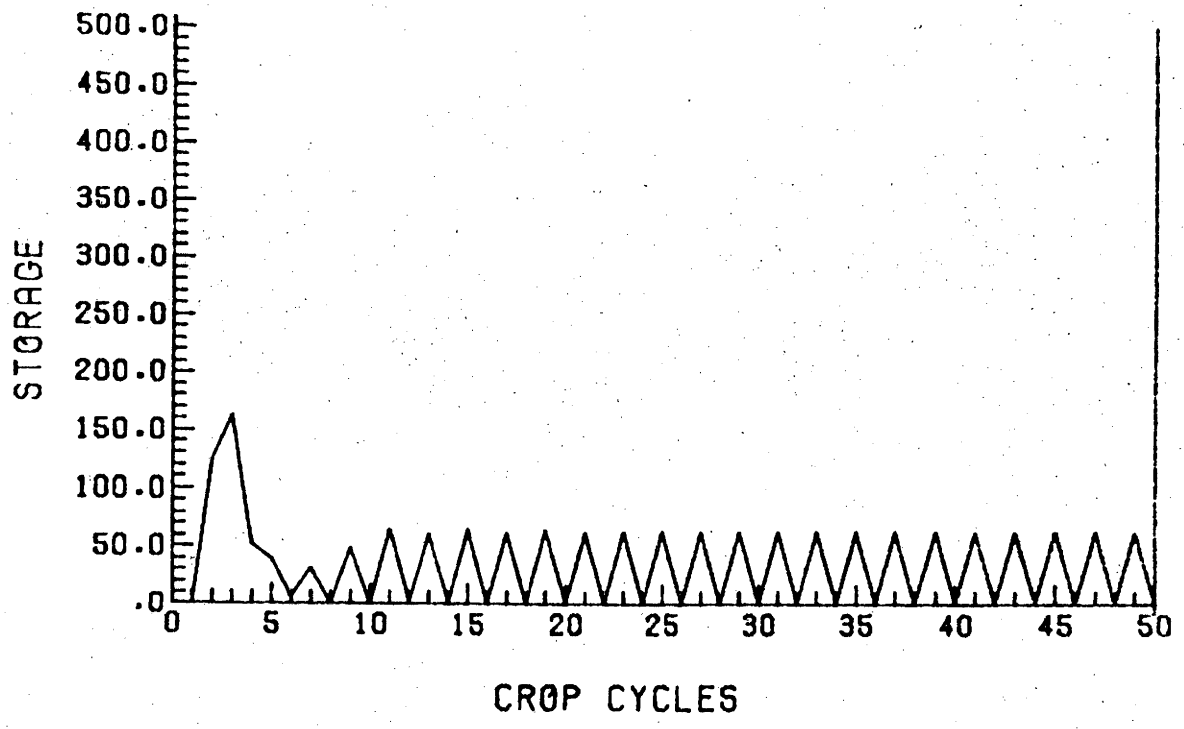

FIG 6.4.57: AGRIBUSINESS + GENERAL RESISTANCE SUPERRACE THROUGH TIME

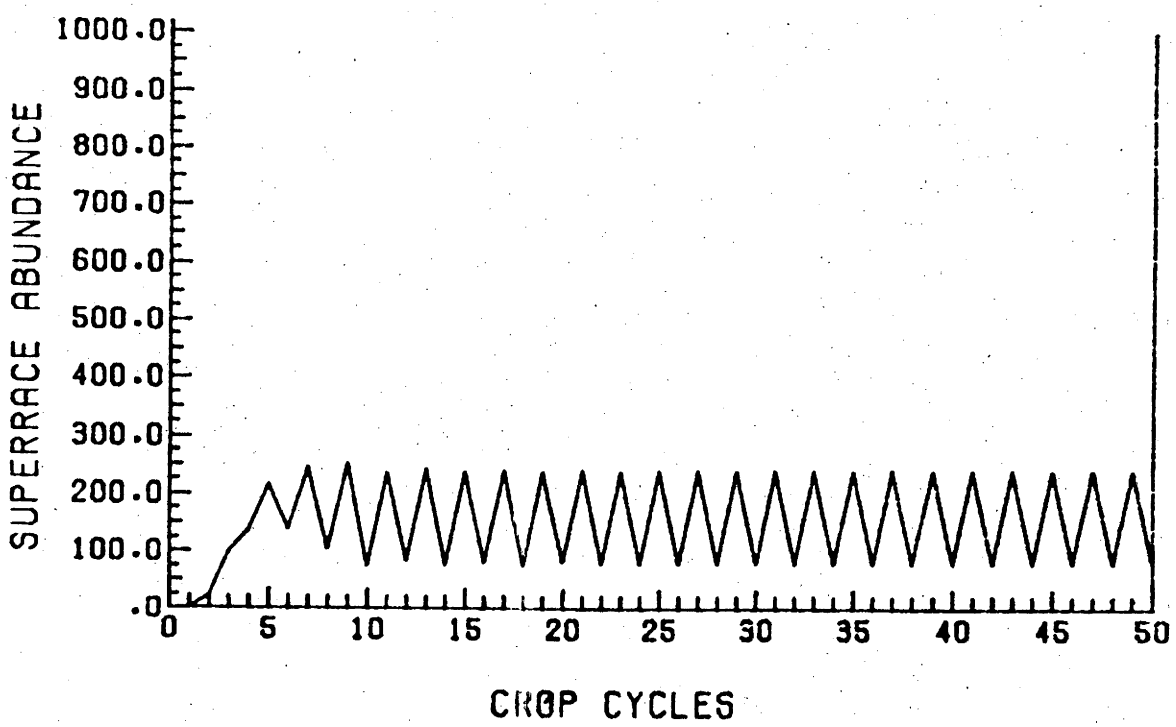

FIG. 6.4.57A: COMPONENT AND TOTAL YIELD UNDER CROSSPROTECTION

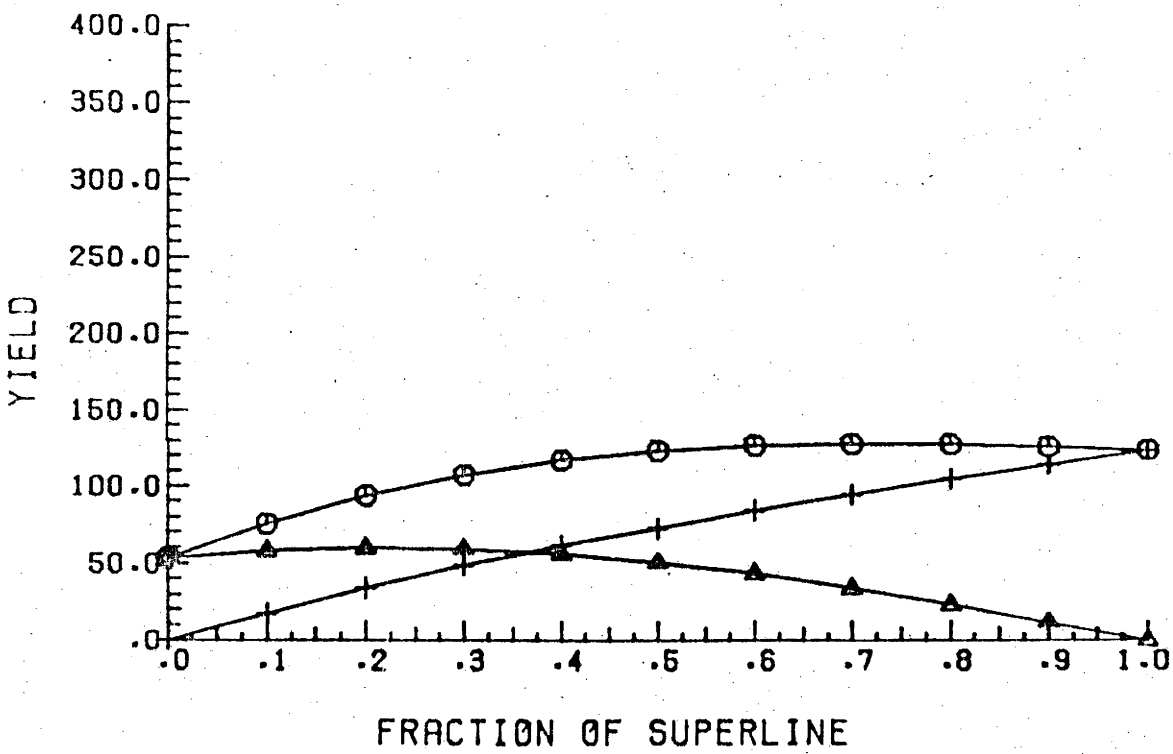

FIG 6.4.58A: AGRIBUSINESS + CROSSPROTECTION
FRACTION SOLD VS STORAGE

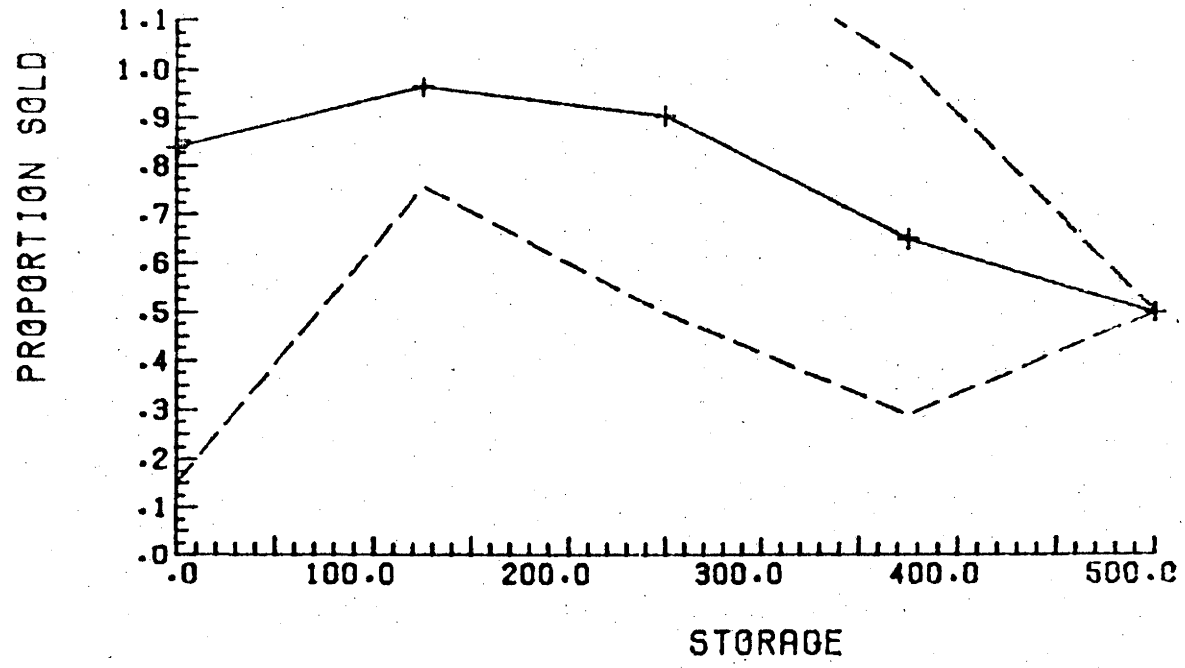

FIG 6.4.58B: AGRIBUSINESS + CROSSPROTECTION
FRACTION SOLD VS SUPERRACE

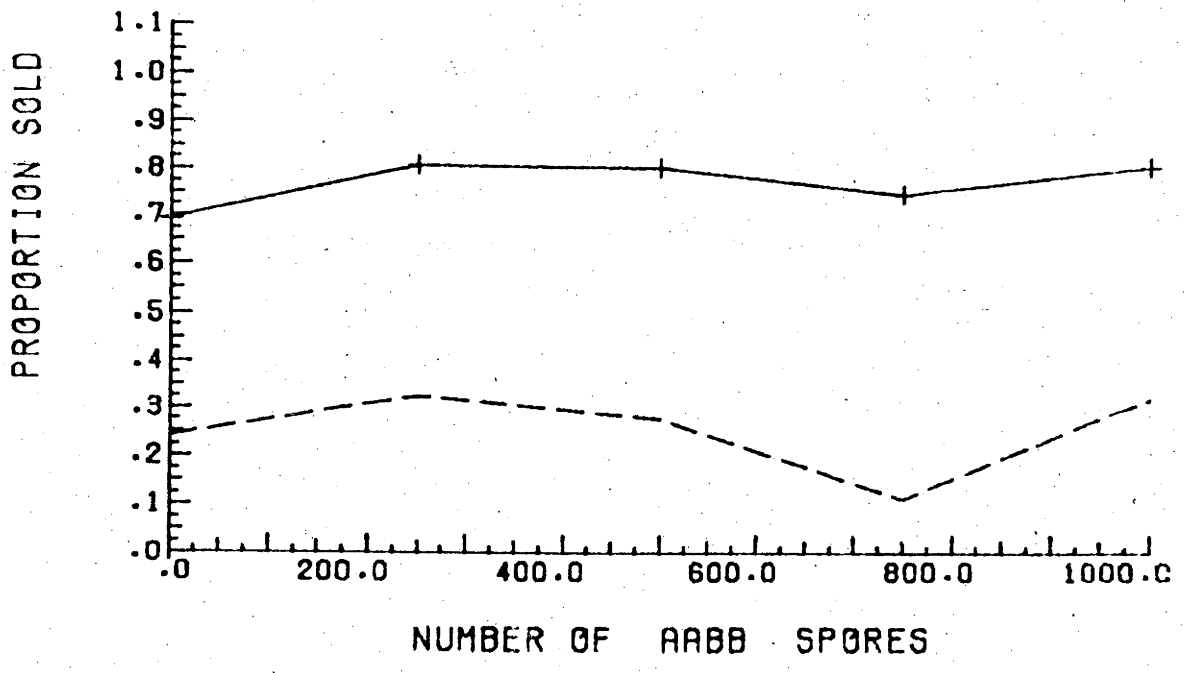

FIG 6.4.59A: AGRIBUSINESS + CROSSPROTECTION
AREA PLANTED VS STORAGE

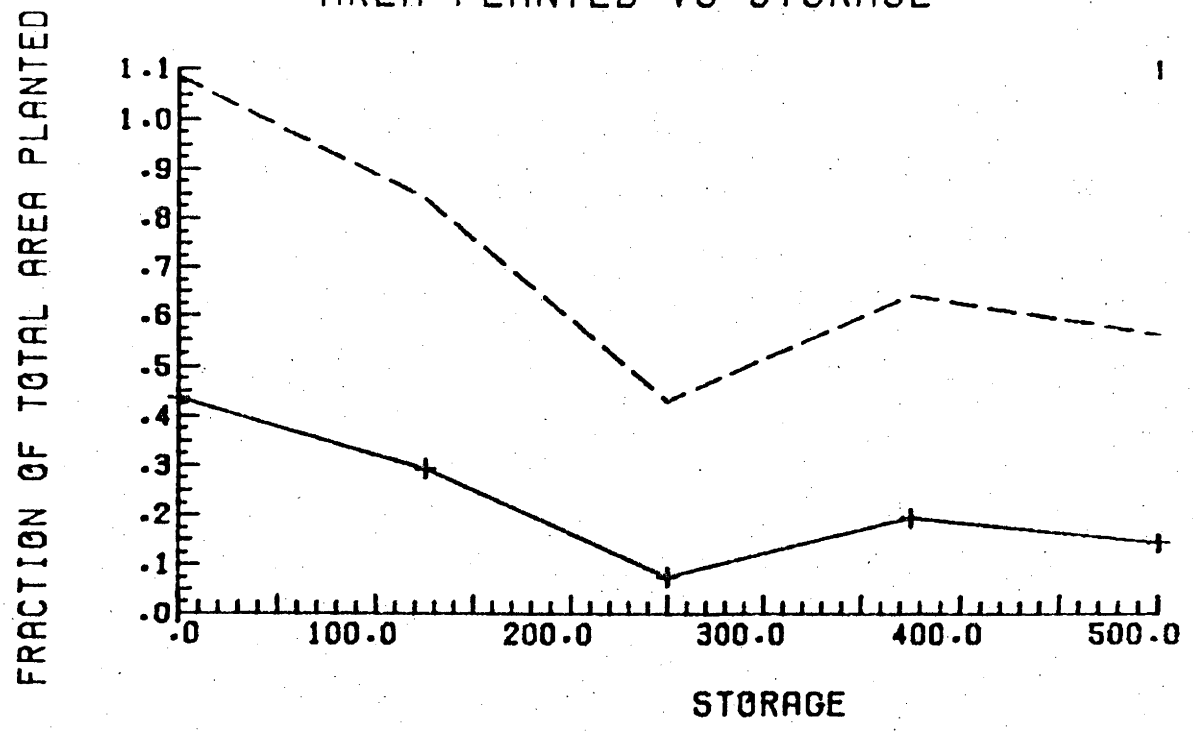

FIG 6.4.59B: AGRIBUSINESS + CROSSPROTECTION
AREA PLANTED VS SUPERRACE

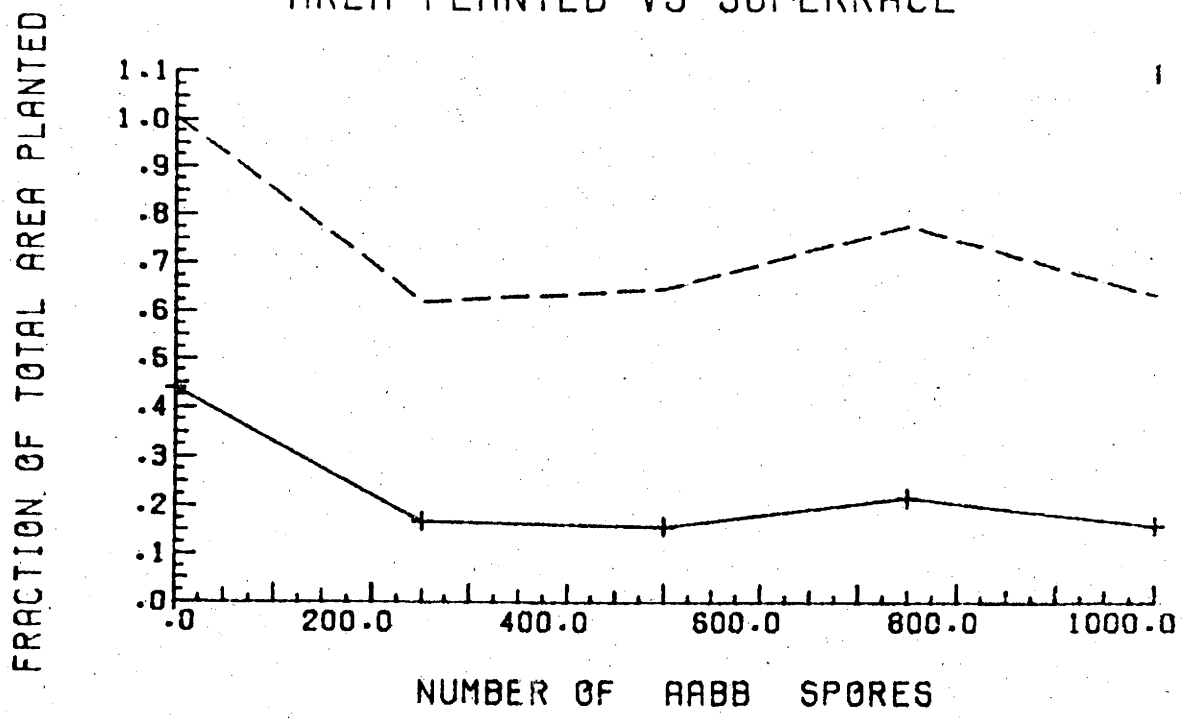

FIG 6.4.60A: AGRIBUSINESS + CROSSPROTECTION
E(RETURN) VS STORAGE

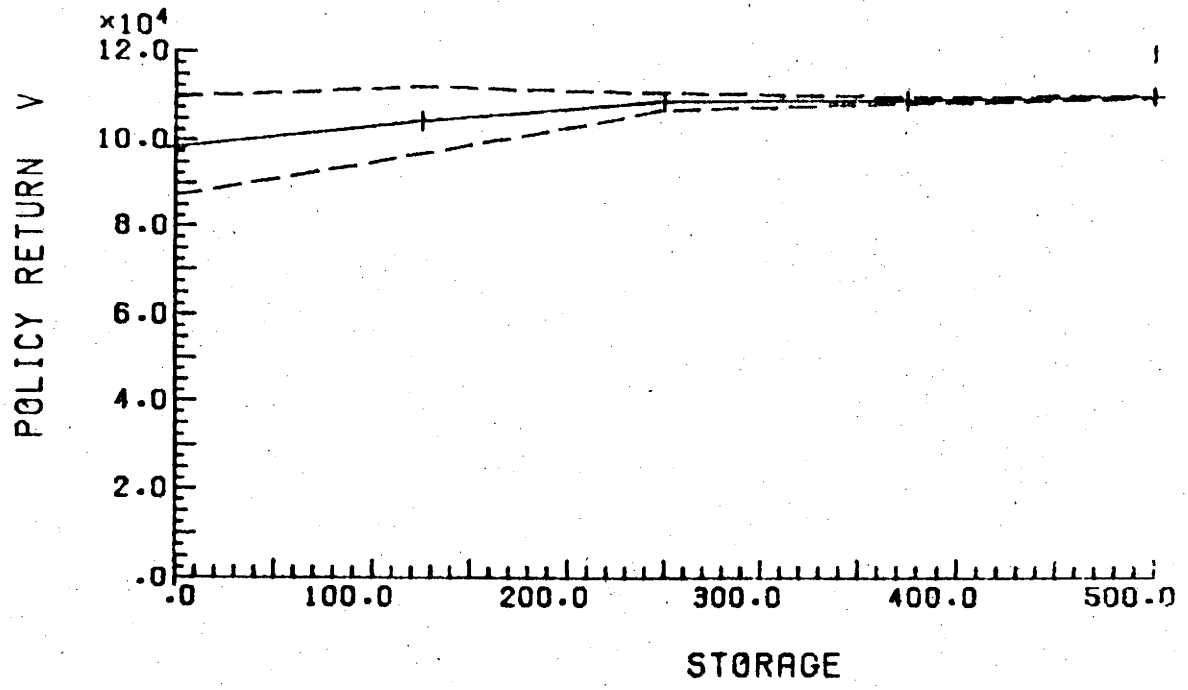

FIG 6.4.60B: AGRIBUSINESS + CROSSPROTECTION
E(RETURN) VS SUPERRACE

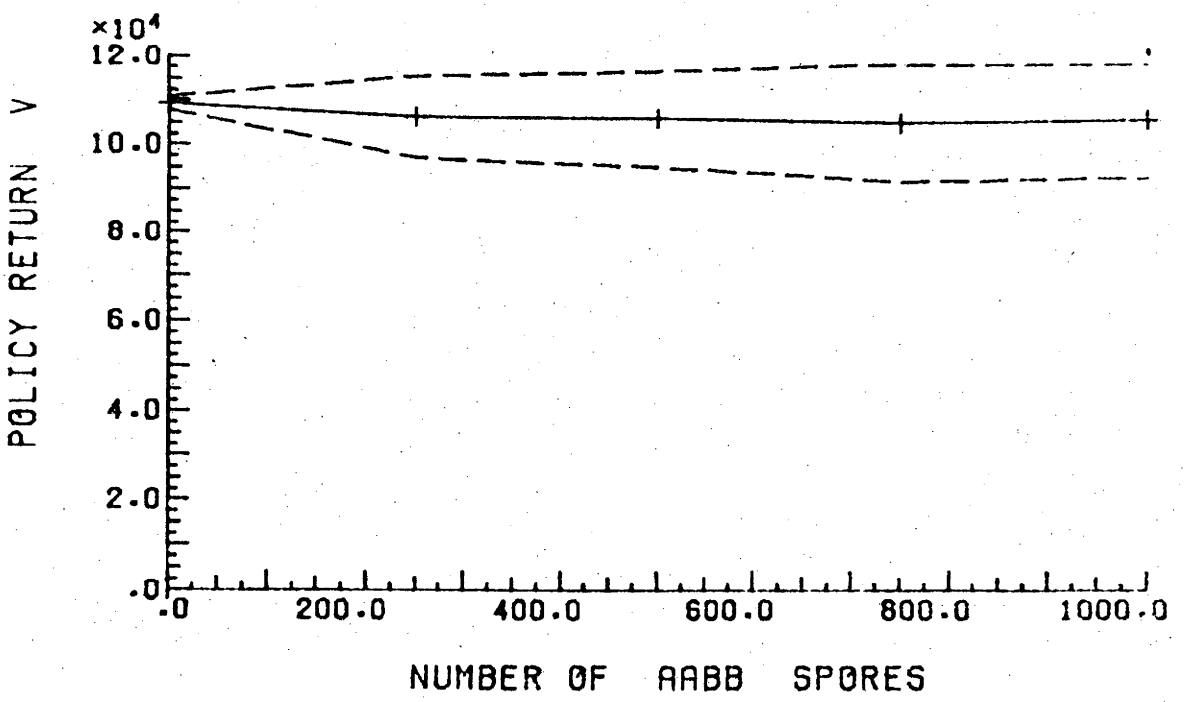

FIG 6.4.61A: AGRIBUSINESS + CROSSPROTECTION
STORAGE LEFT VS STORAGE

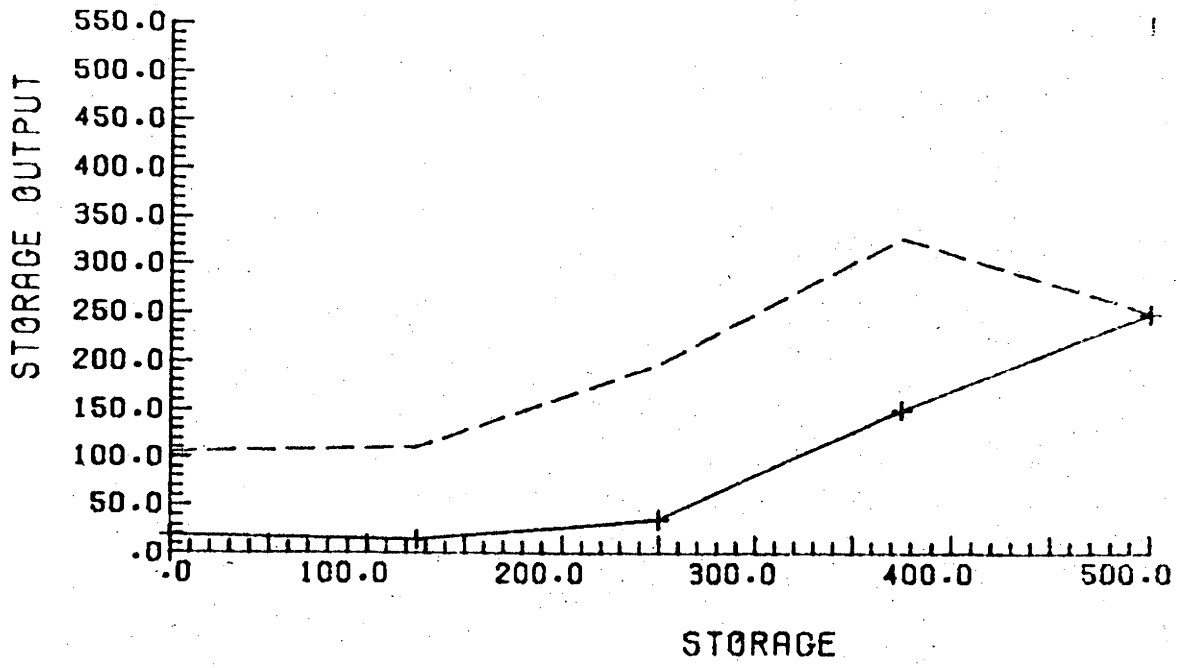

FIG 6.4.61B: AGRIBUSINESS + CROSSPROTECTION
STORAGE LEFT VS SUPERRACE

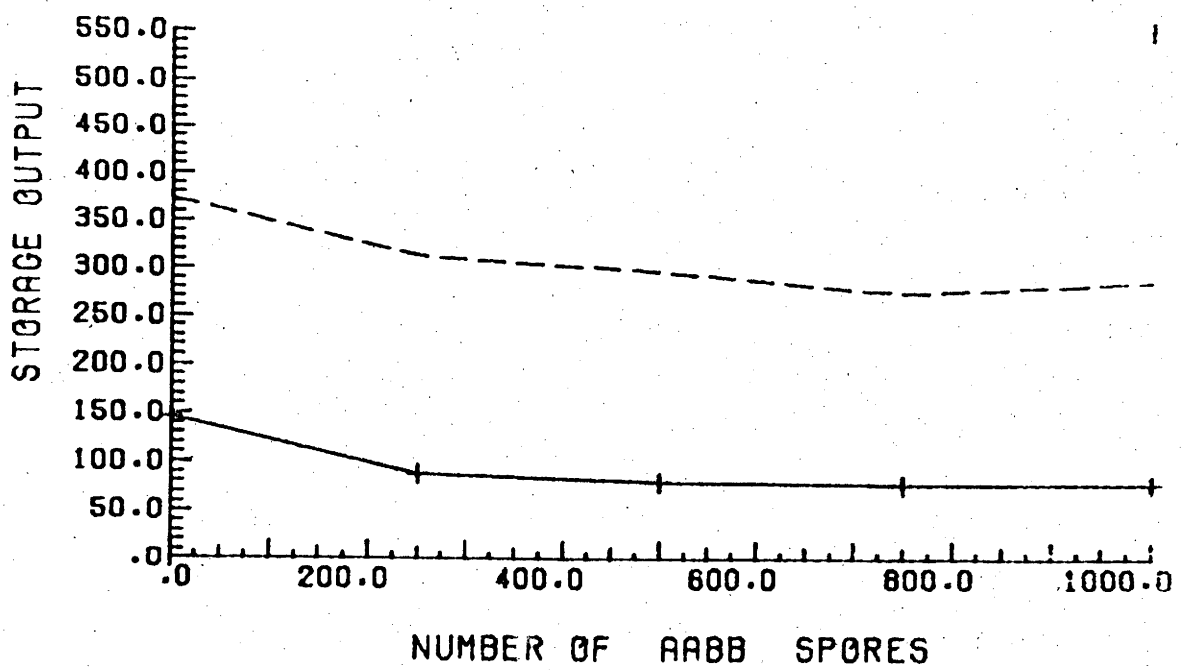

FIG 6.4.62A: AGRIBUSINESS + CROSSPROTECTION
AABB OUTPUT VS STORAGE

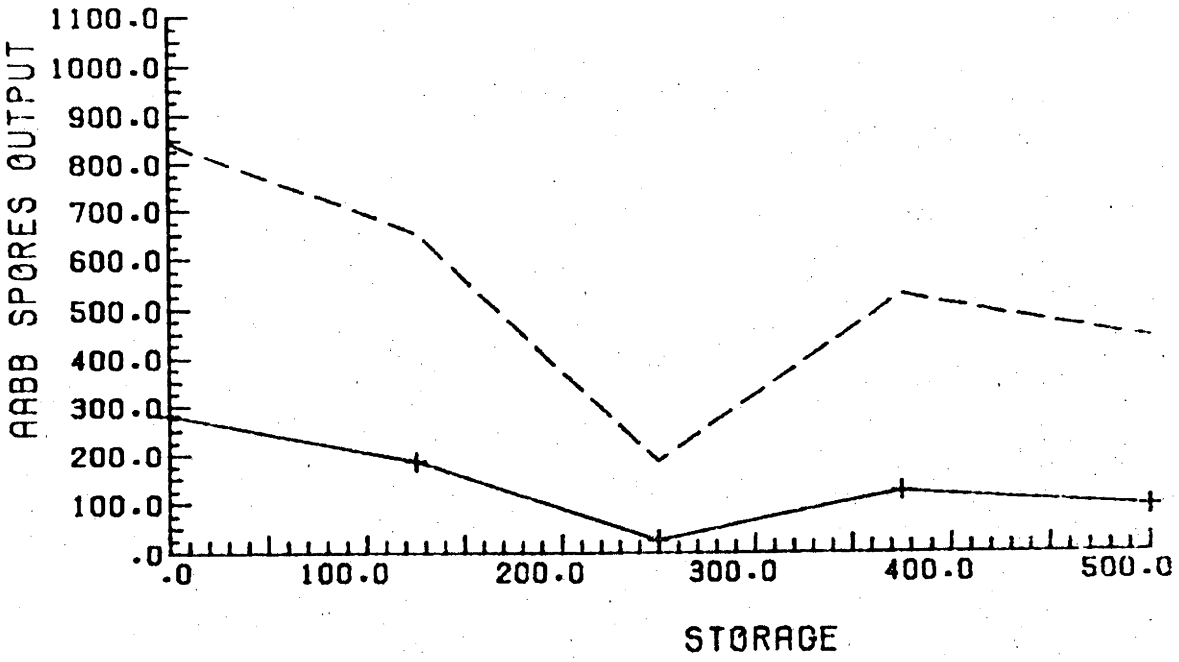

FIG 6.4.62B: AGRIBUSINESS + CROSSPROTECTION
AABB OUTPUT VS SUPERRACE

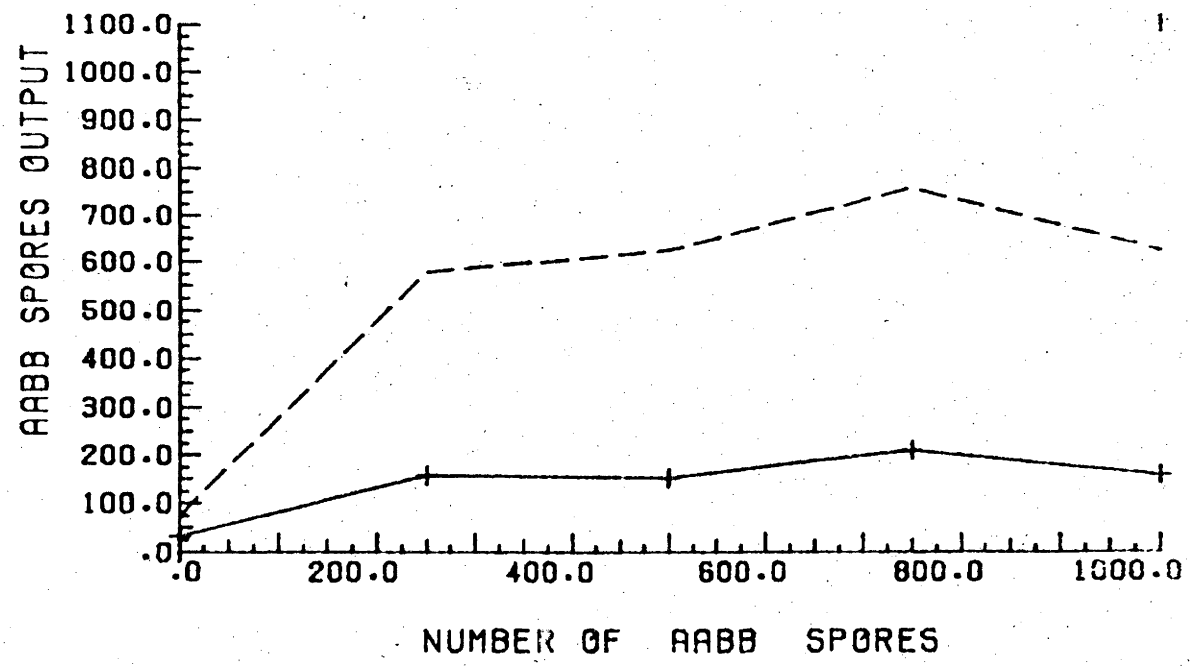

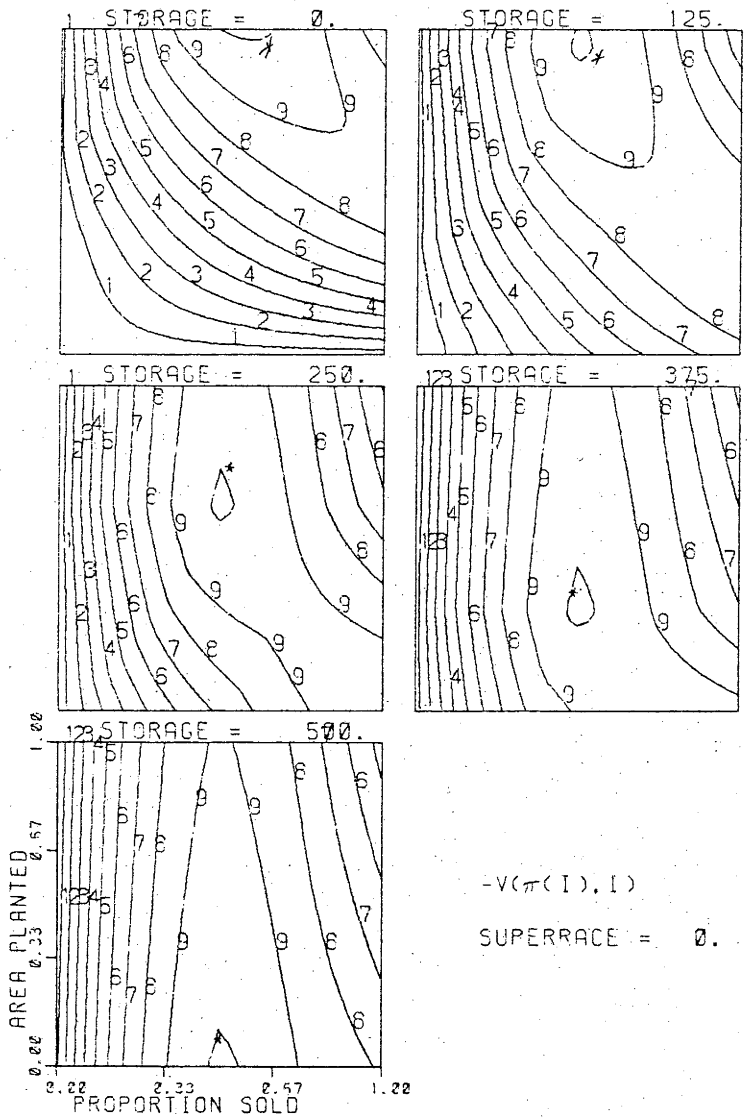

FIG. 6.4.63: RETURN FOR AGRIBUSINESS + CROSSPROTECTION

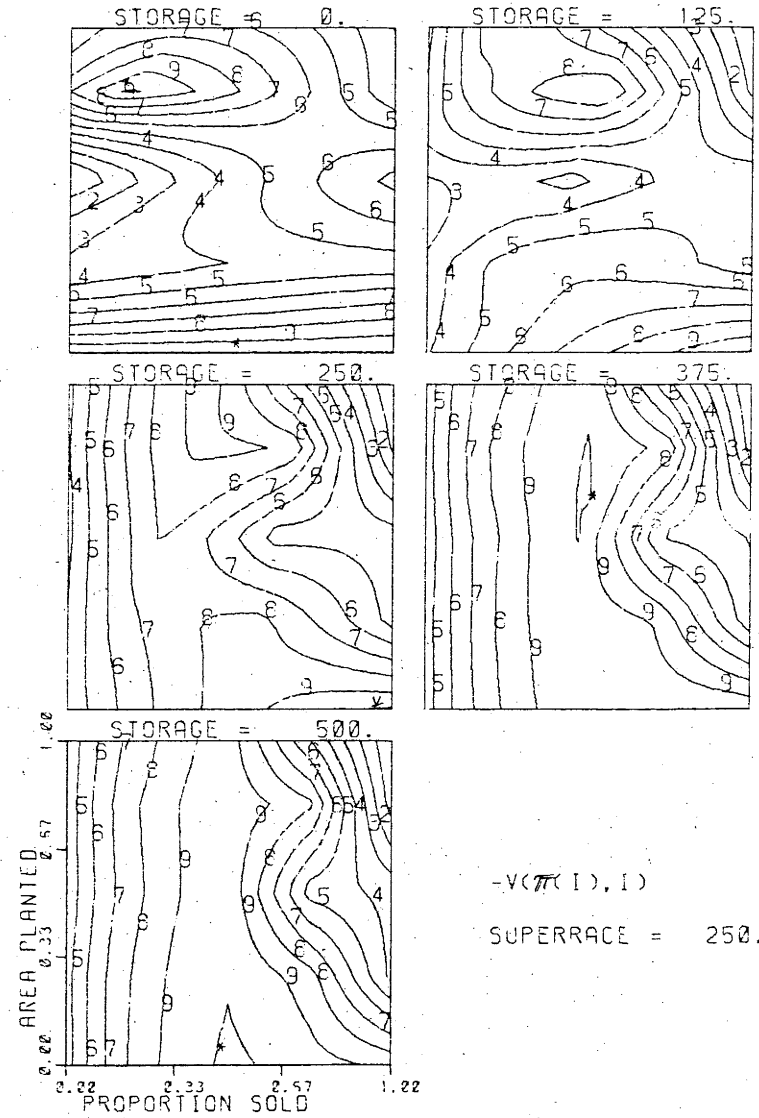

FIG. 6.4.64: RETURN FOR AGRIBUSINESS + CROSSPROTECTION

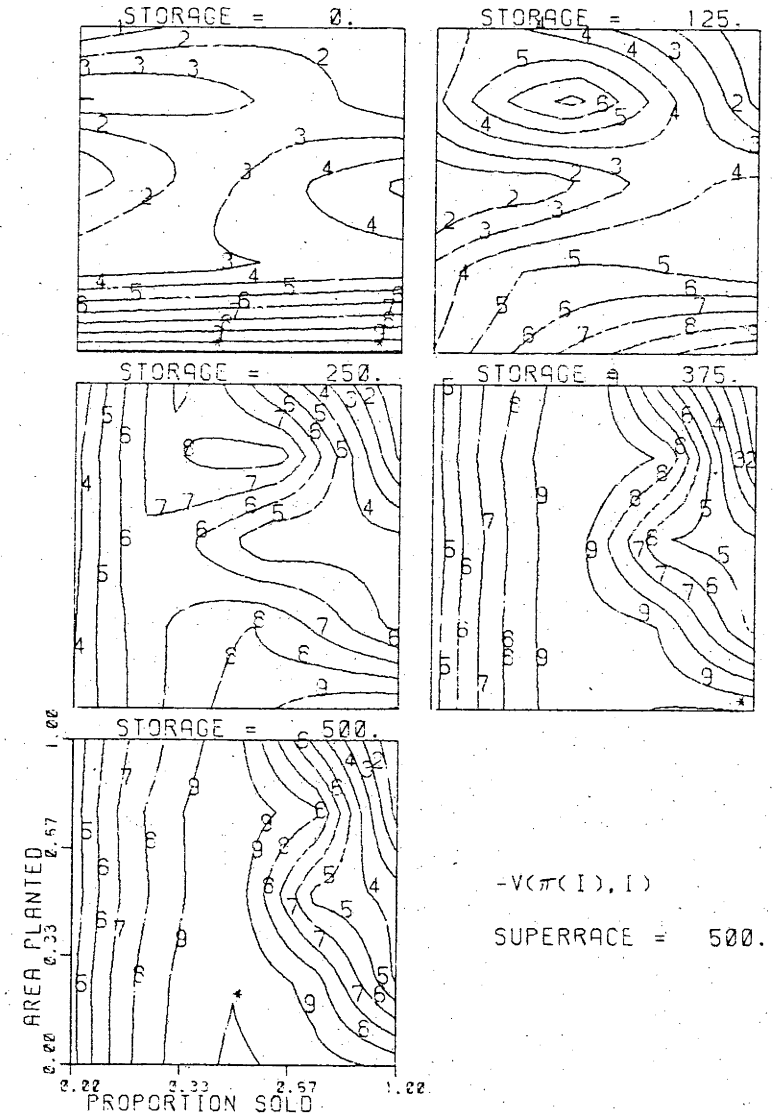

FIG. 6.4.65: RETURN FOR AGRIBUSINESS + CROSSPROTECTION

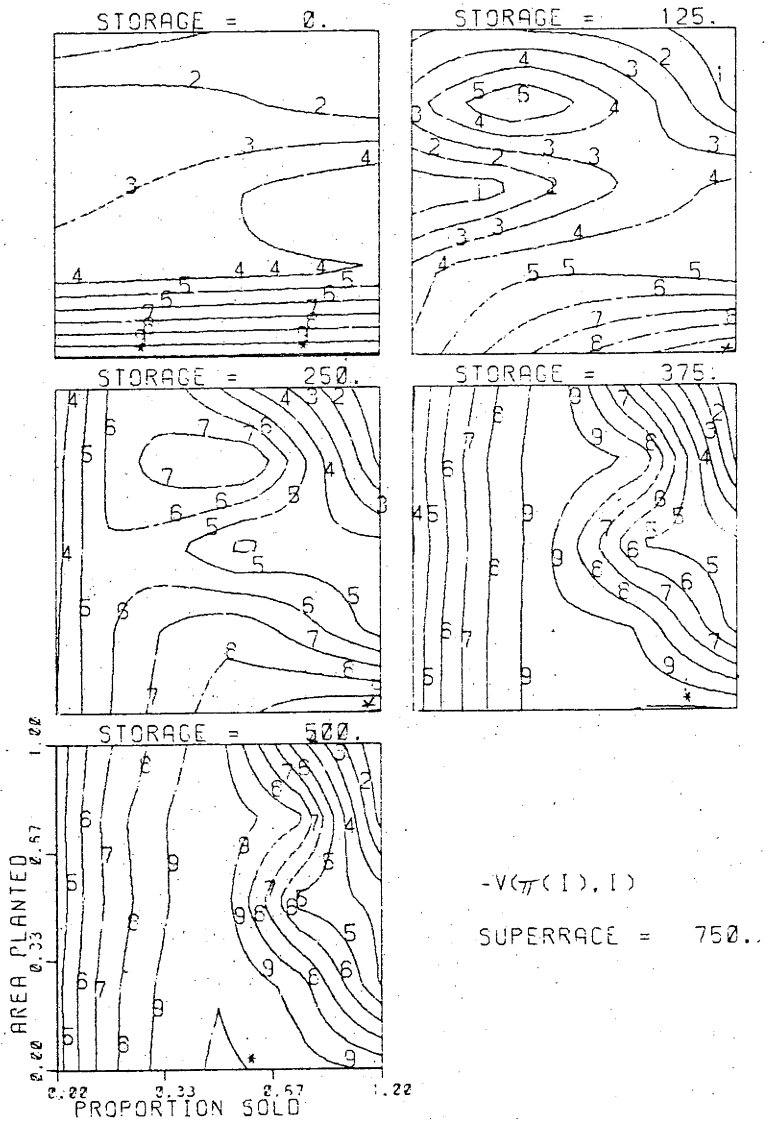

FIG. 6.4.66: RETURN FOR AGRIBUSINESS + CROSSPROTECTION

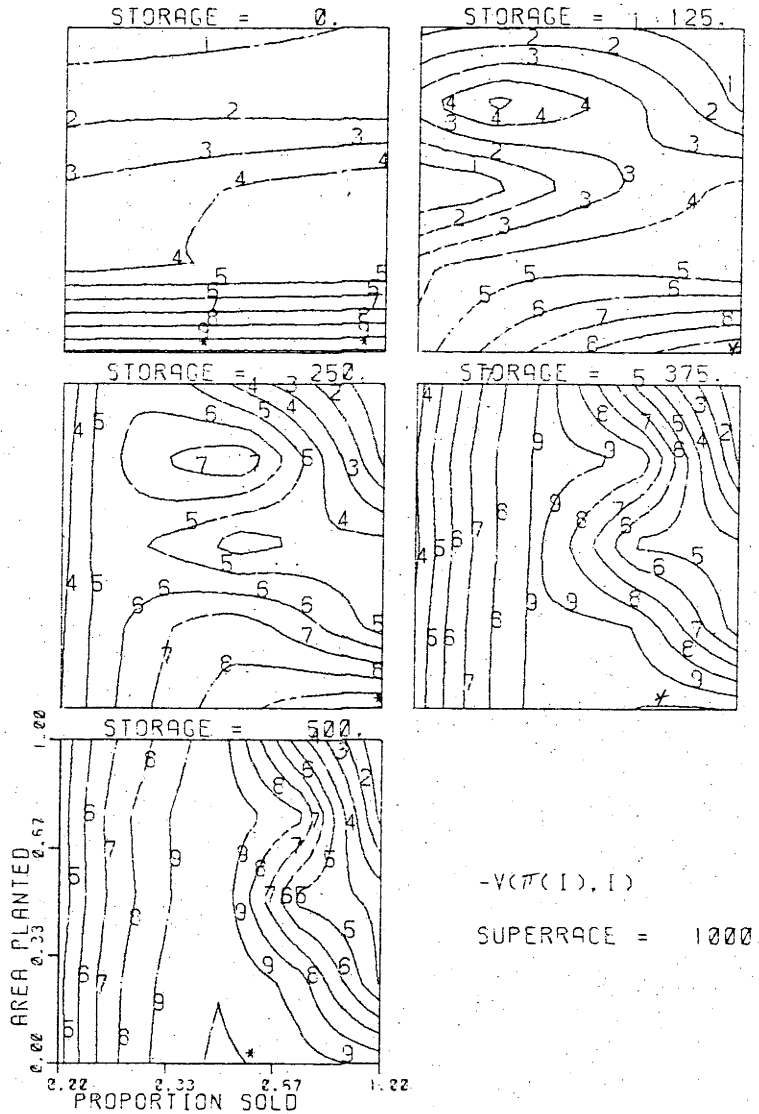

FIG. 6 4.67: RETURN FOR AGRIBUSINESS + CROSSPROTECTION

FIG 6.4.68: AGRIBUSINESS + CROSSPROTECTION
FRACTION SOLD THROUGH TIME

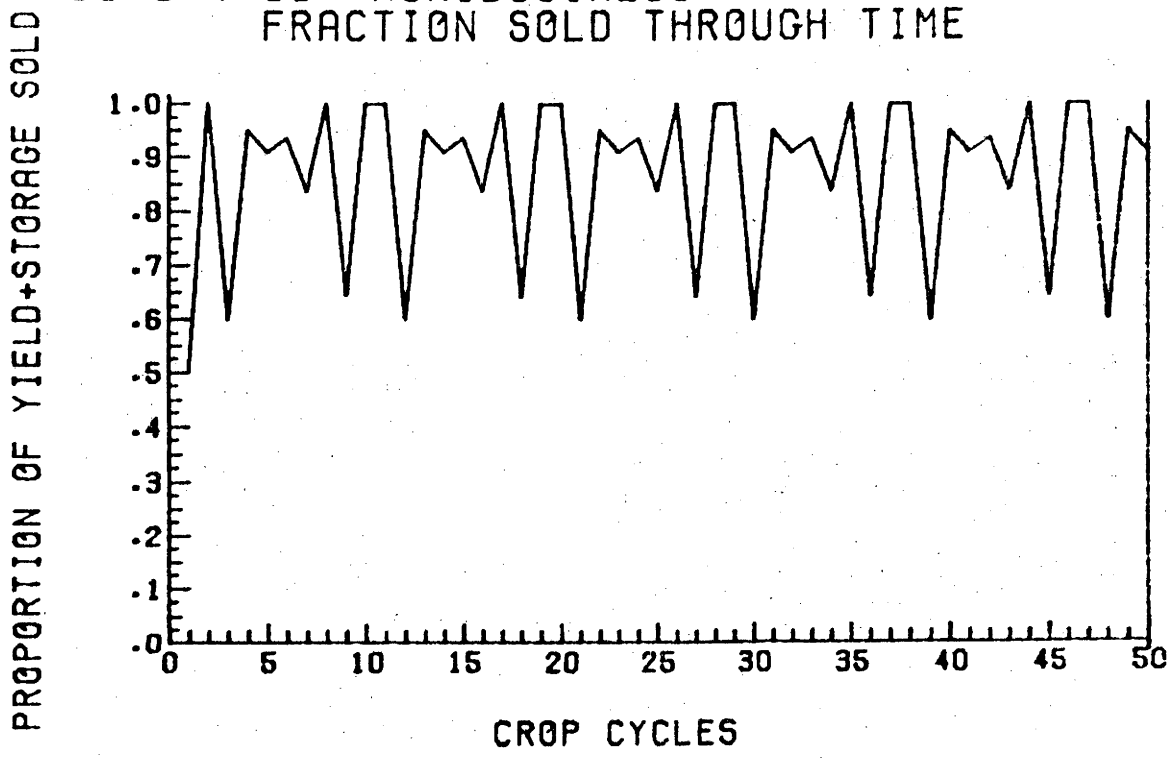

FIG 6.4.69: AGRIBUSINESS + CROSSPROTECTION
AREA PLANTED THROUGH TIME

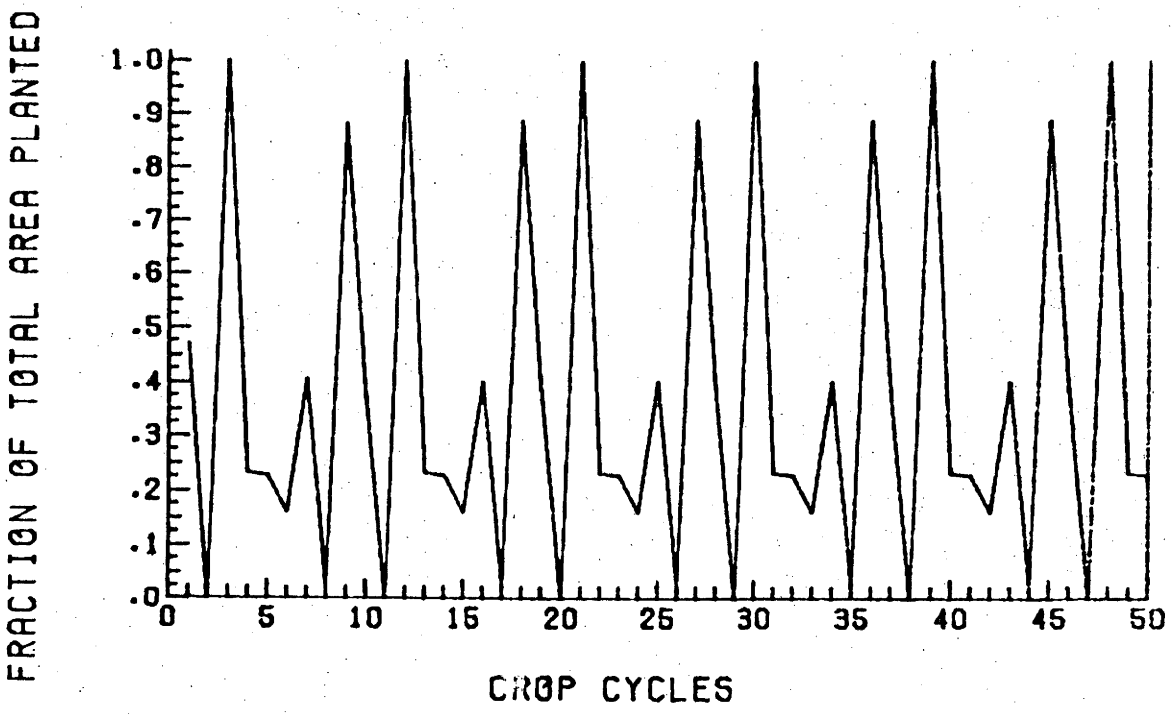

FIG 6.4.70: AGRIBUSINESS + CROSSPROTECTION
RETURN THROUGH TIME

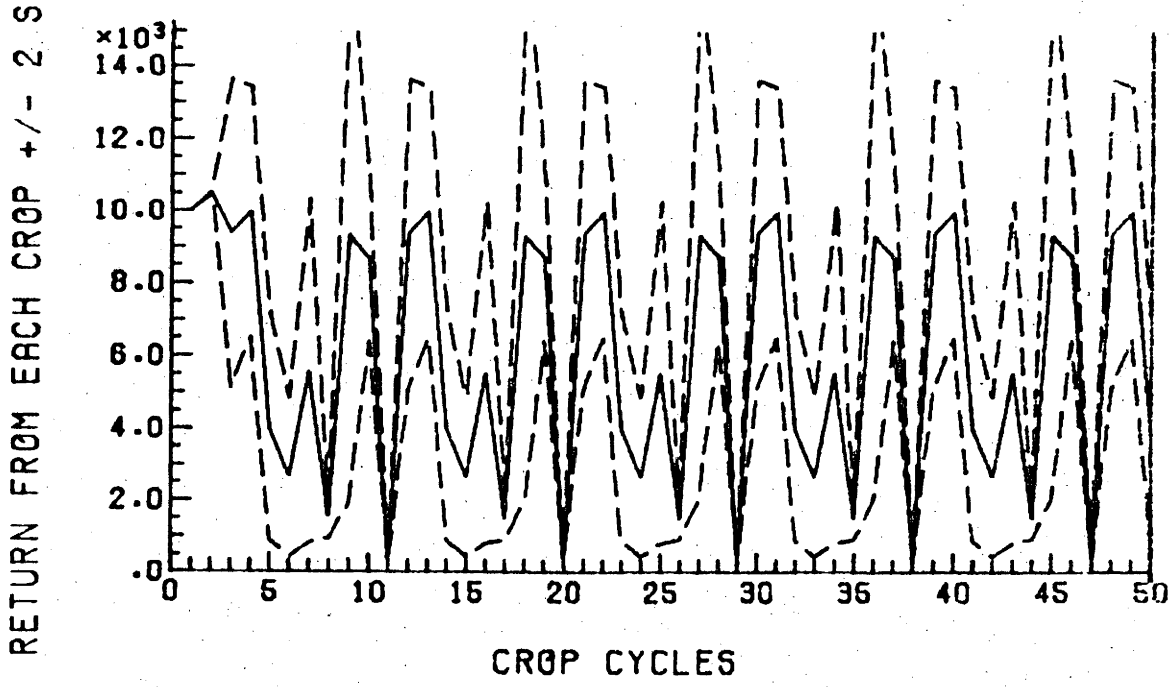

FIG 6.4.71: AGRIBUSINESS + CROSSPROTECTION
YIELD THROUGH TIME

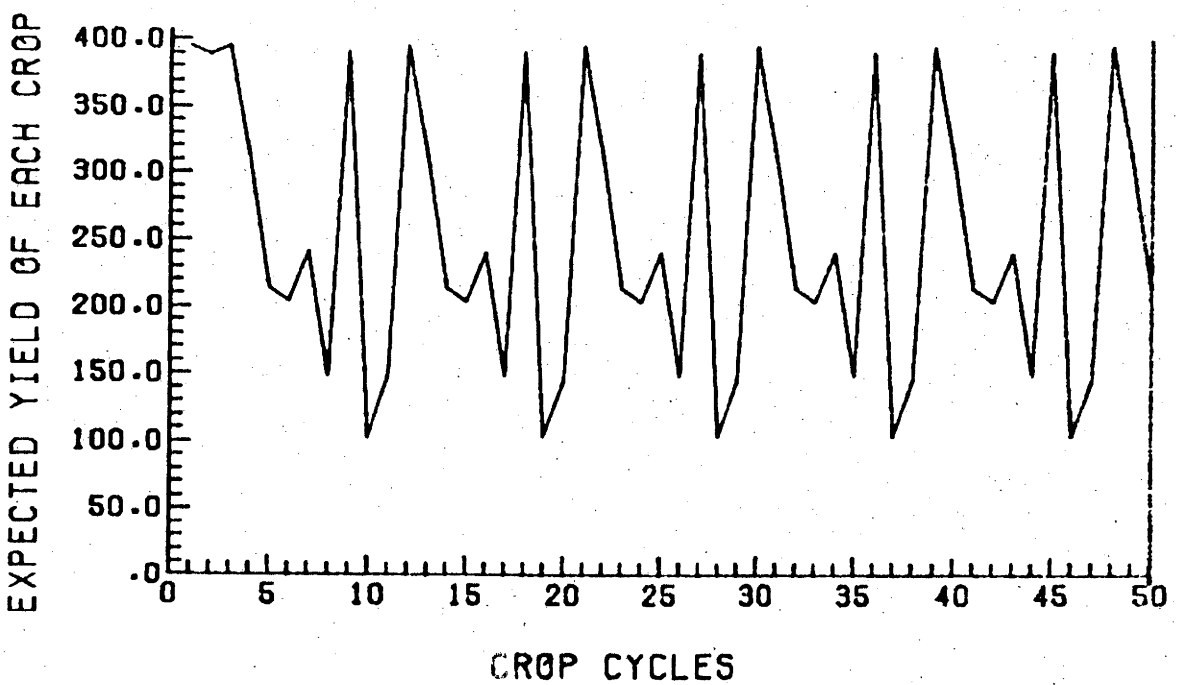

FIG 6.4.72: AGRIBUSINESS + CROSSPROTECTION
STORAGE THROUGH TIME

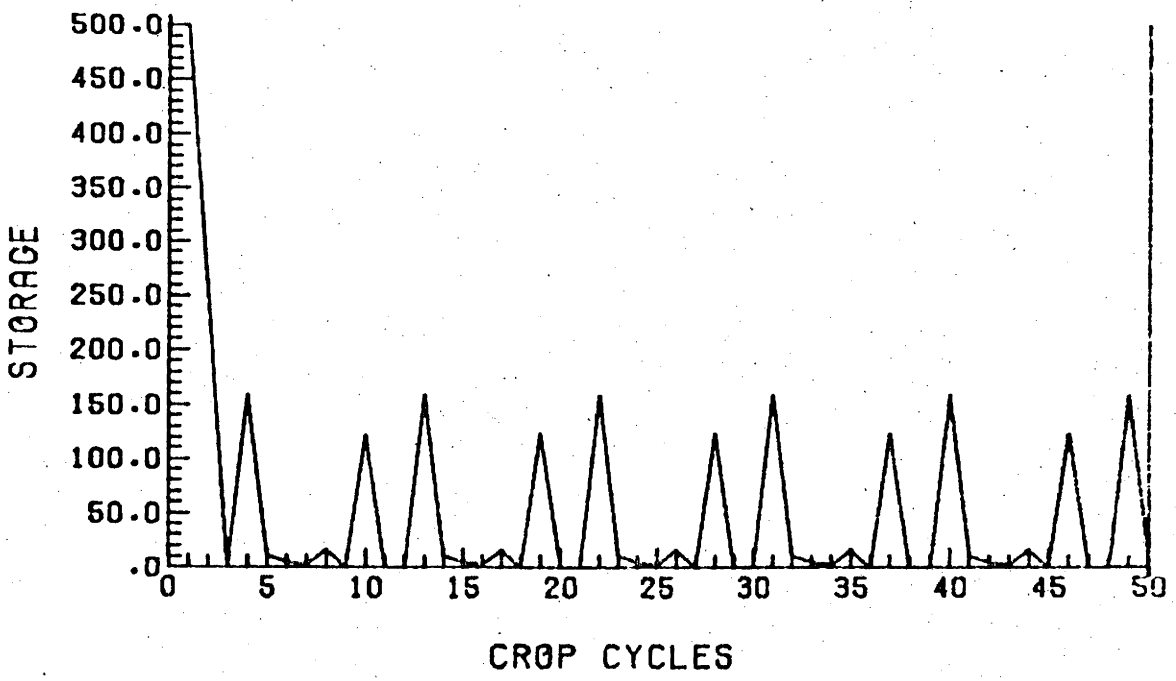

FIG 6.4.73: AGRIBUSINESS + CROSSPROTECTION
SUPERRACE THROUGH TIME

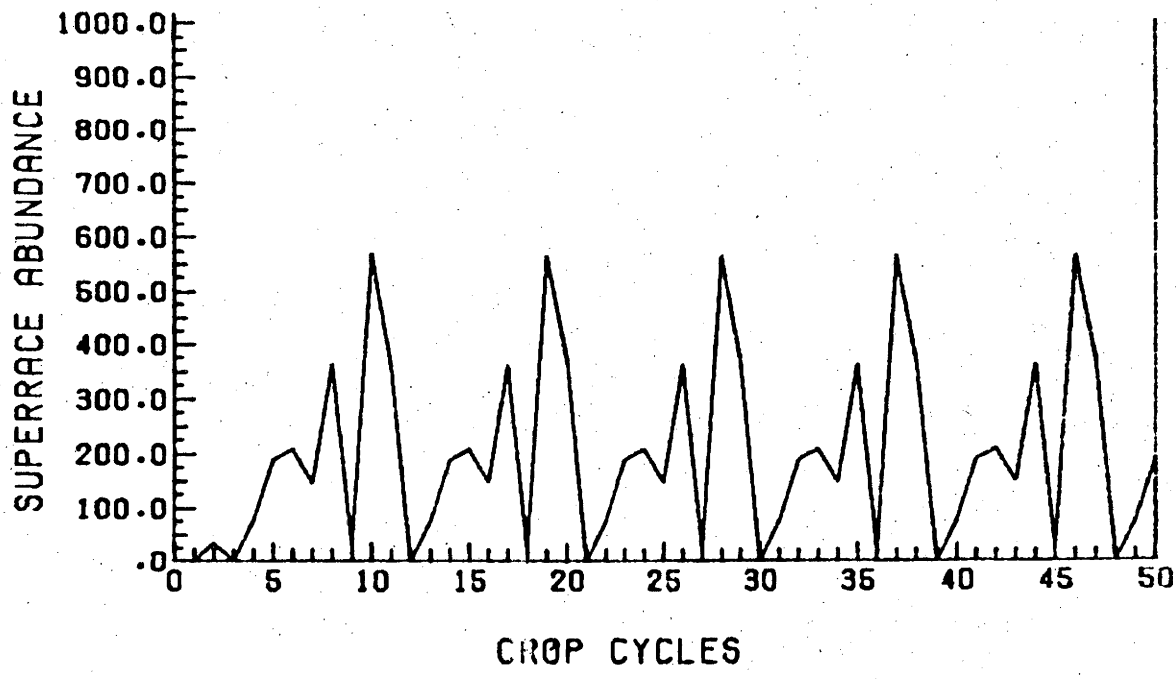

FIG. 6.4.74A: NATURAL INCREASE OF *FEEBLE* SUPERRACE WITHOUT COMPETITION

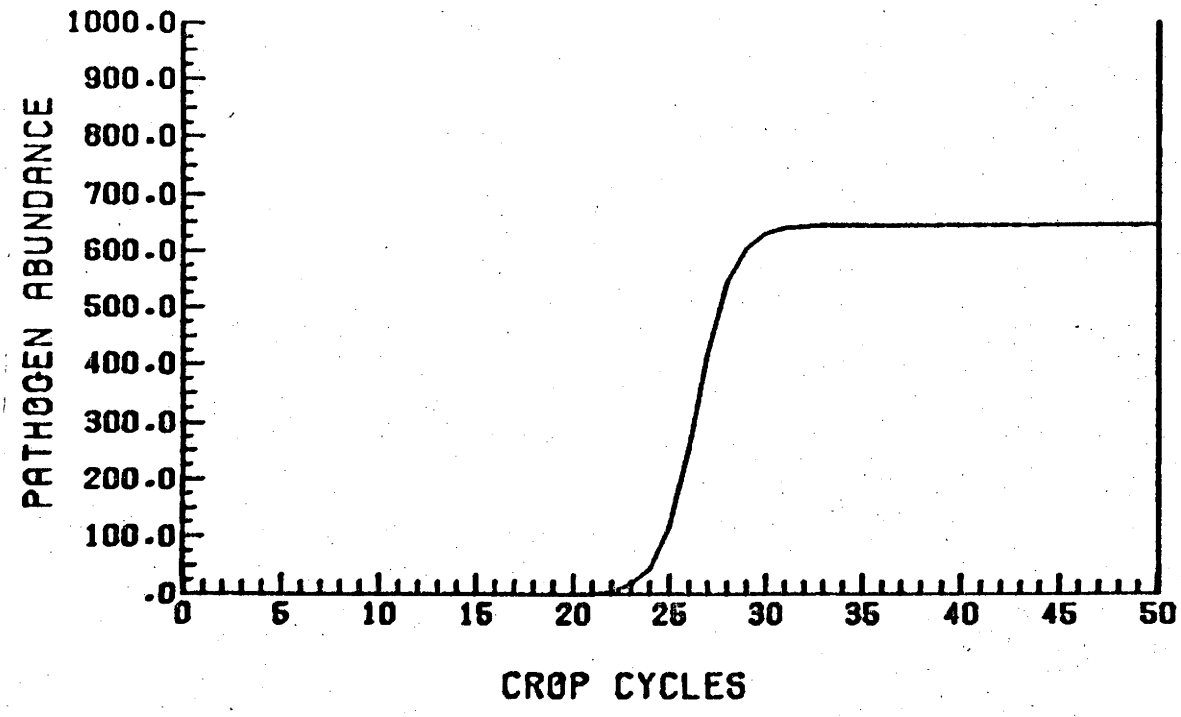

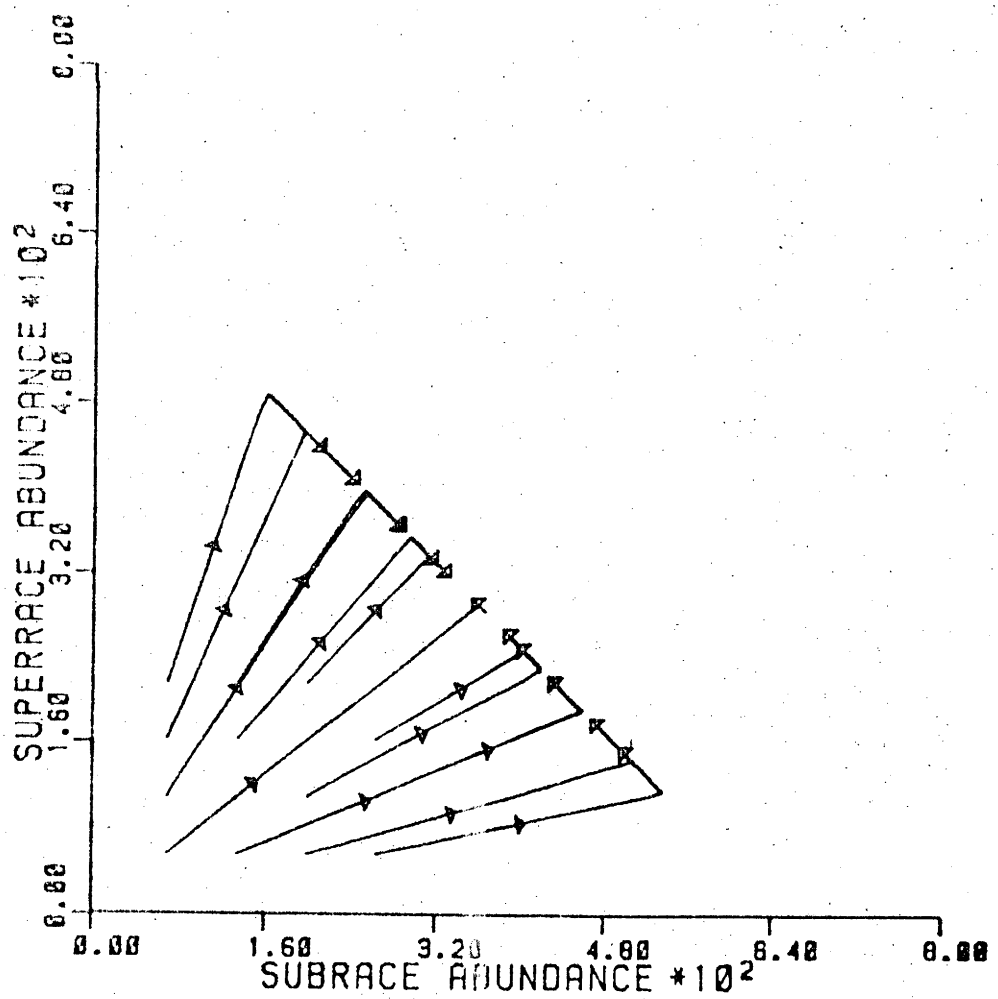

FIG. 6.4.74B: TRAJECTORIES OF PATHOGEN ABUNDANCE

FIG. 6.4.74C: LONG-TERM YIELD UNDER STABILISING SELECTION

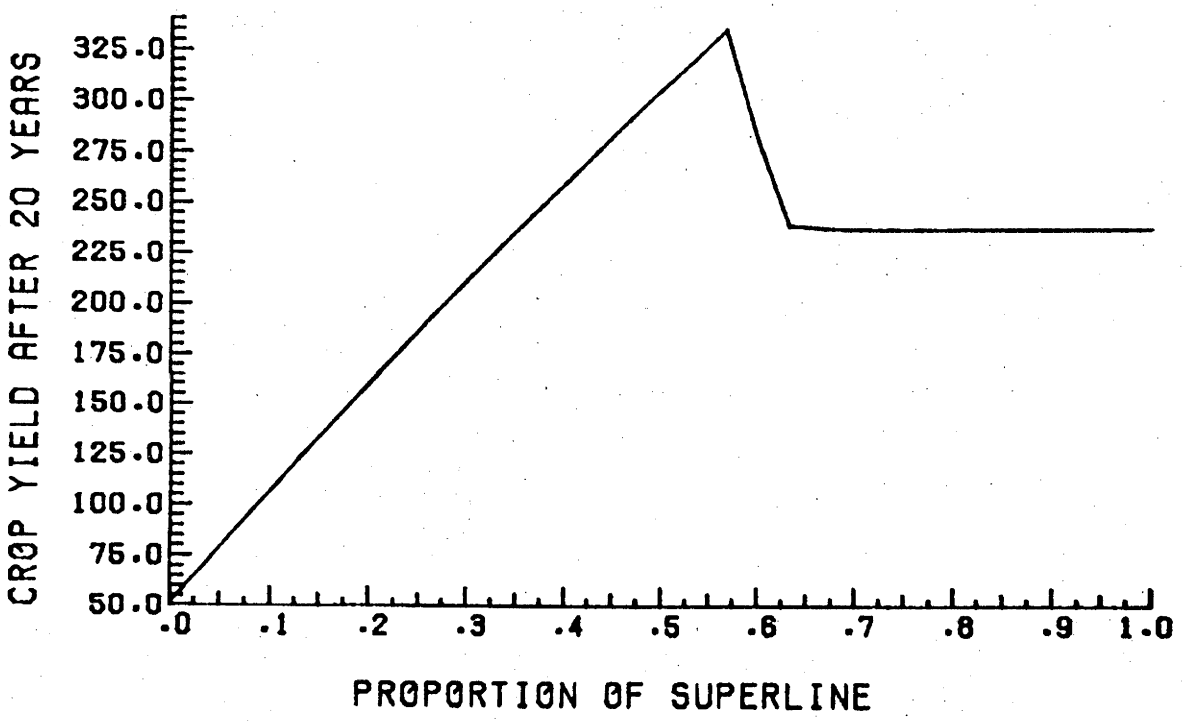

FIG 6.4.75A: AGRIBUSINESS + STABILISING SELECTION
FRACTION SOLD VS STORAGE

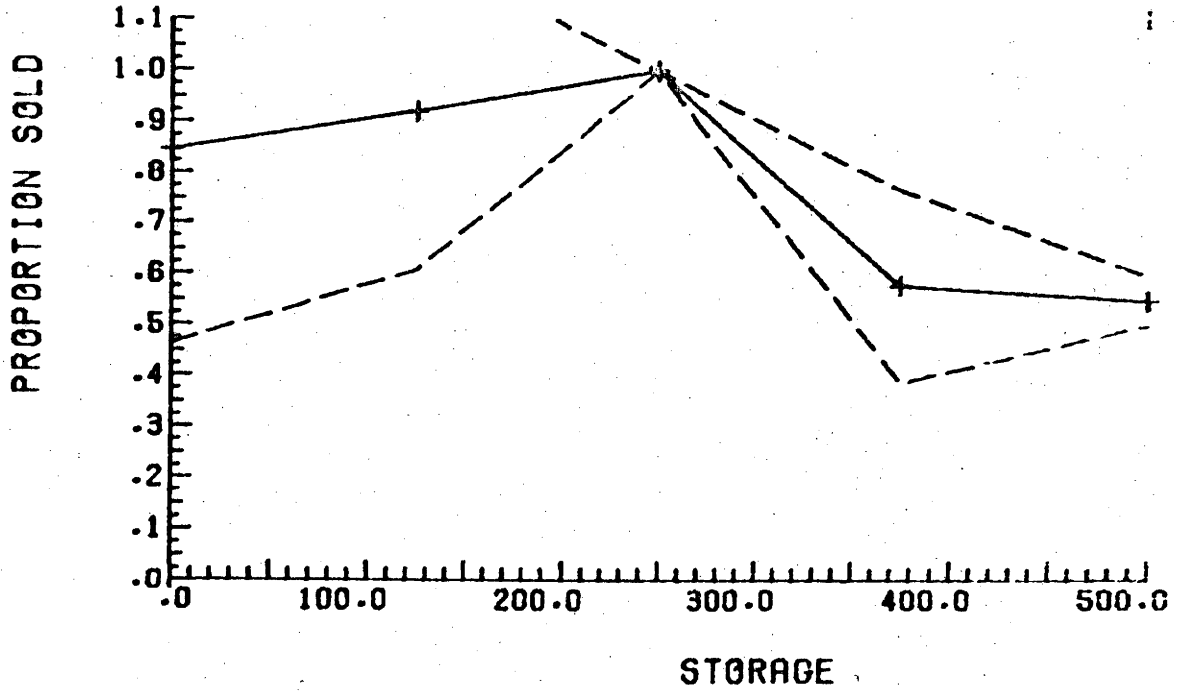

FIG 6.4.75B: AGRIBUSINESS + STABILISING SELECTION
FRACTION SOLD VS AABB

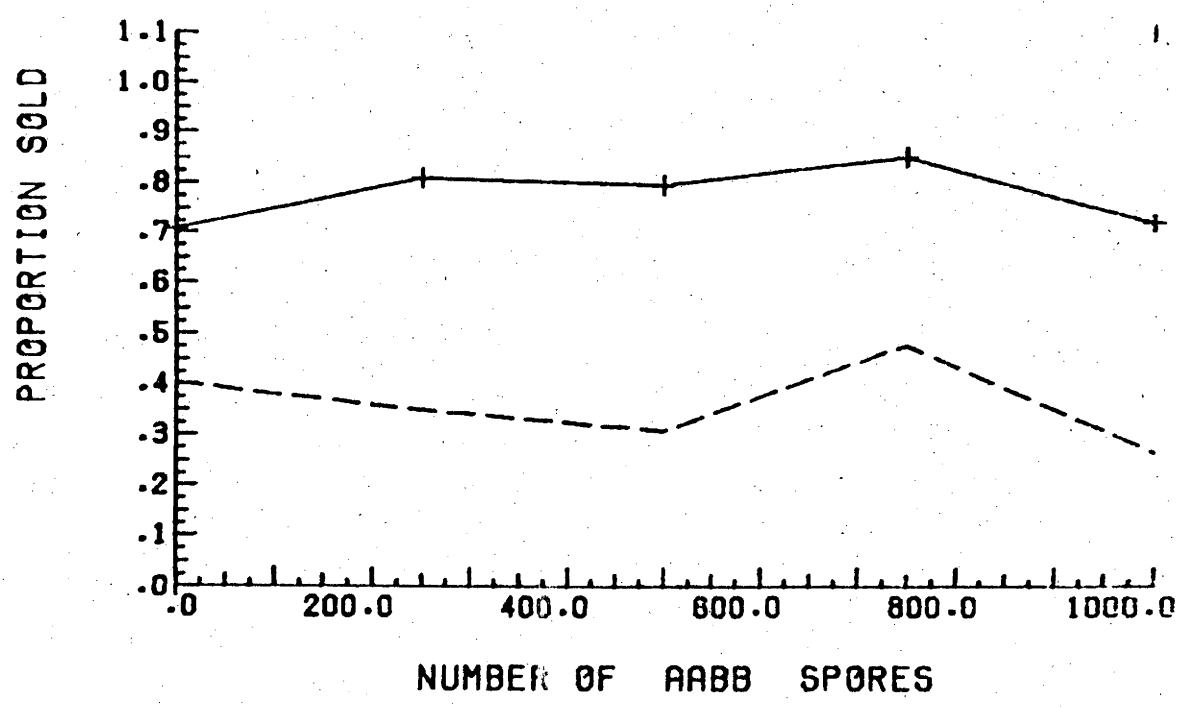

FIG 6.4.76A: AGRIBUSINESS + STABILISING SELECTION
AREA PLANTED VS STORAGE

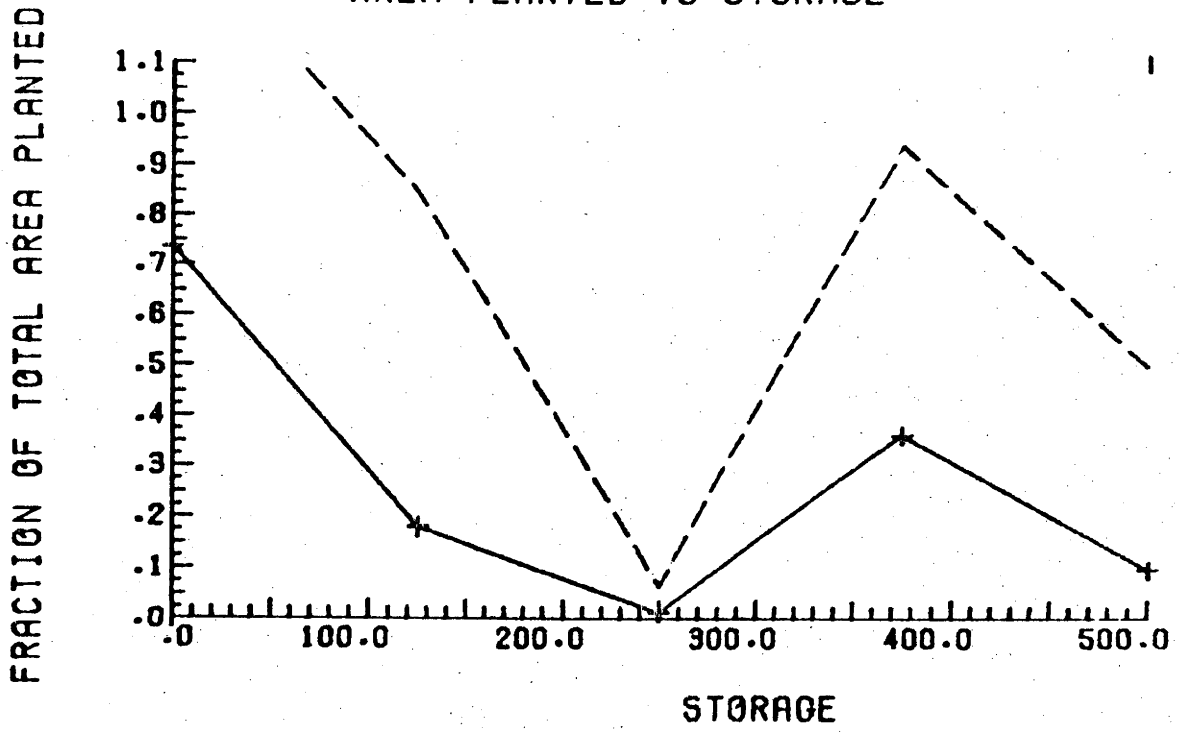

FIG 6.4.76B: AGRIBUSINESS + STABILISING SELECTION
AREA PLANTED VS AABB

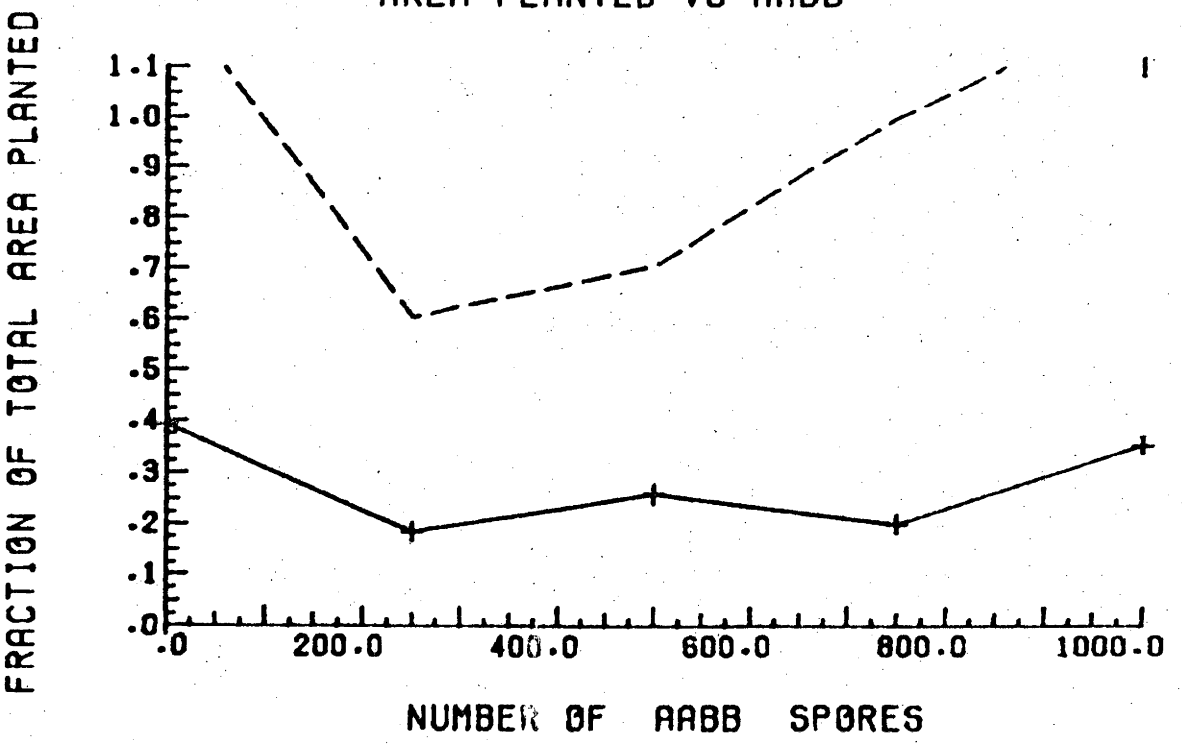

FIG 6.4.77A: AGRIBUSINESS + STABILISING SELECTION
E(RETURN) VS STORAGE

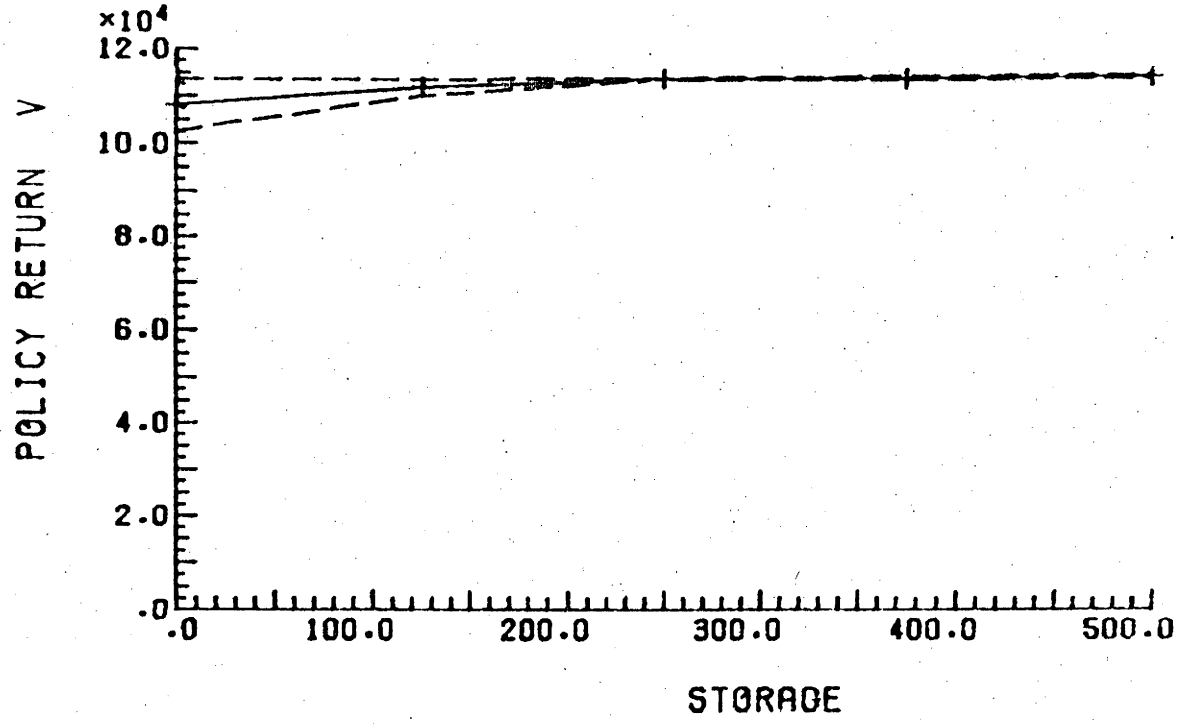

FIG 6.4.77B: AGRIBUSINESS + STABILISING SELECTION
E(RETURN) VS SUPERRACE

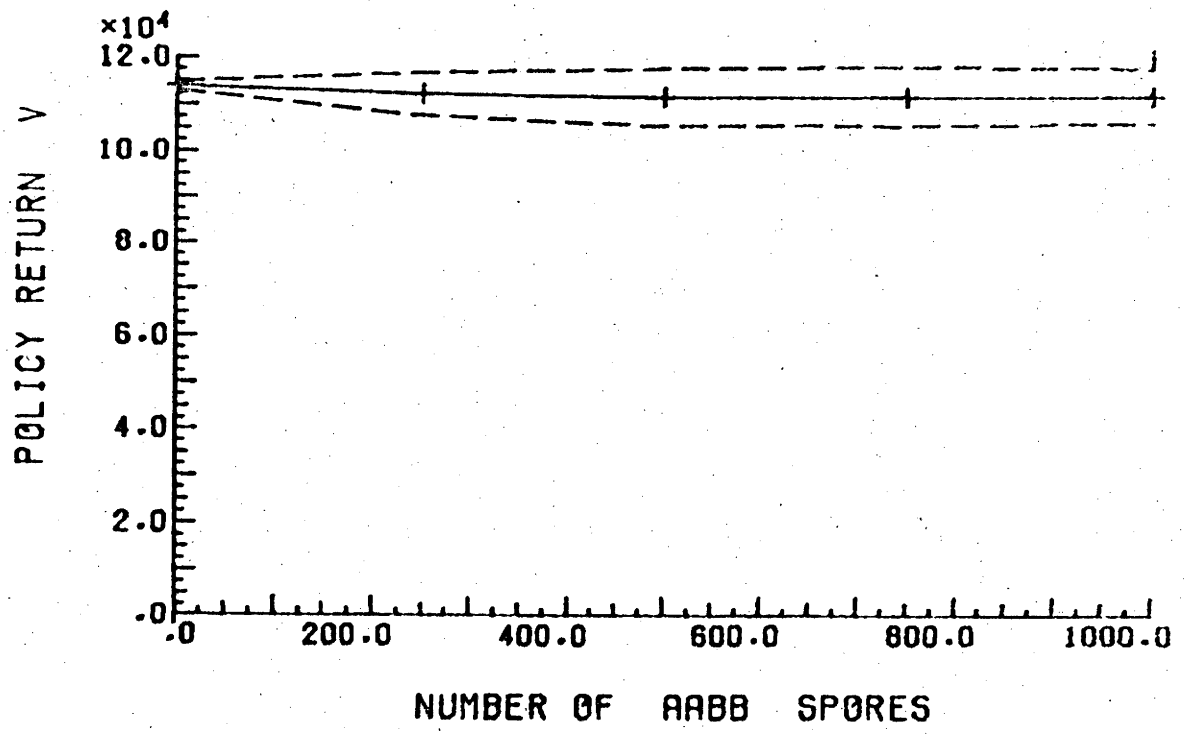

FIG 6.4.78A: AGRIBUSINESS + STABILISING SELECTION
STORAGE LEFT VS STORAGE

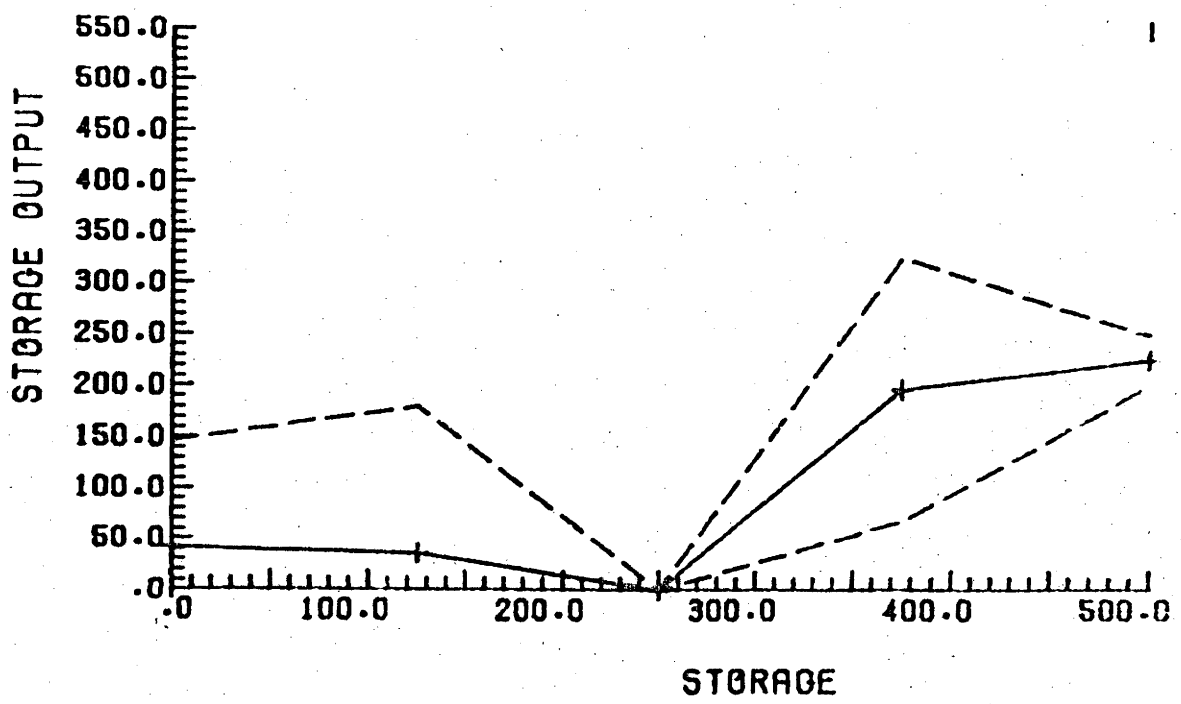

FIG 6.4.78B: AGRIBUSINESS + STABILISING SELECTION
STORAGE LEFT VS AABB

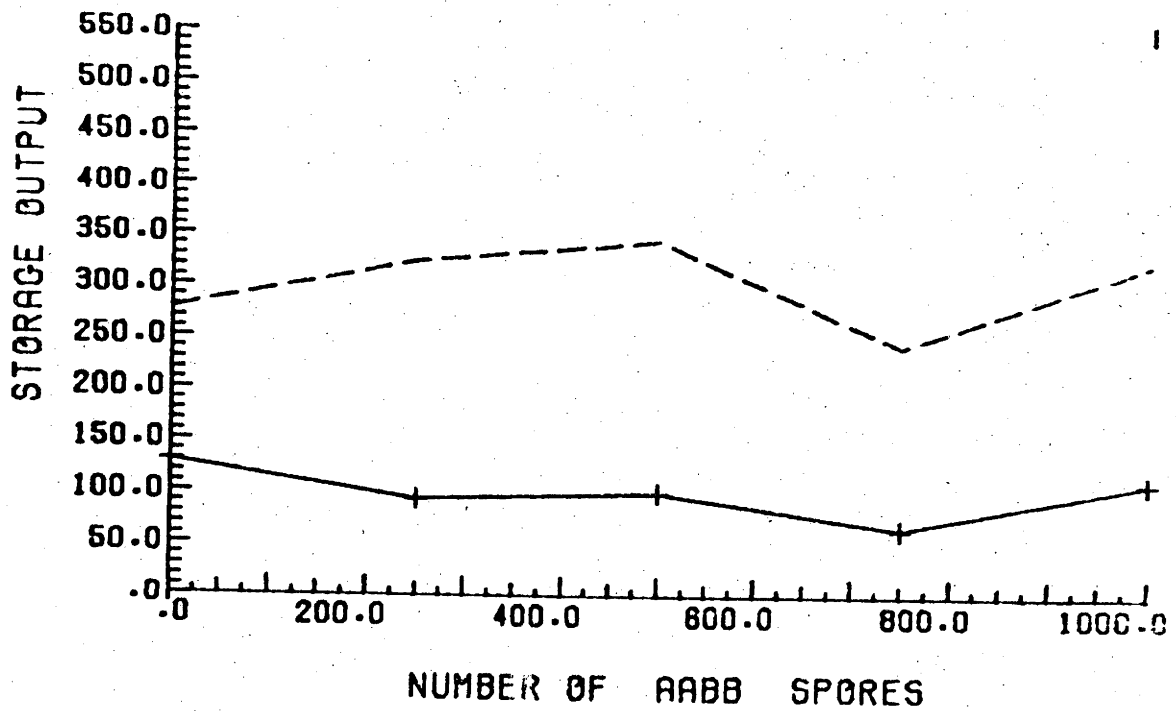

FIG 6.4.79A: AGRIBUSINESS + STABILISING SELECTION
AABB OUTPUT VS STORAGE

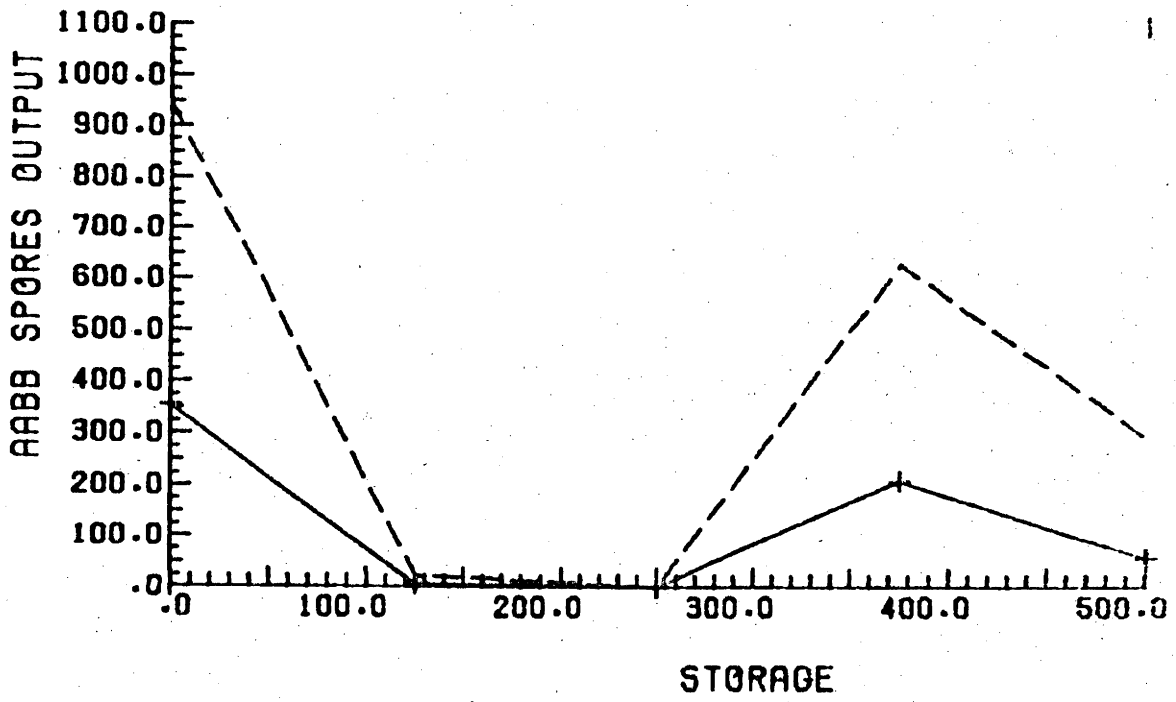

FIG 6.4.79B: AGRIBUSINESS + STABILISING SELECTION
AABB OUTPUT VS AABB

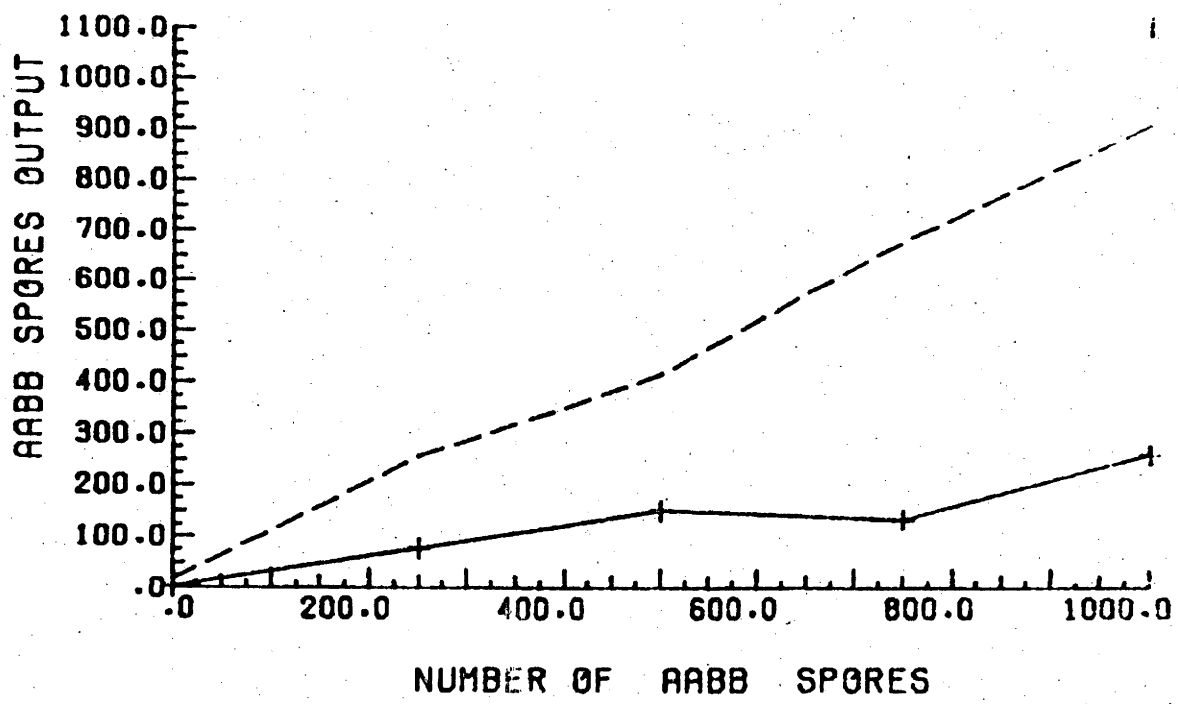

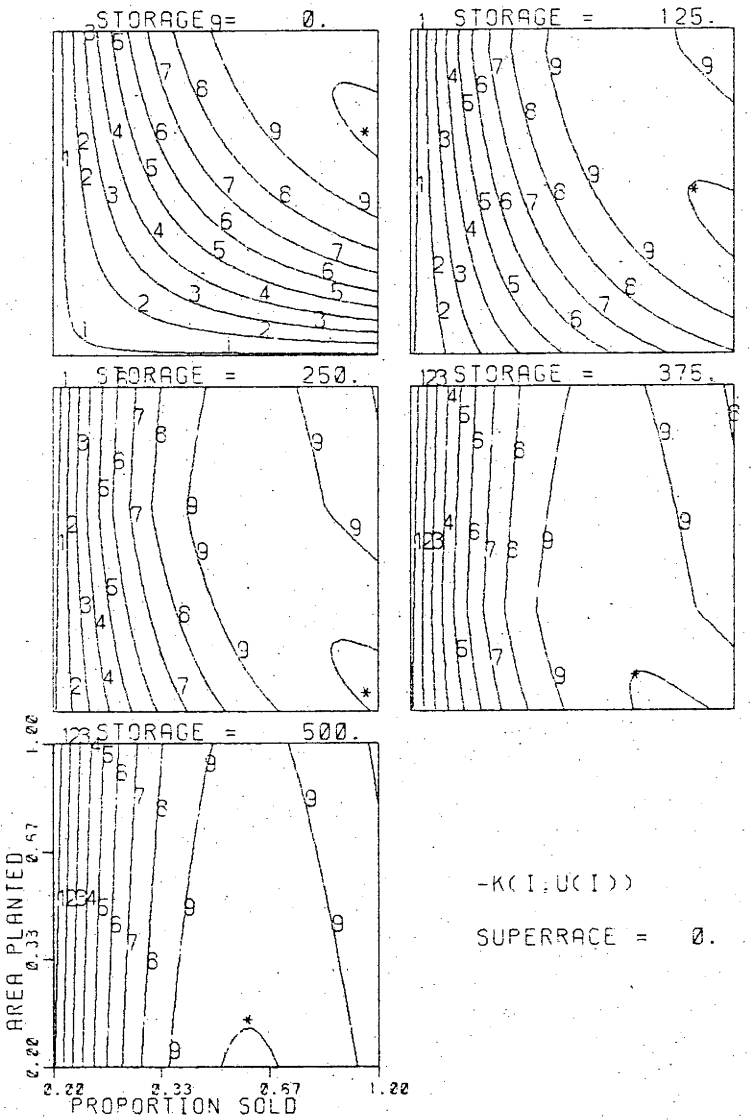

FIG. 6.4.80A: RETURN WITH STABILISING SELECTION

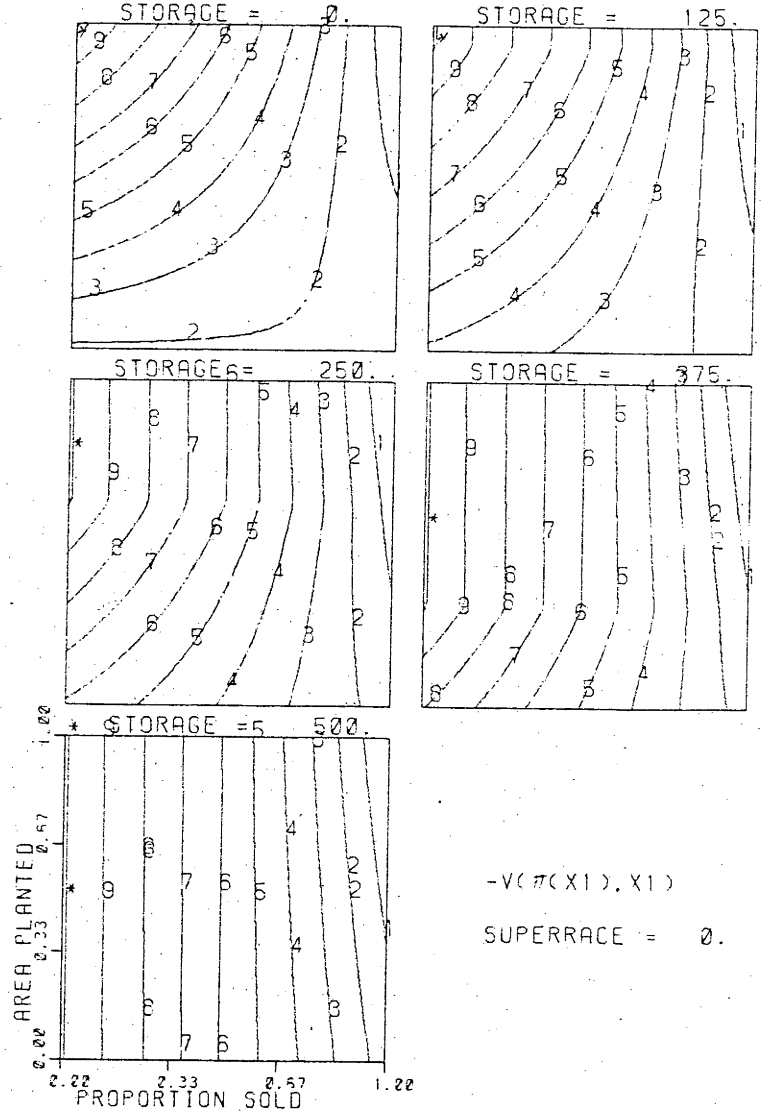

FIG. 6.4.80B: RETURN WITH STABILISING SELECTION

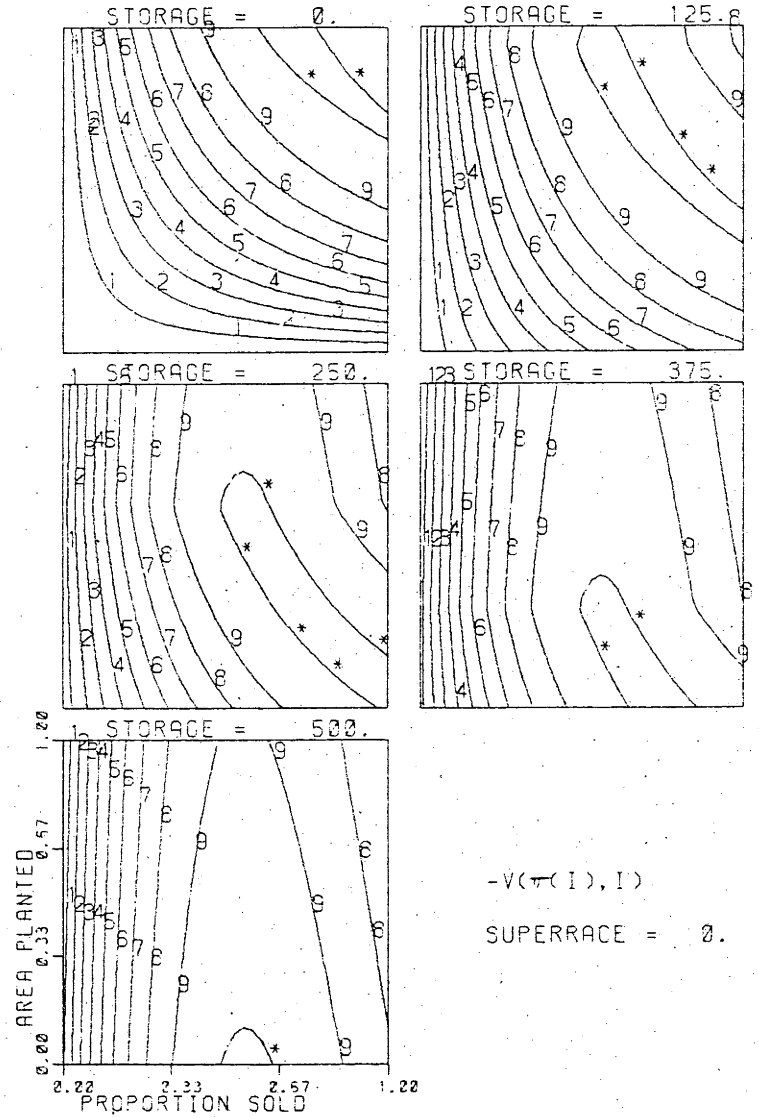

-V(π(I),I)

SUPERRACE = 0.

FIG. 5.4.20C: RETURN WITH STABILISING SELECTION

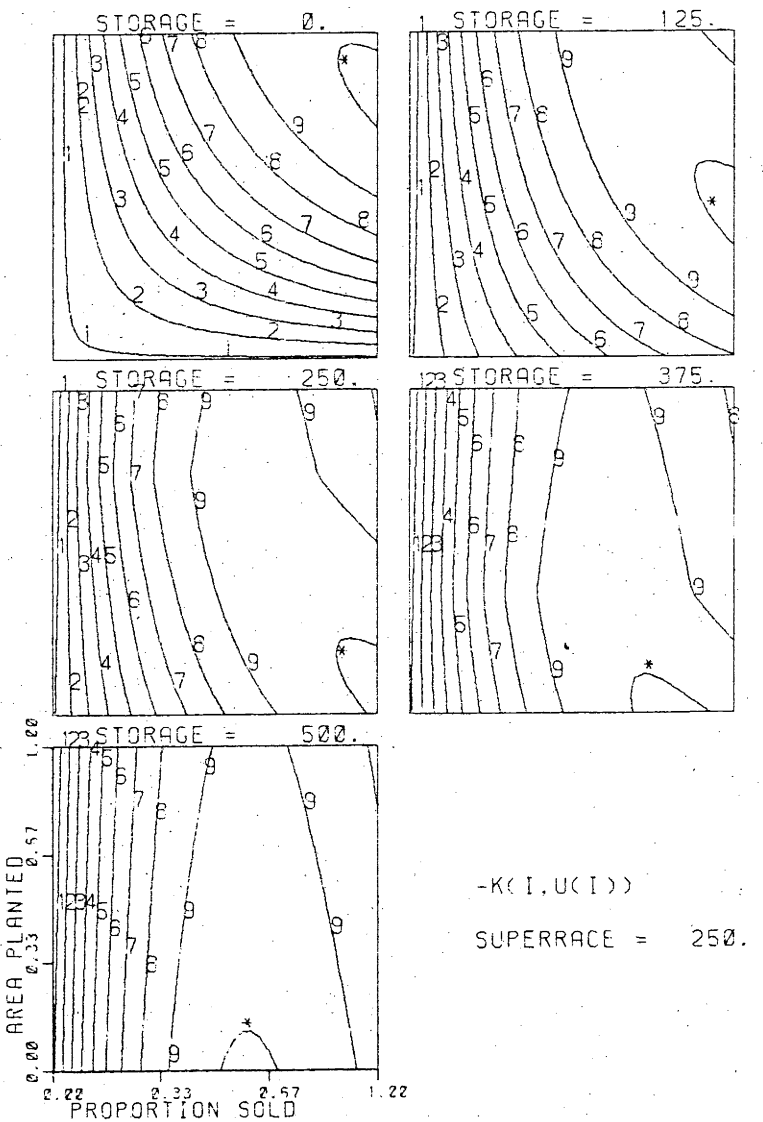

-K(I,U(I))

SUPERRACE = 250.

FIG. 5.4.21A: RETURN WITH STABILISING SELECTION

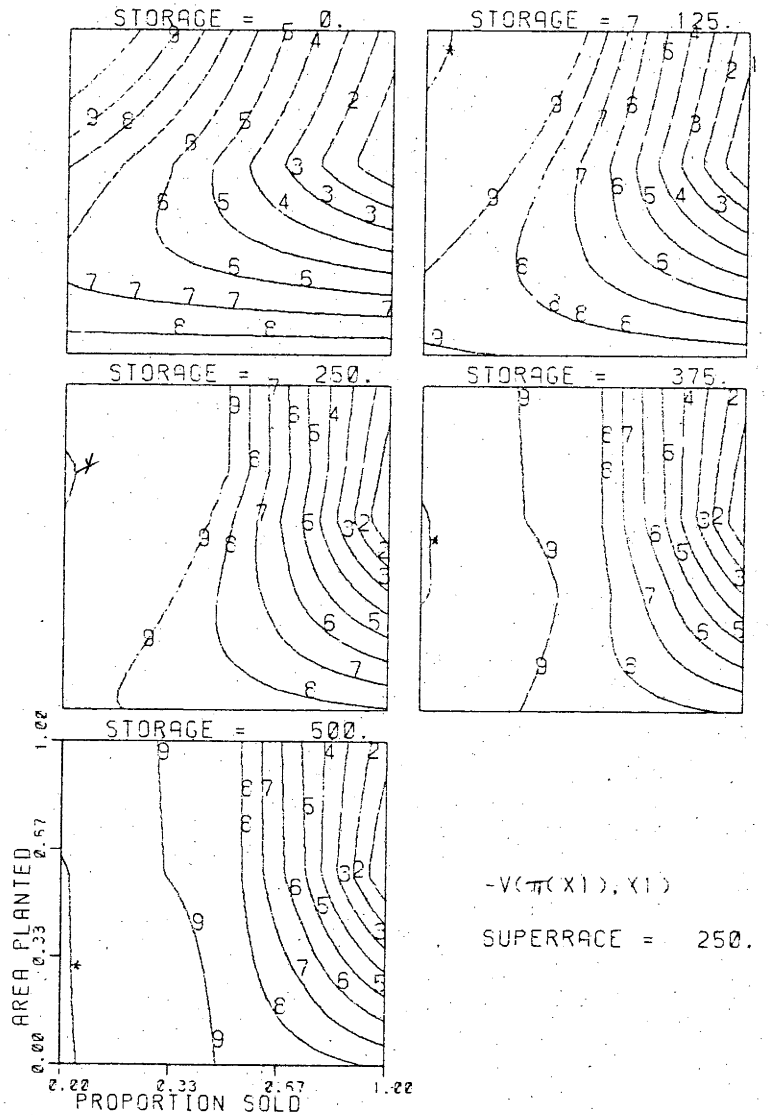

FIG. 6.4.8B: RETURN WITH STABILISING SELECTION

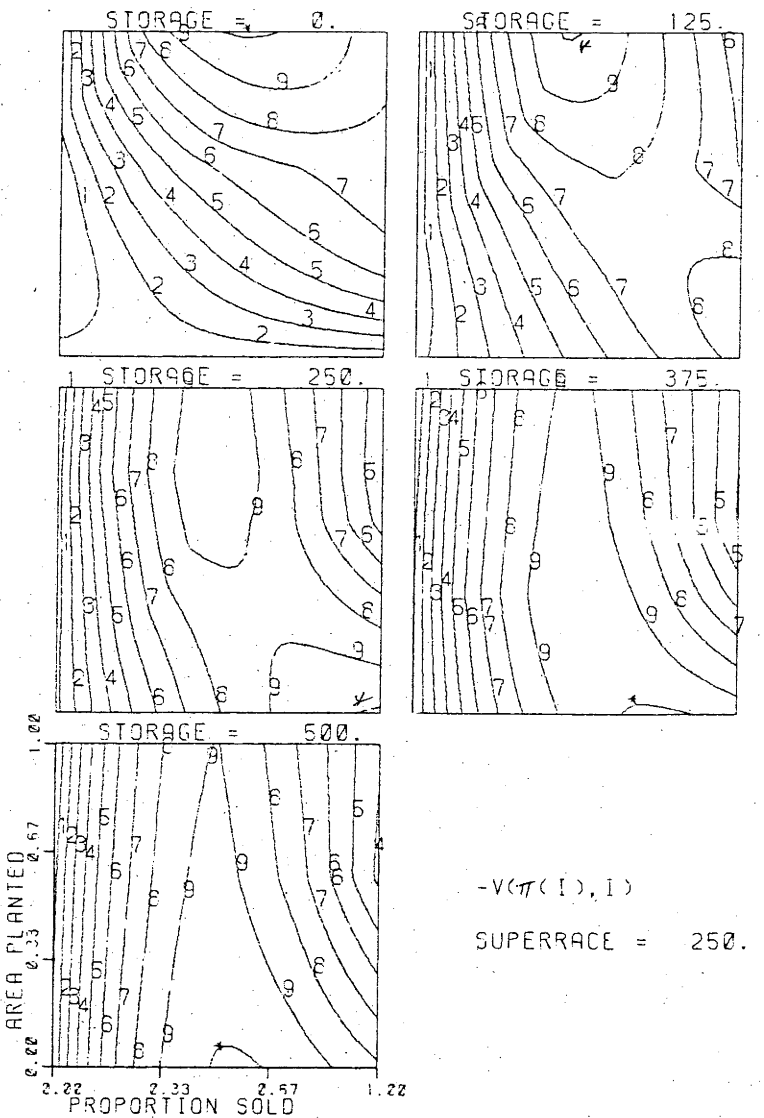

FIG. 6.4.8C: RETURN WITH STABILISING SELECTION

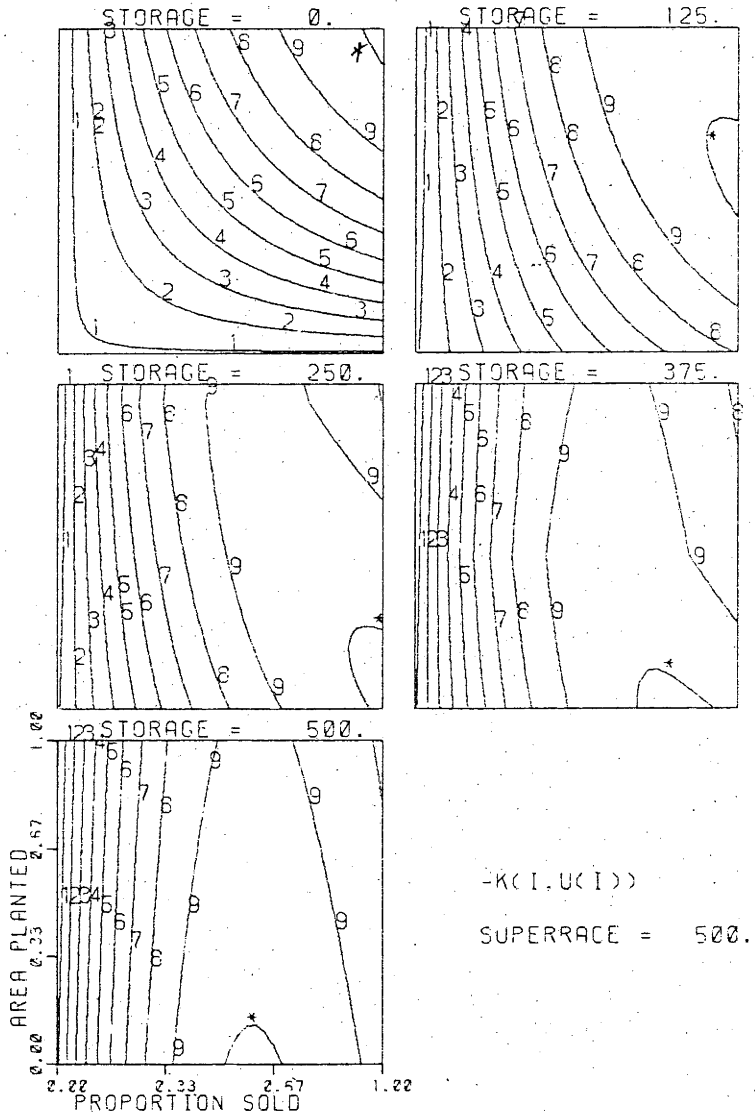

FIG. 6.4.82A: RETURN WITH STABILISING SELECTION

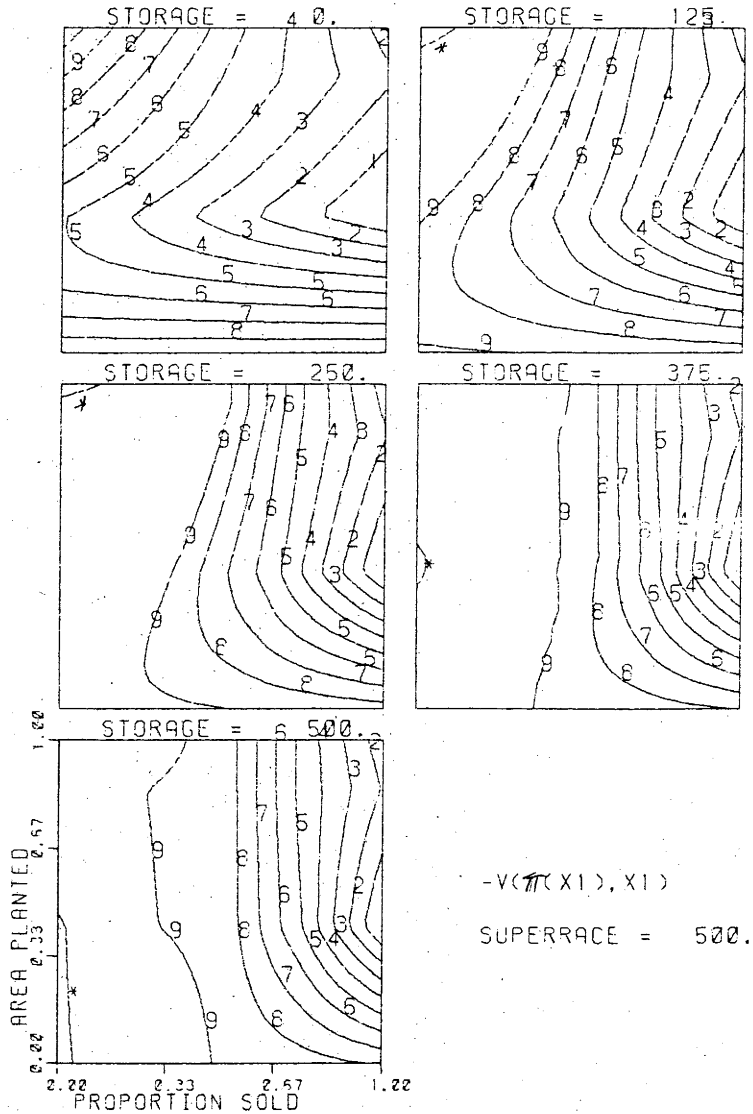

FIG. 6.4.82B: RETURN WITH STABILISING SELECTION

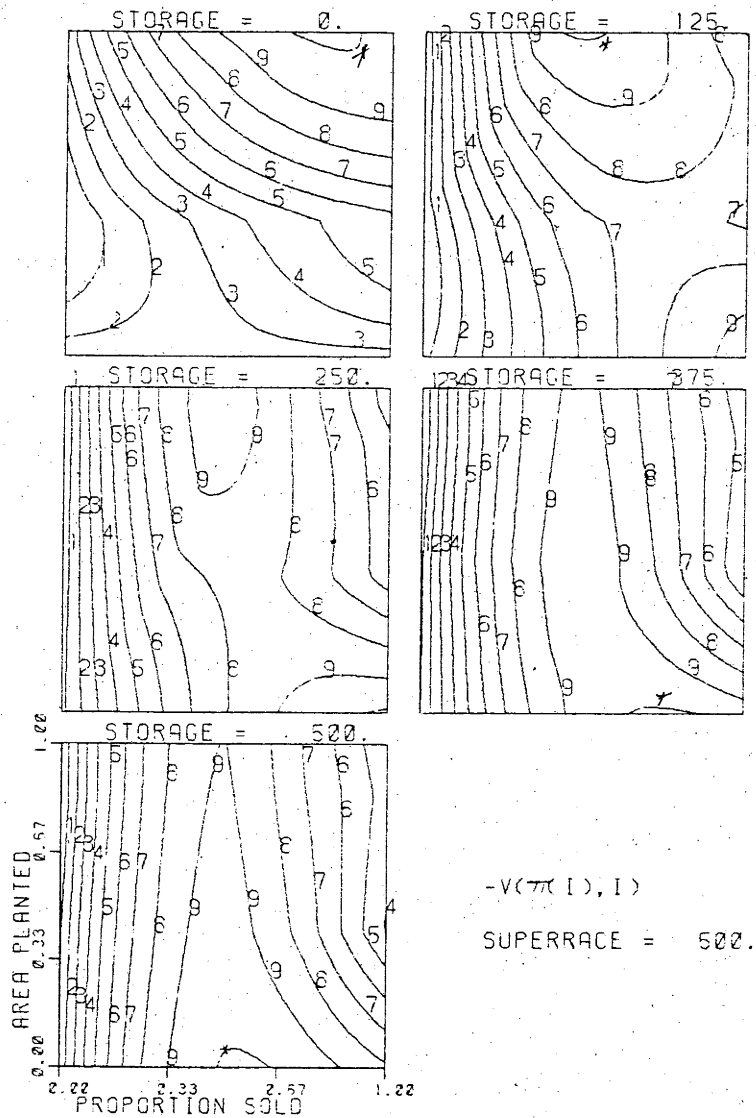

FIG. 6.4.22C: RETURN WITH STABILISING SELECTION

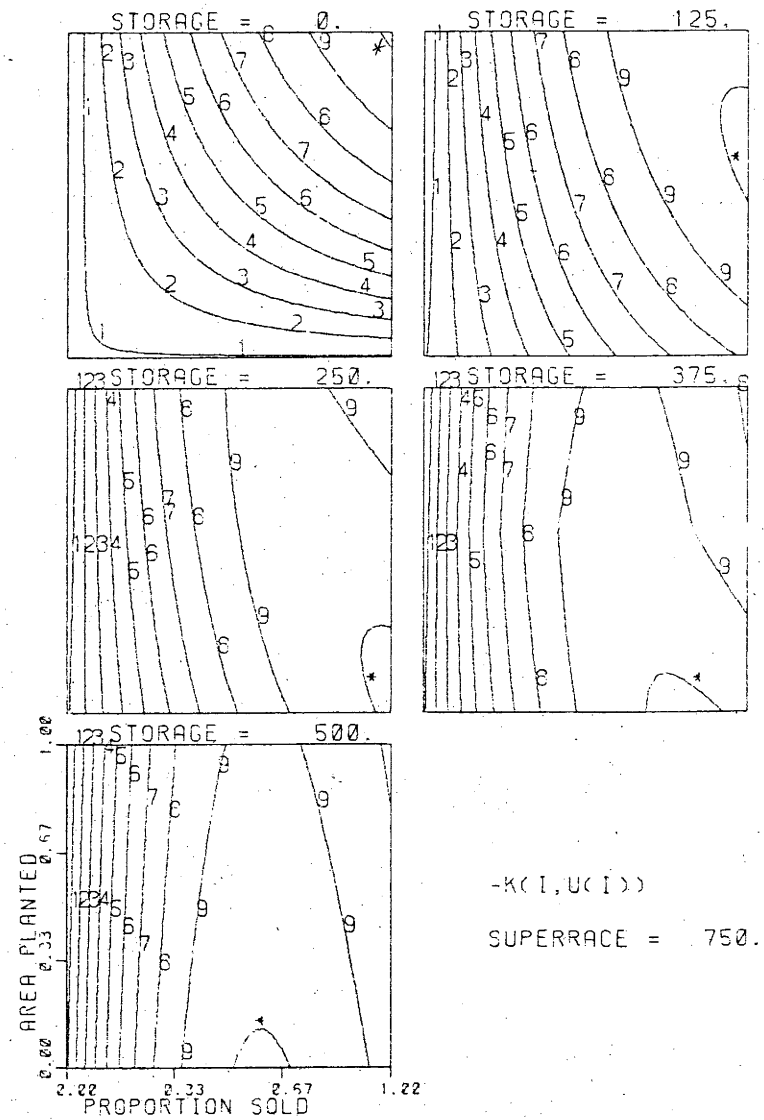

FIG. 6.4.23A: RETURN WITH STABILISING SELECTION

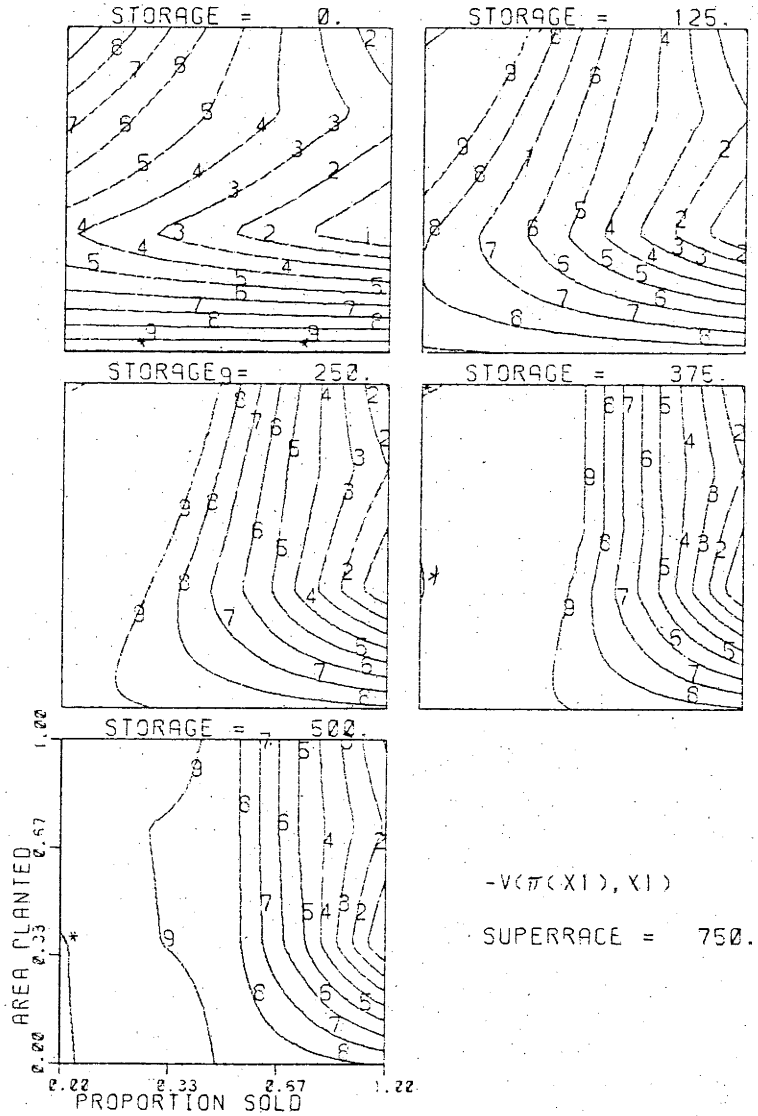

-V(π(X1),X1)

SUPERRACE = 750.

FIG. 6.4.83B: RETURN WITH STABILISING SELECTION

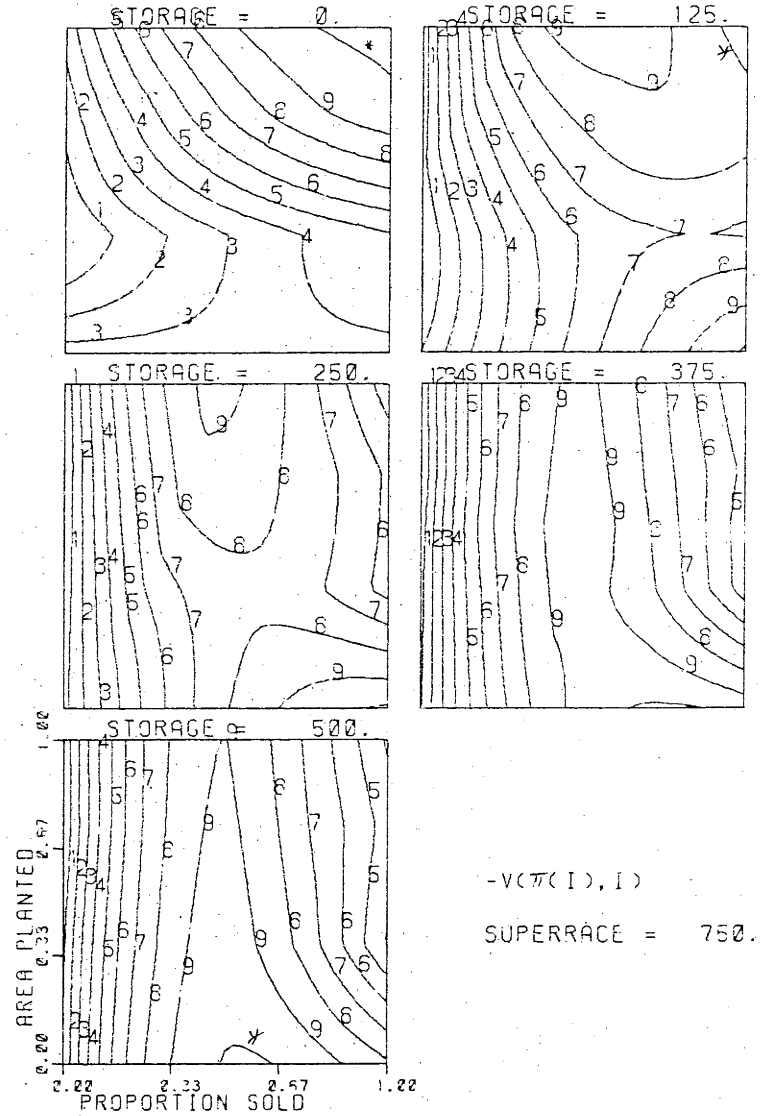

-V(π(I),I)

SUPERRACE = 750.

FIG. 6.4.83C: RETURN WITH STABILISING SELECTION

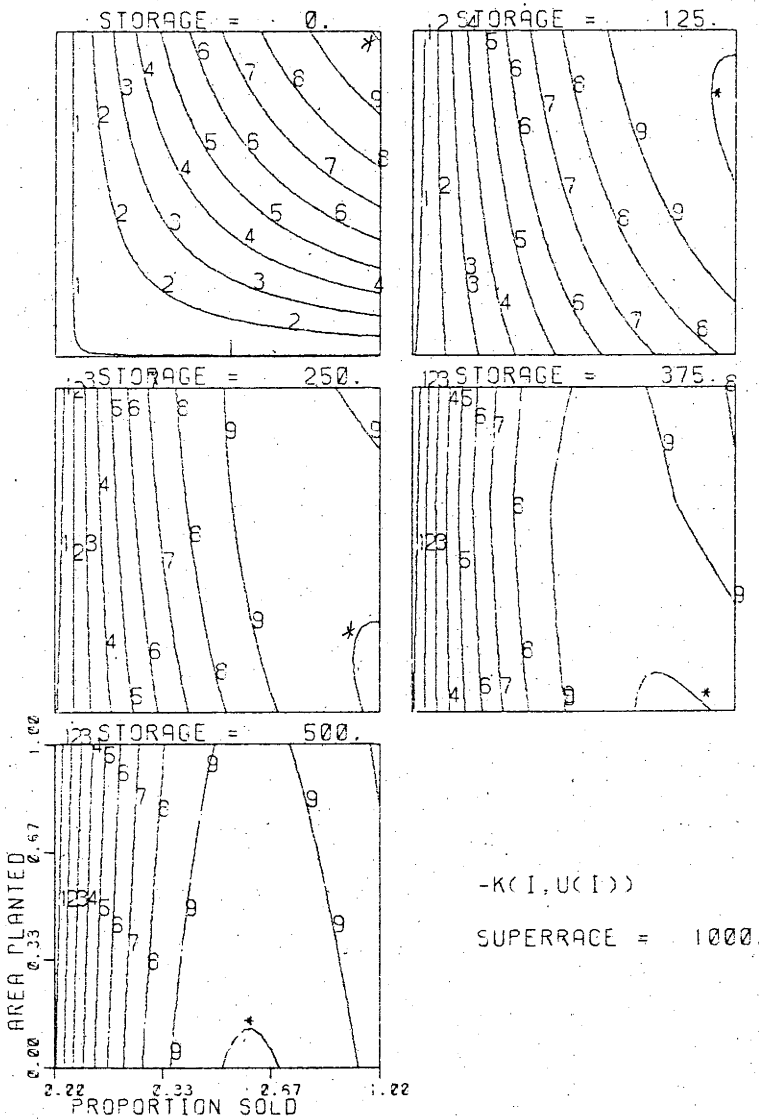

FIG. 6.4.24A: RETURN WITH STABILISING SELECTION

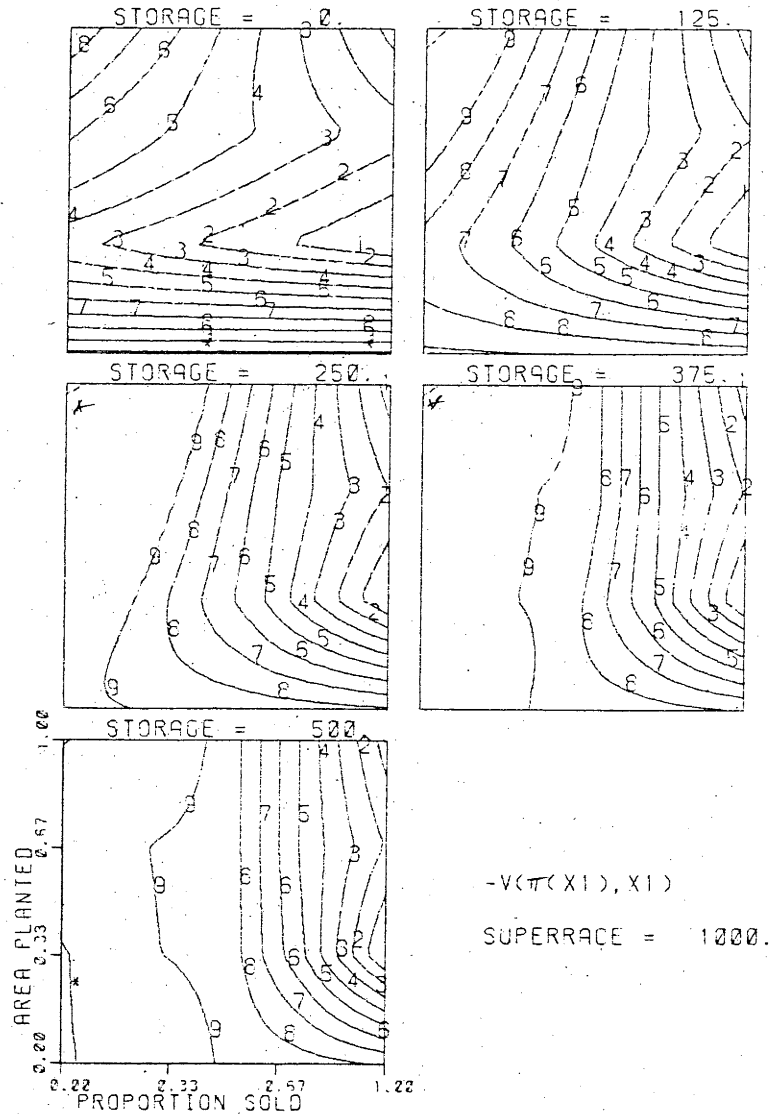

FIG. 6.4.24B: RETURN WITH STABILISING SELECTION

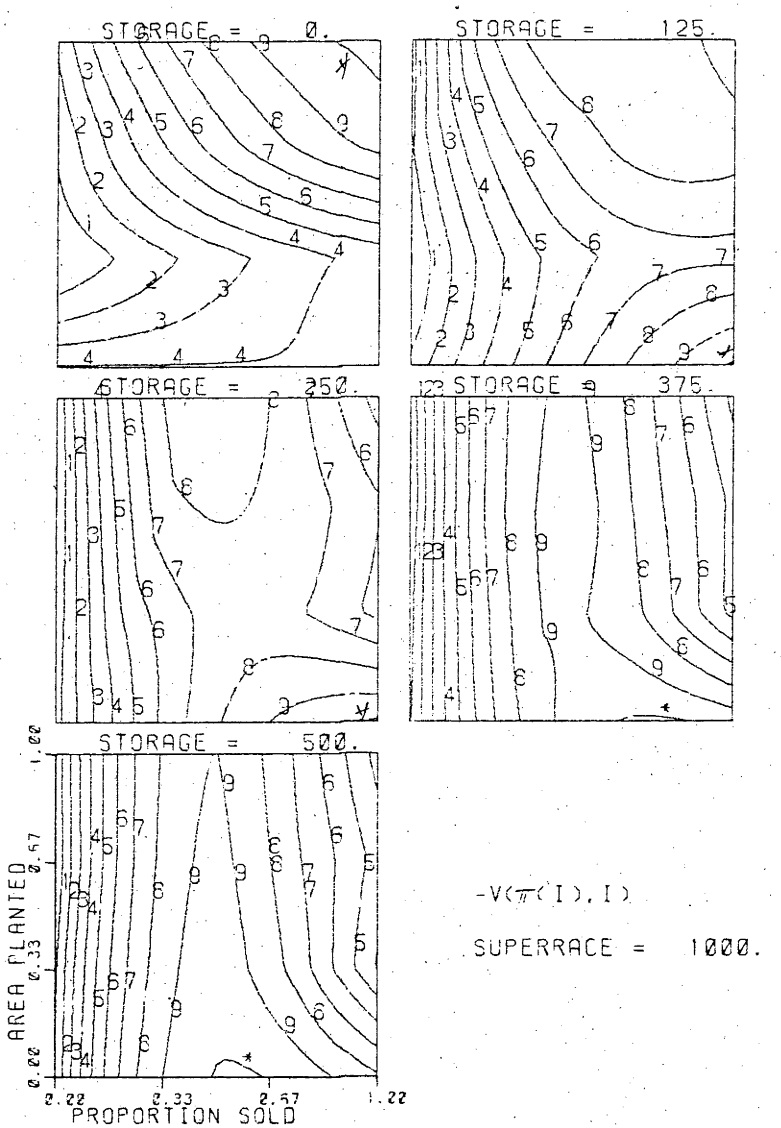

FIG. 6.4.84C: RETURN WITH STABILISING SELECTION

FIG 6.4.85: AGRIBUSINESS + STABILISING SELECTION: FRACTION SOLD VS TIME

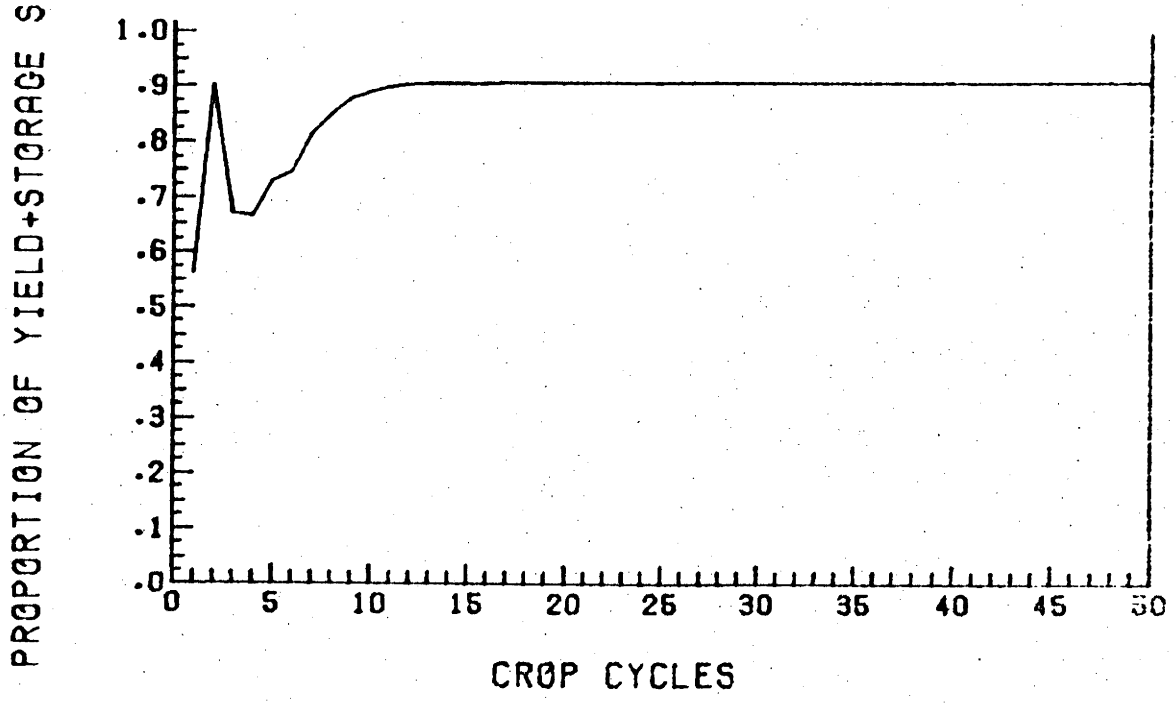

FIG 6.4.86: AGRIBUSINESS + STABILISING SELECTION: AREA PLANTED VS TIME

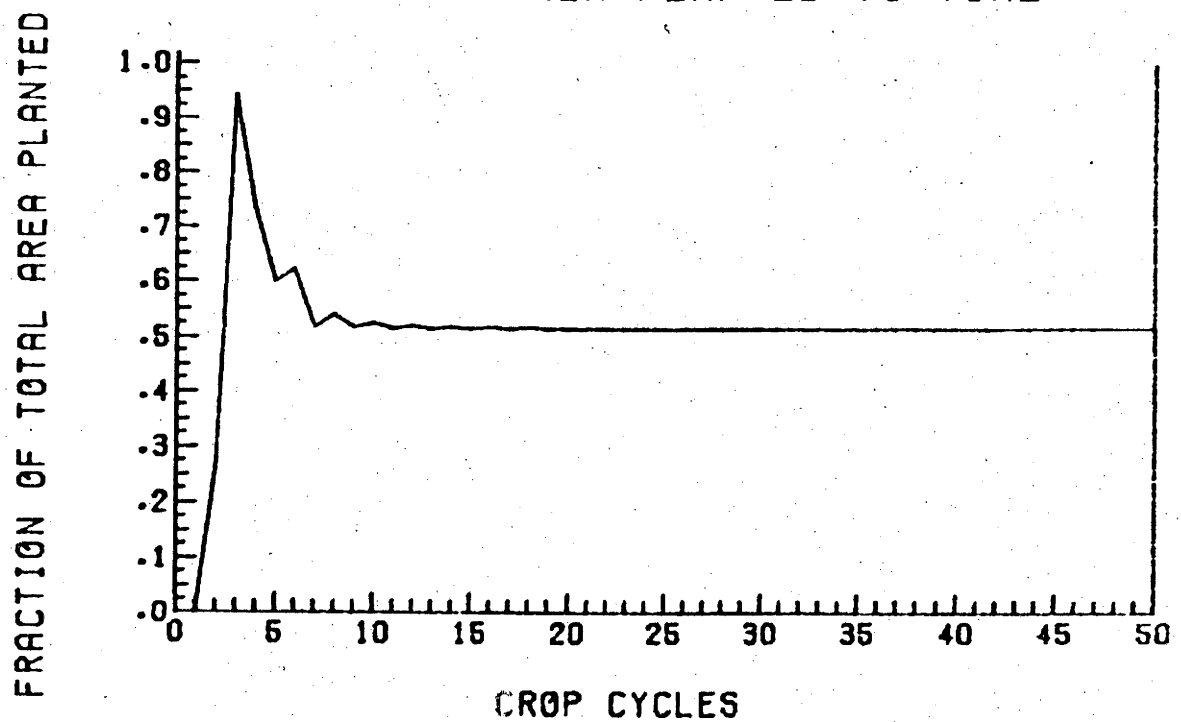

FIG 6.4.87: AGRIBUSINESS + STABILISING SELECTION: RETURN VS TIME

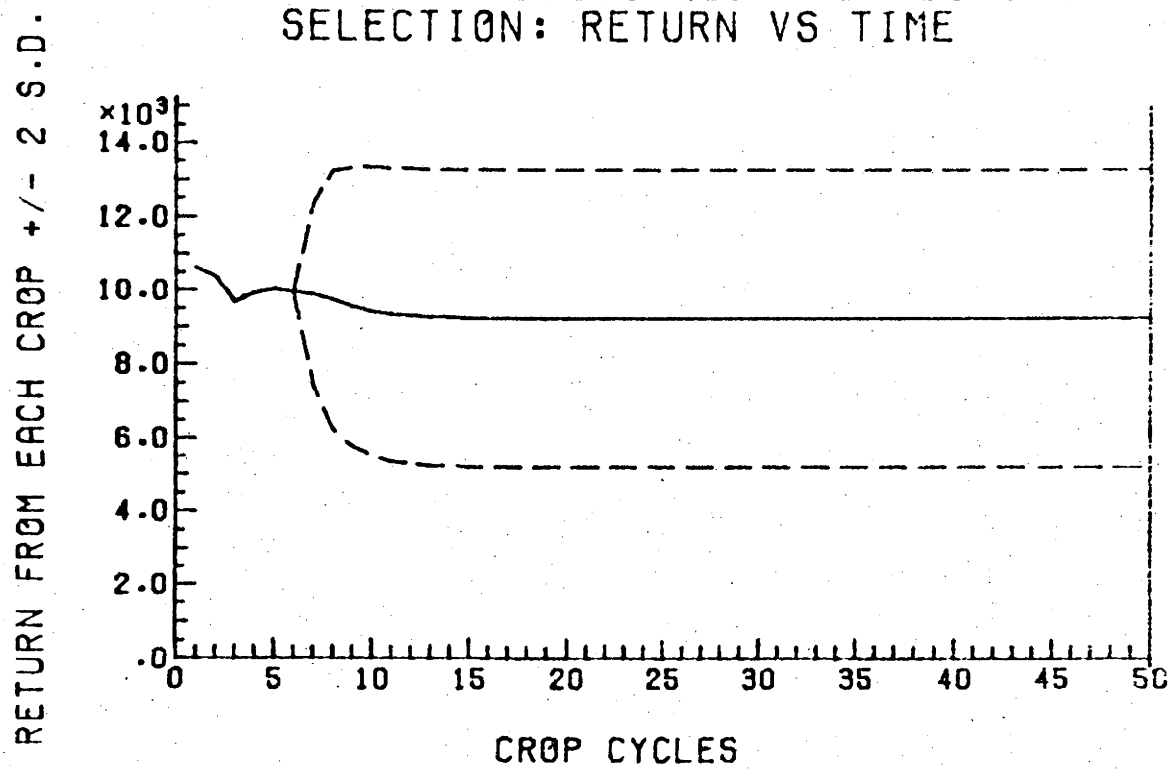

FIG 6.4.88: AGRIBUSINESS + STABILISING SELECTION: YIELD VS TIME

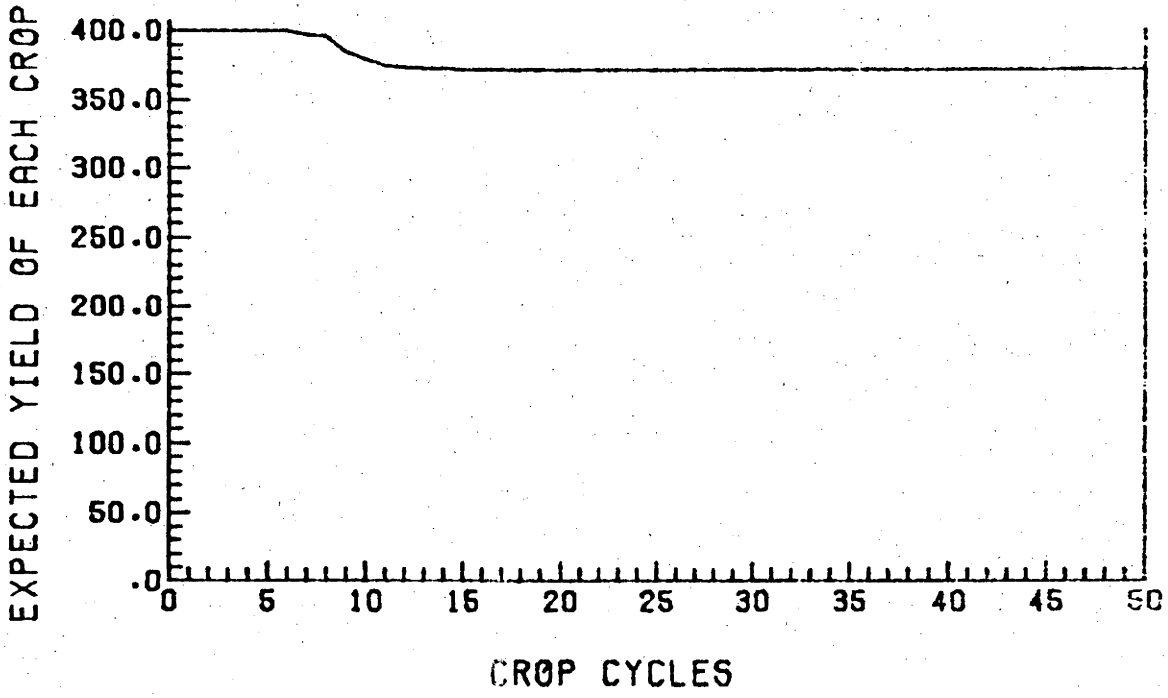

FIG 6.4.89: AGRIBUSINESS + STABILISING SELECTION: STORAGE VS TIME

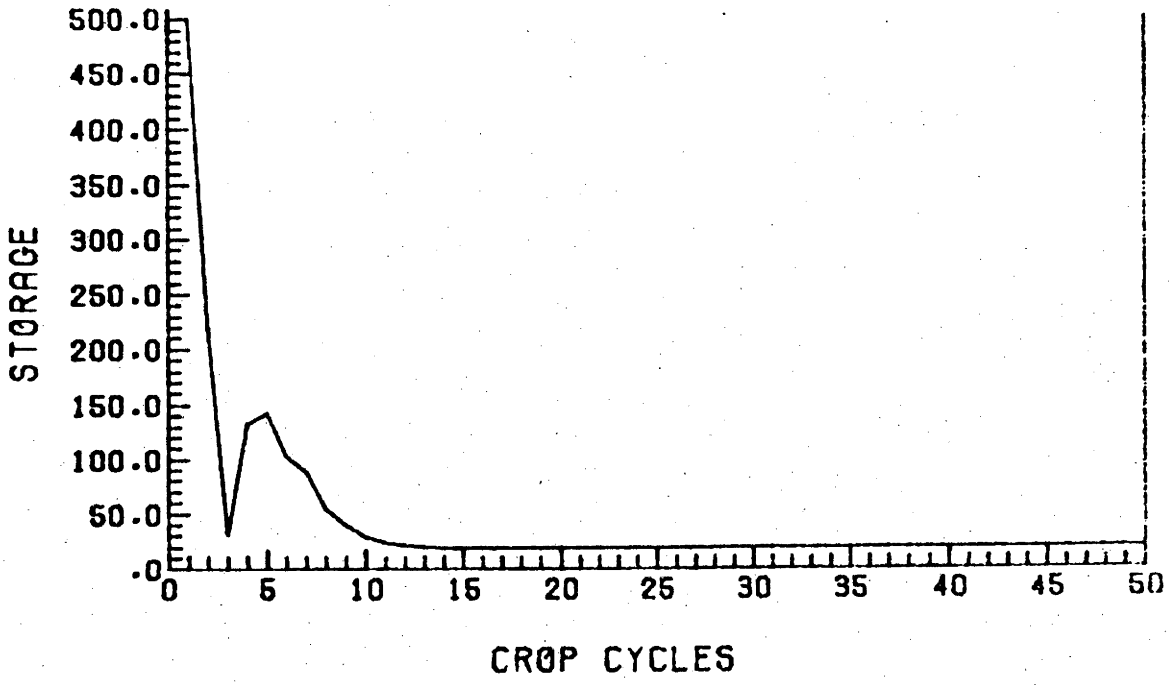

FIG 6.4.90: AGRIBUSINESS + STABILISING SELECTION: SUPERRACE VS TIME

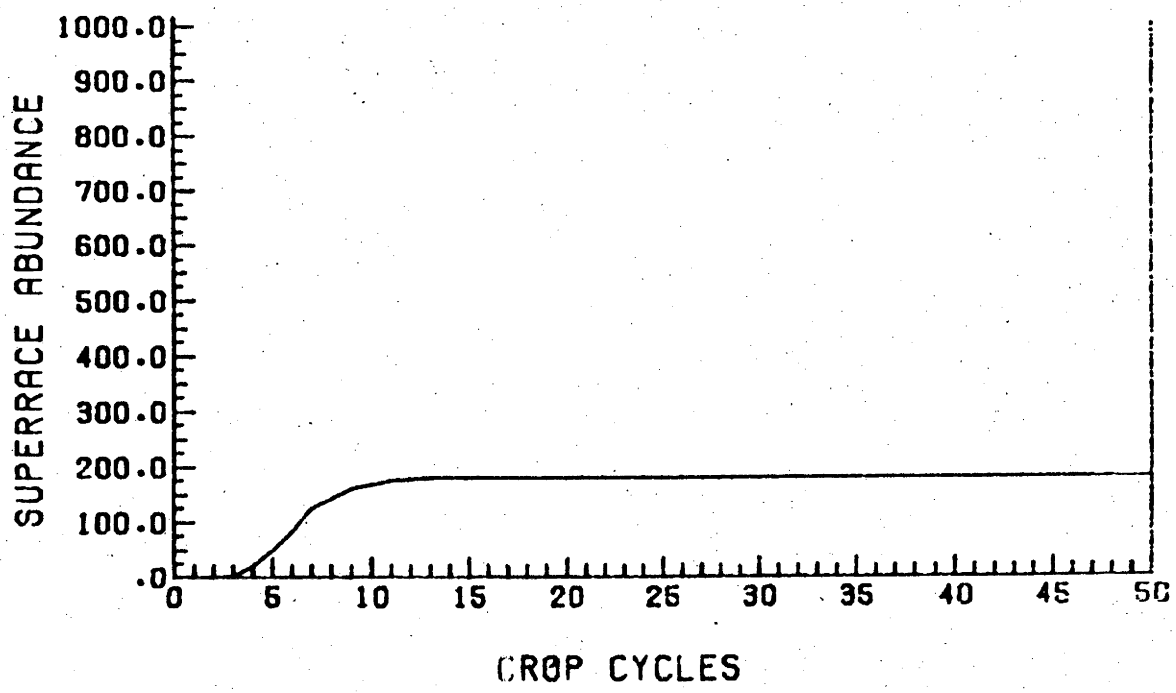

FIG. 6.4.90.1A: AGRIBUSINESS EQUILIBRIUM SEARCH: YIELD VS TIME

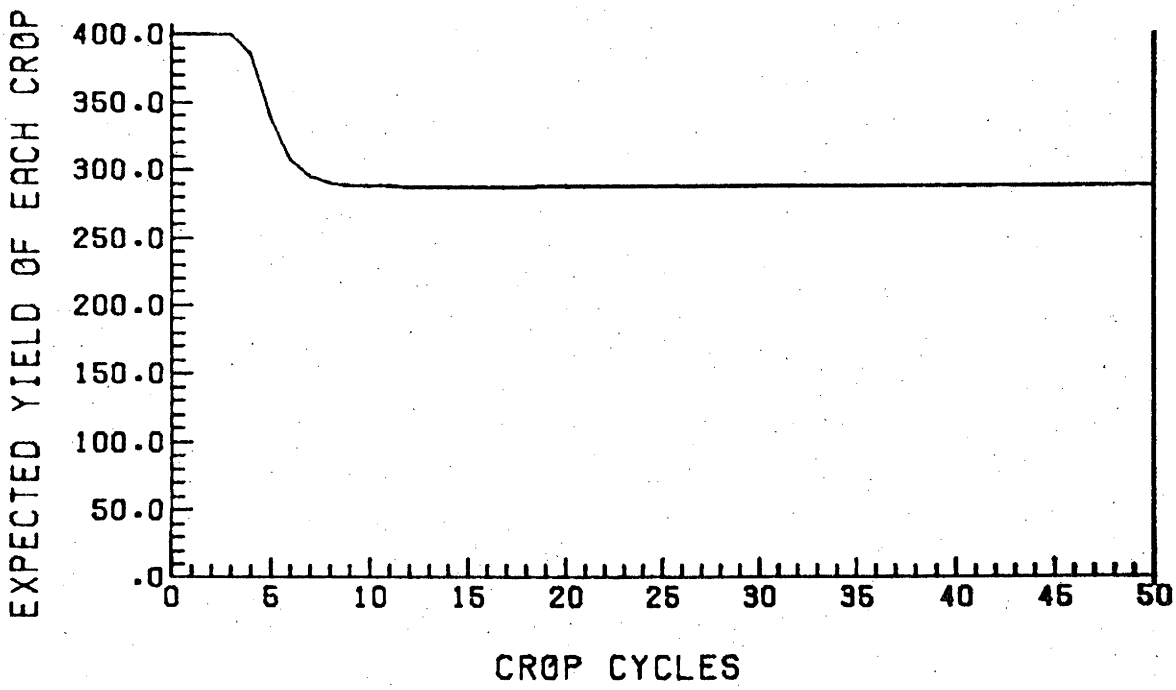

FIG 6.4.90.1B: AGRIBUSINESS EQUILIBRIUM SEARCH: SUPERRACE VS TIME

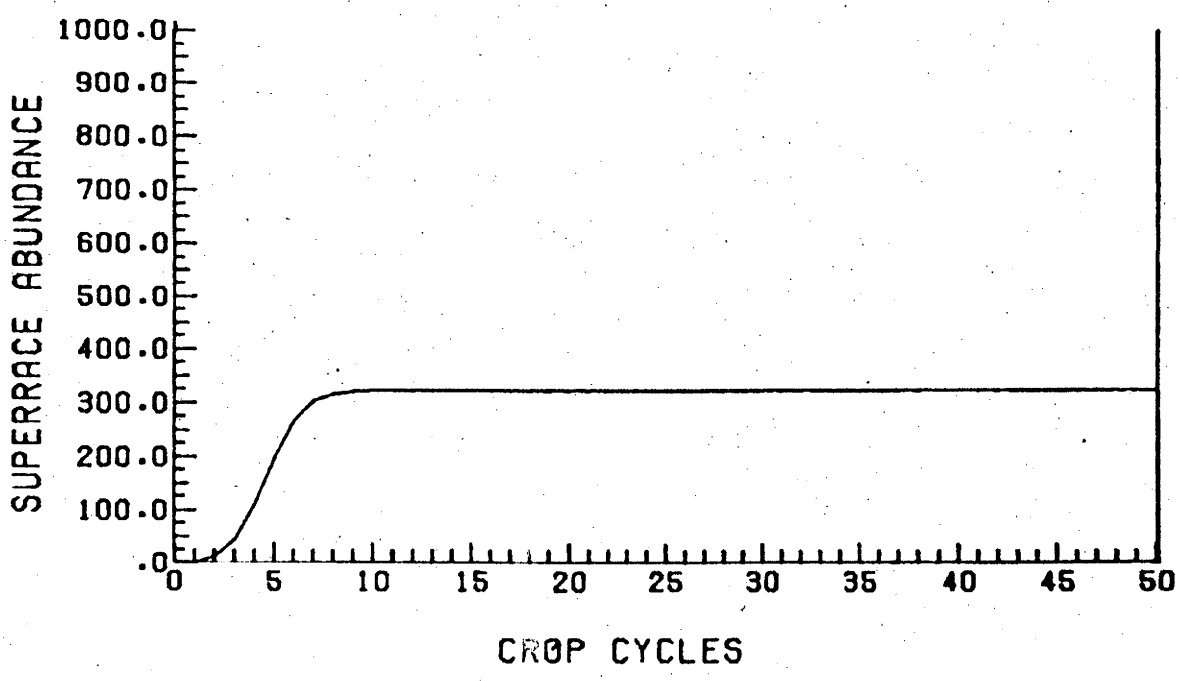

FIG 6.4.90.1C: AGRIBUSINESS EQUILIBRIUM SEARCH: SUBRACE VS TIME

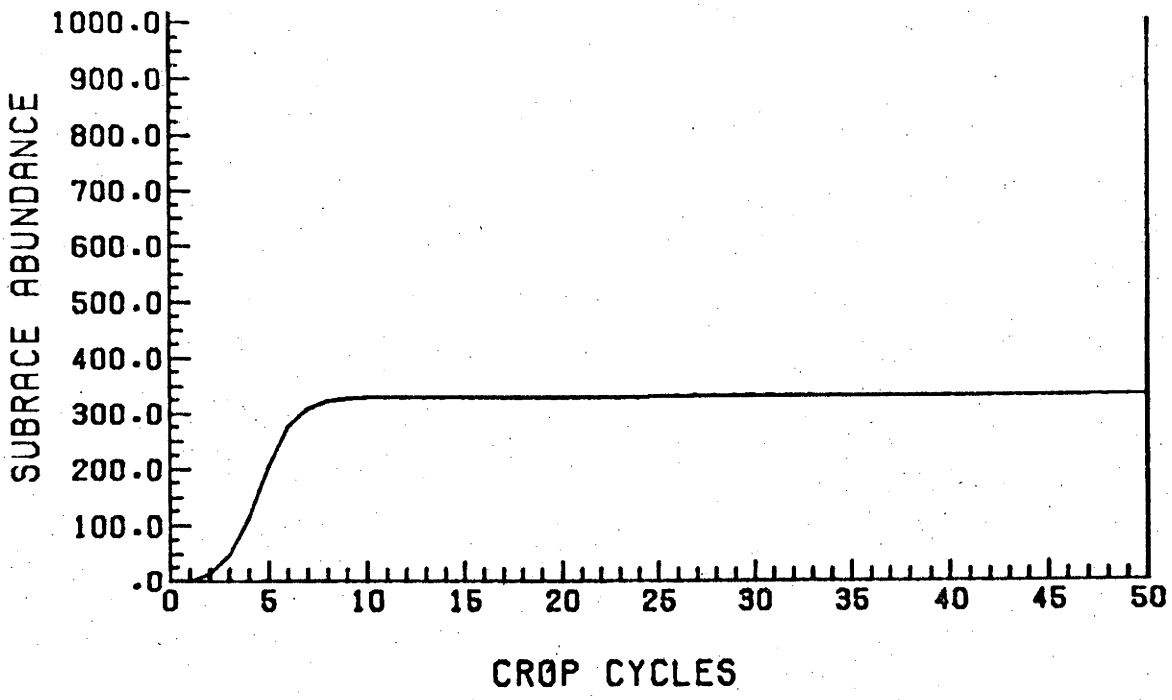

FIG. 6.4.90.1D: AGRIBUSINESS EQUILIBRIUM SEARCH: RETURN VS TIME

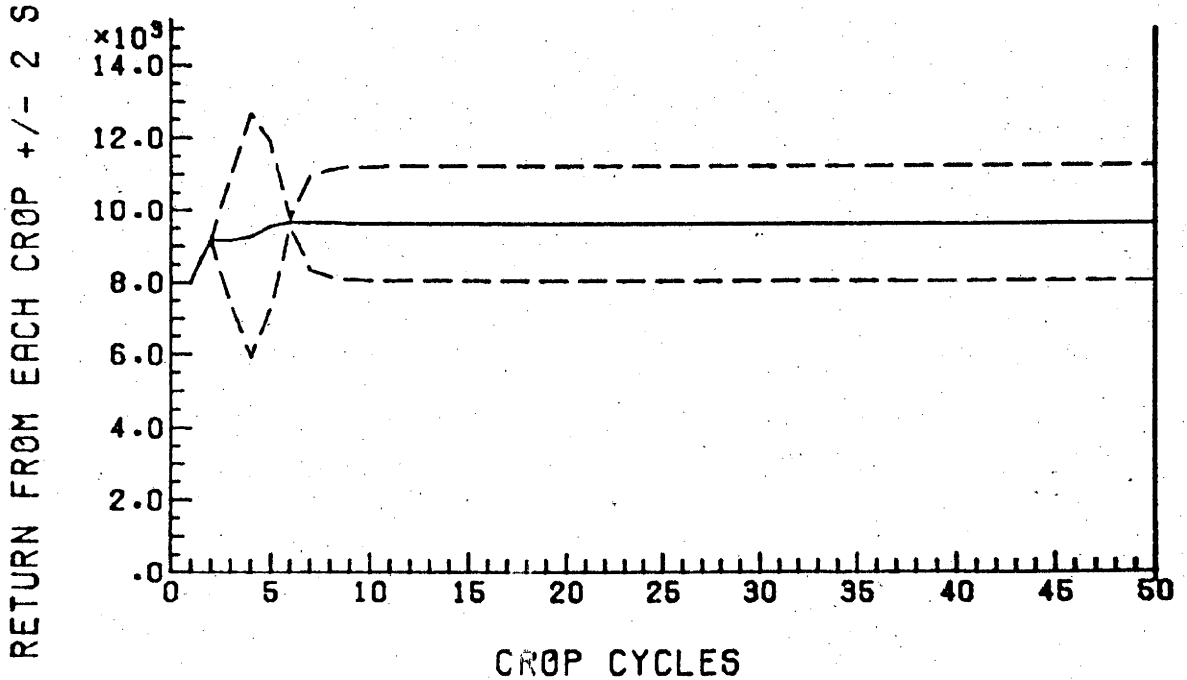

FIG. 6.4.91A: LINE OF PARAMETER VALUES SATISFYING THE YIELD DEPRESSION CRITERION

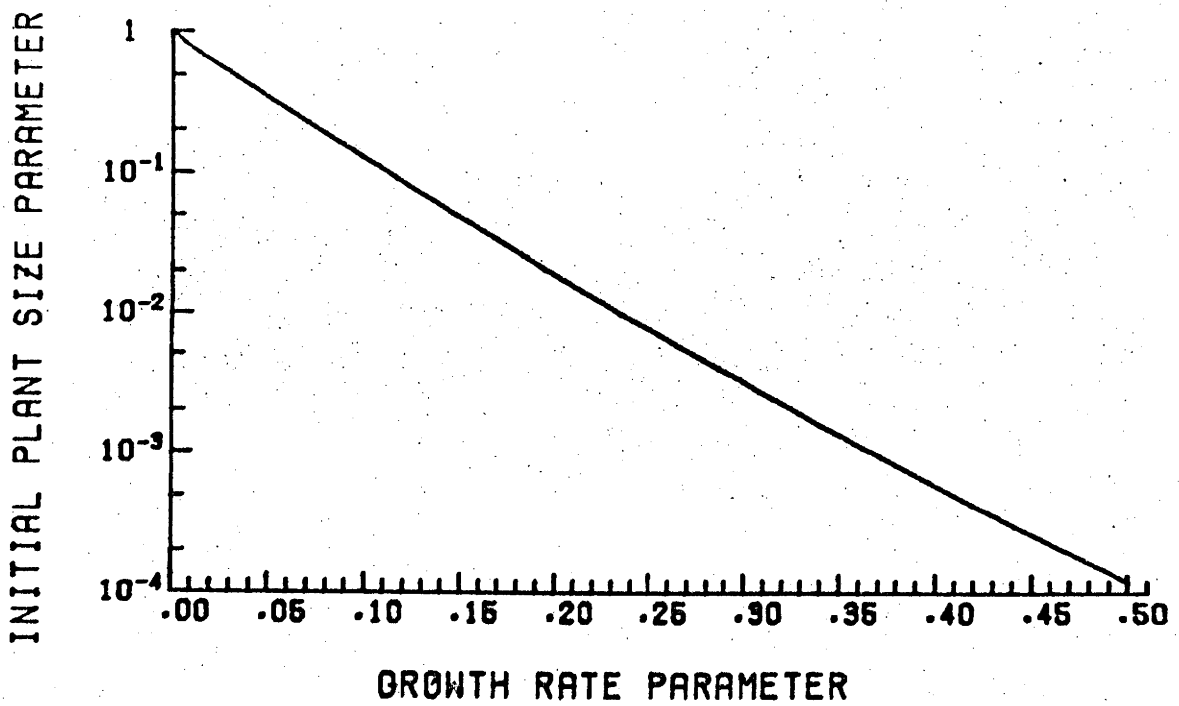

FIG. 6.4.91B: NATURAL INCREASE OF PATHOGEN ON GROWING PLANT

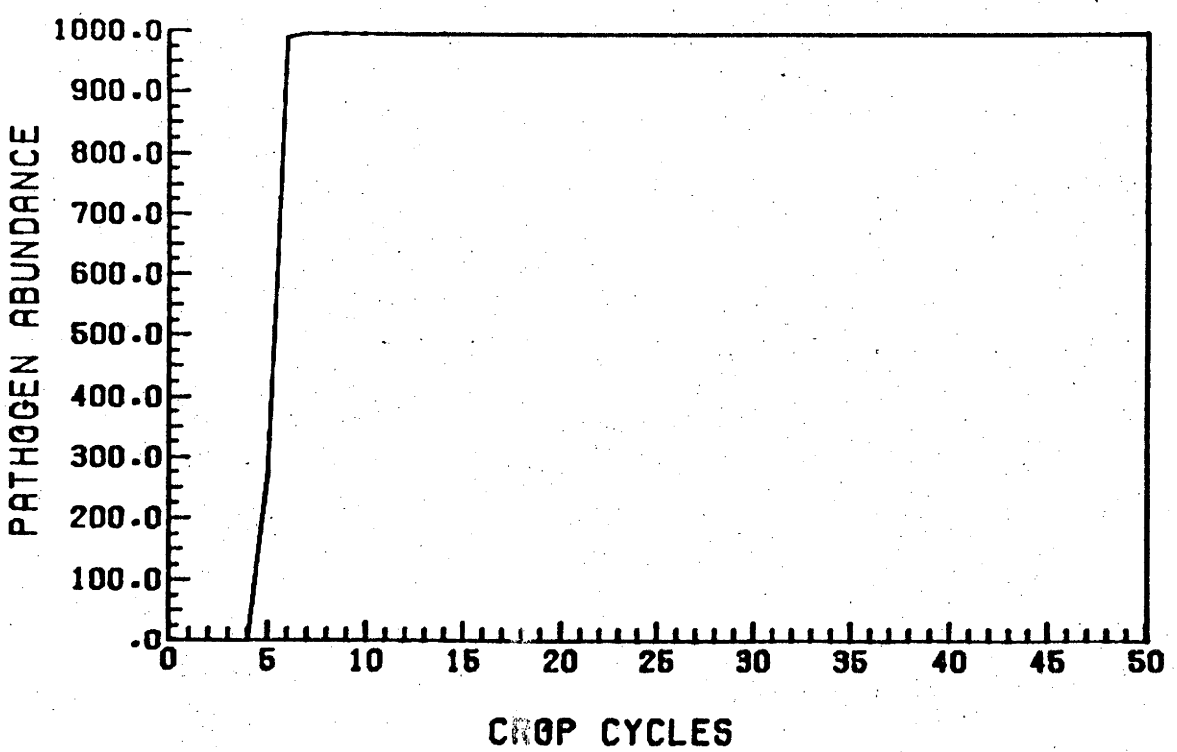

FIG. 6.4.91C: GROWING PLANT
DEPENDENCE OF YIELD ON PATHOGEN INPUT

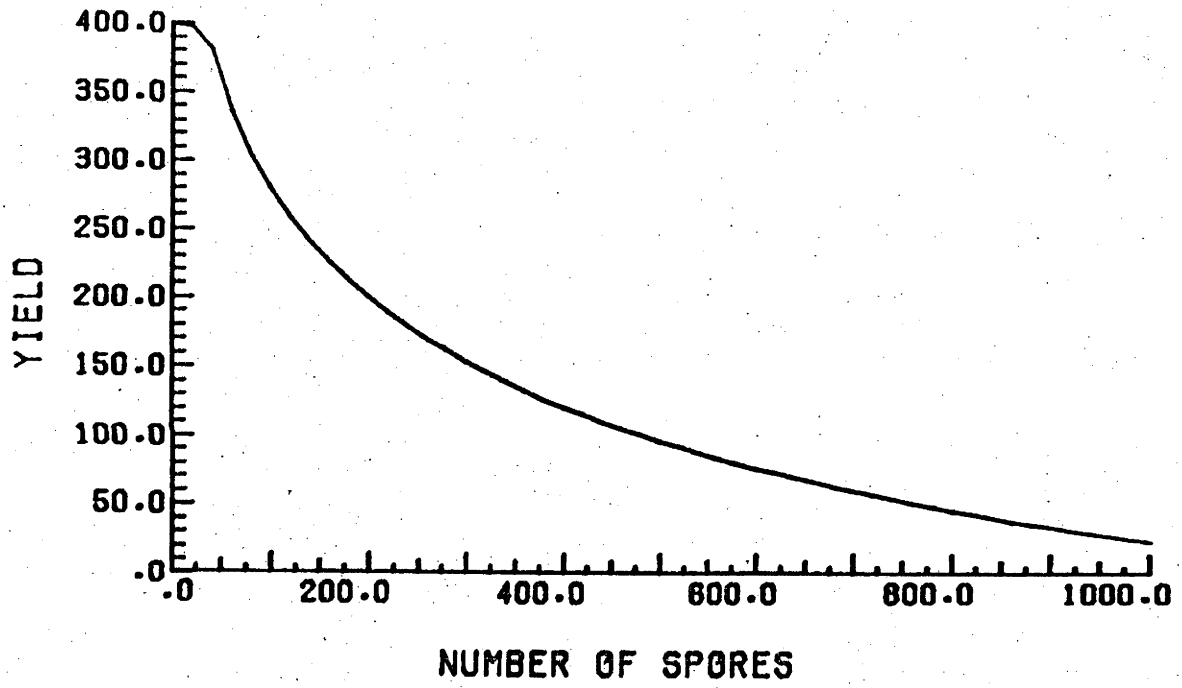

FIG 6.4.93A: AGRIBUSINESS WITH GROWING PLANT AREA PLANTED VS STORAGE

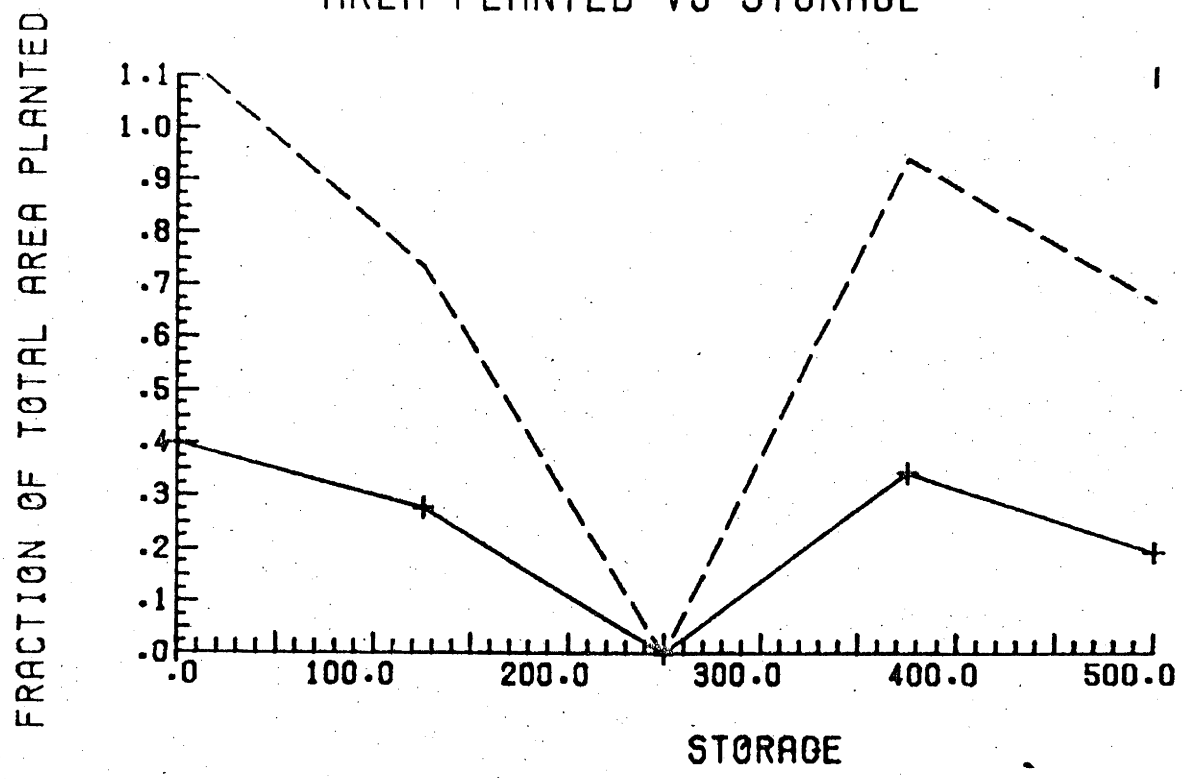

FIG 6.4.93B: AGRIBUSINESS WITH GROWING PLANT AREA PLANTED VS AABB

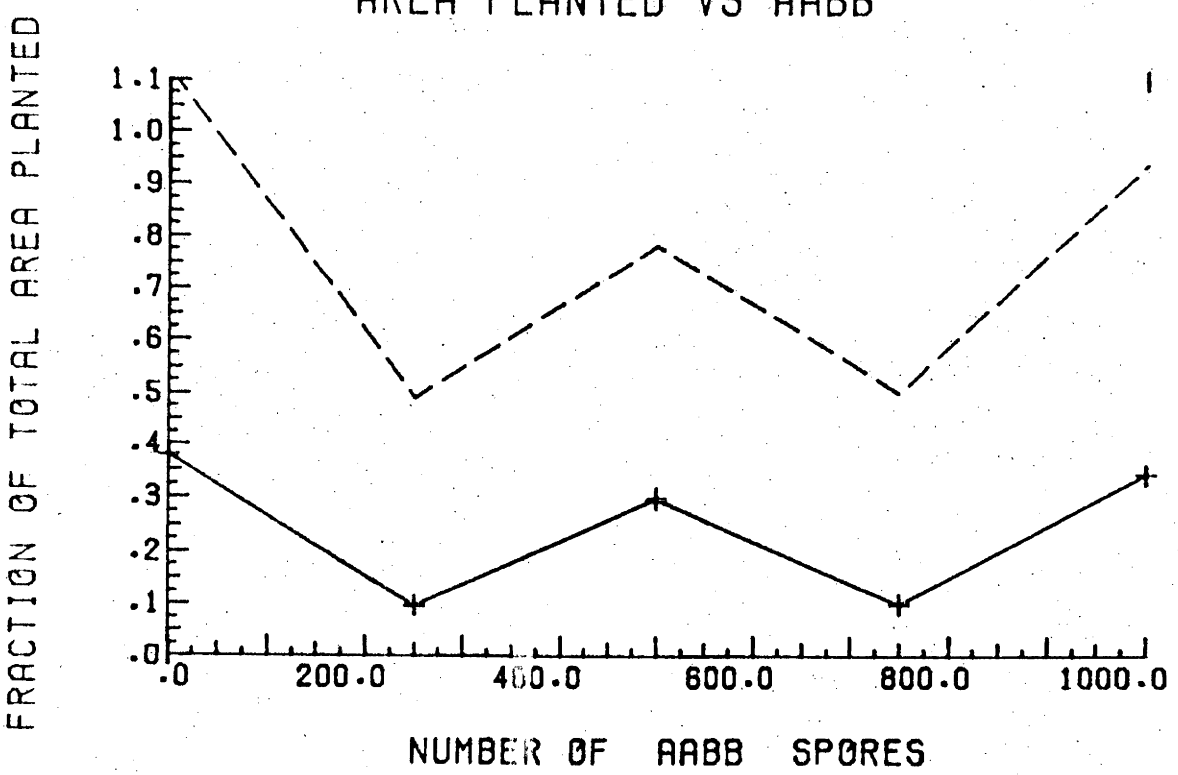

FIG 6.4.94A: AGRIBUSINESS WITH GROWING PLANT
E(RETURN) VS STORAGE

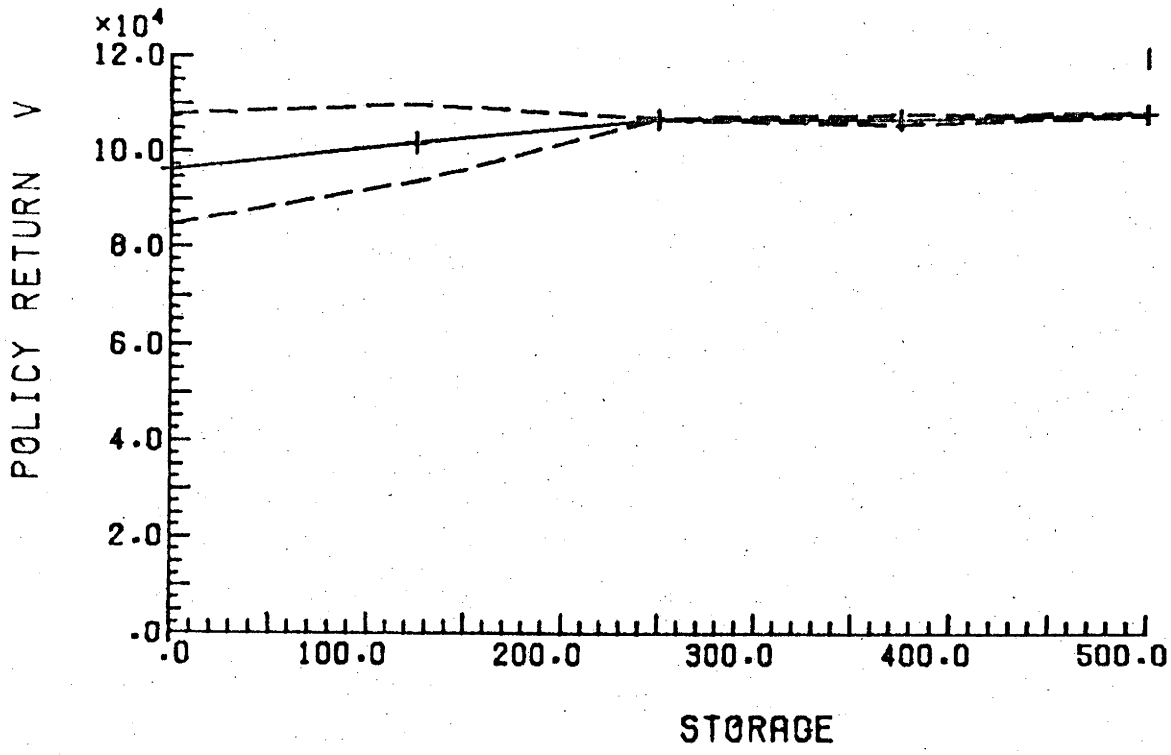

FIG 6.4.94B: AGRIBUSINESS WITH GROWING PLANT
E(RETURN) VS AABB

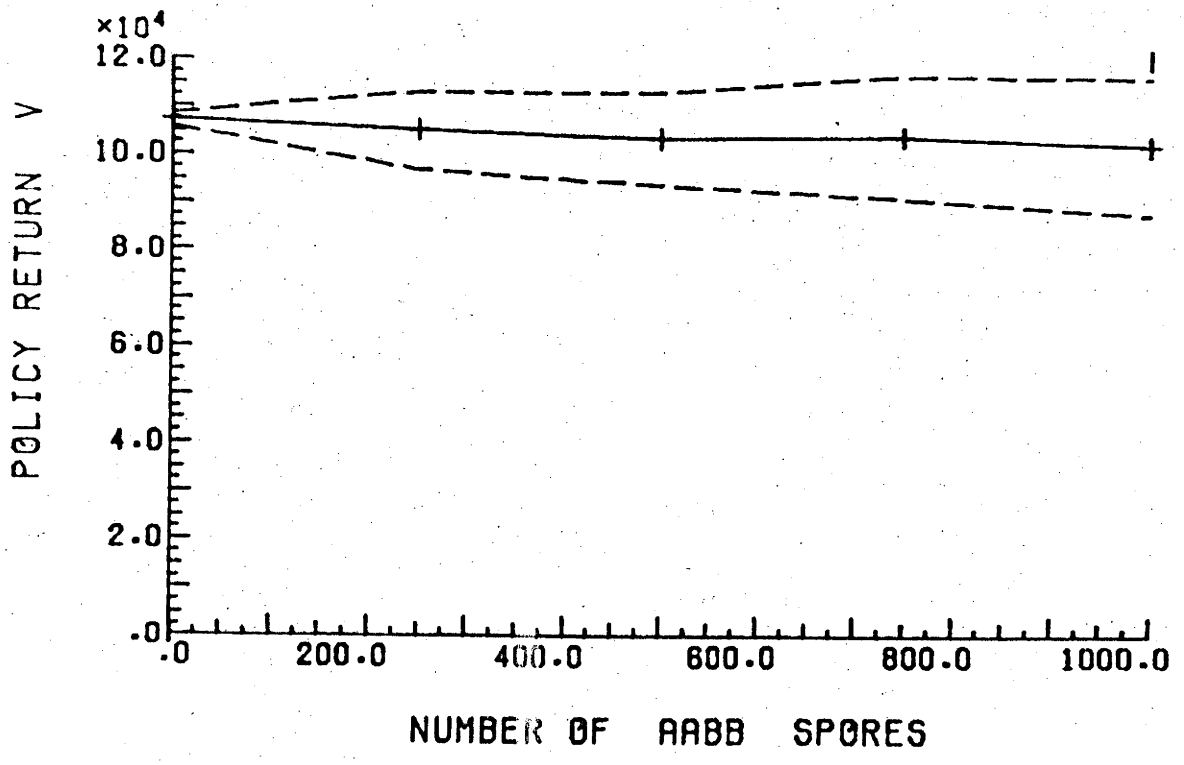

FIG 6.4.95A: AGRIBUSINESS WITH GROWING PLANT
STORAGE OUTPUT VS STORAGE

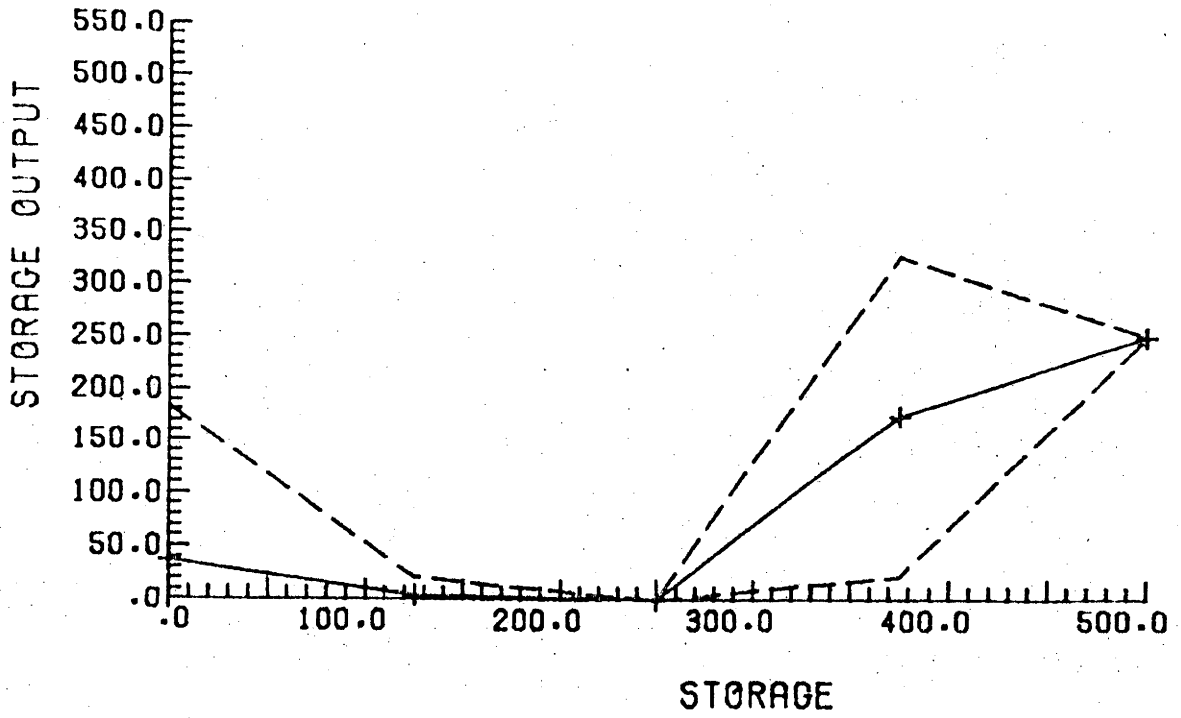

FIG 6.4.95B: AGRIBUSINESS WITH GROWING PLANT
STORAGE OUTPUT VS AABB

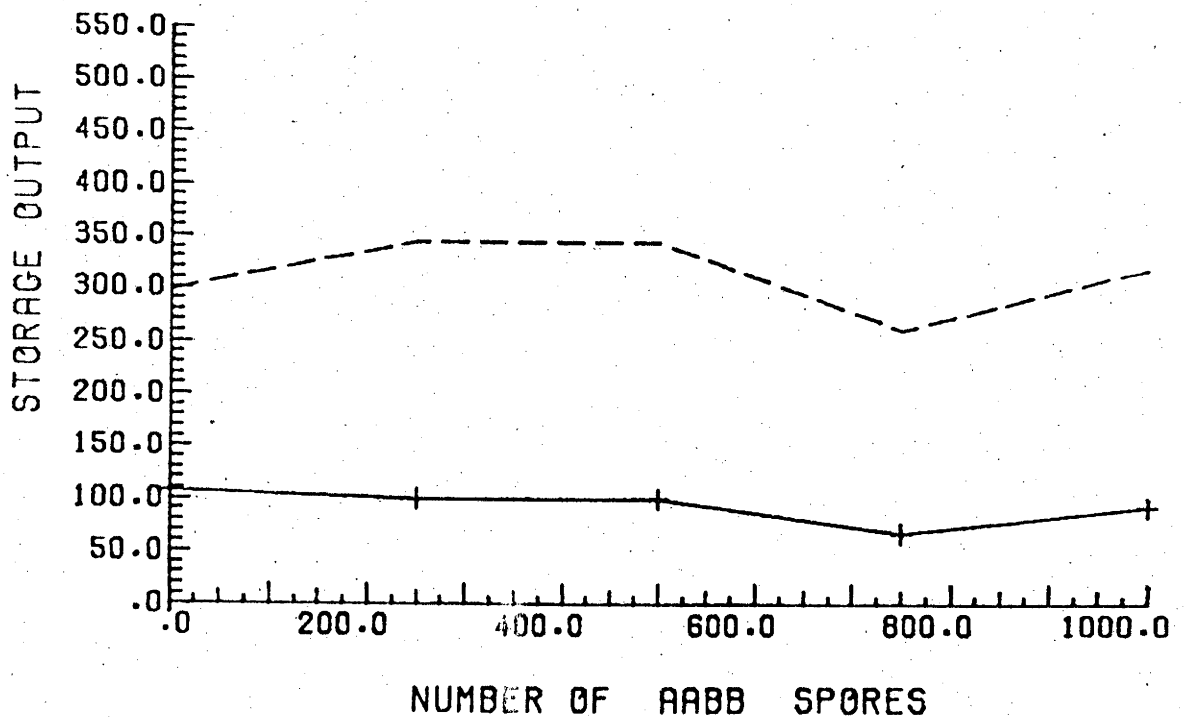

FIG 6.4.96A: AGRIBUSINESS WITH GROWING PLANT
AABB OUTPUT VS STORAGE

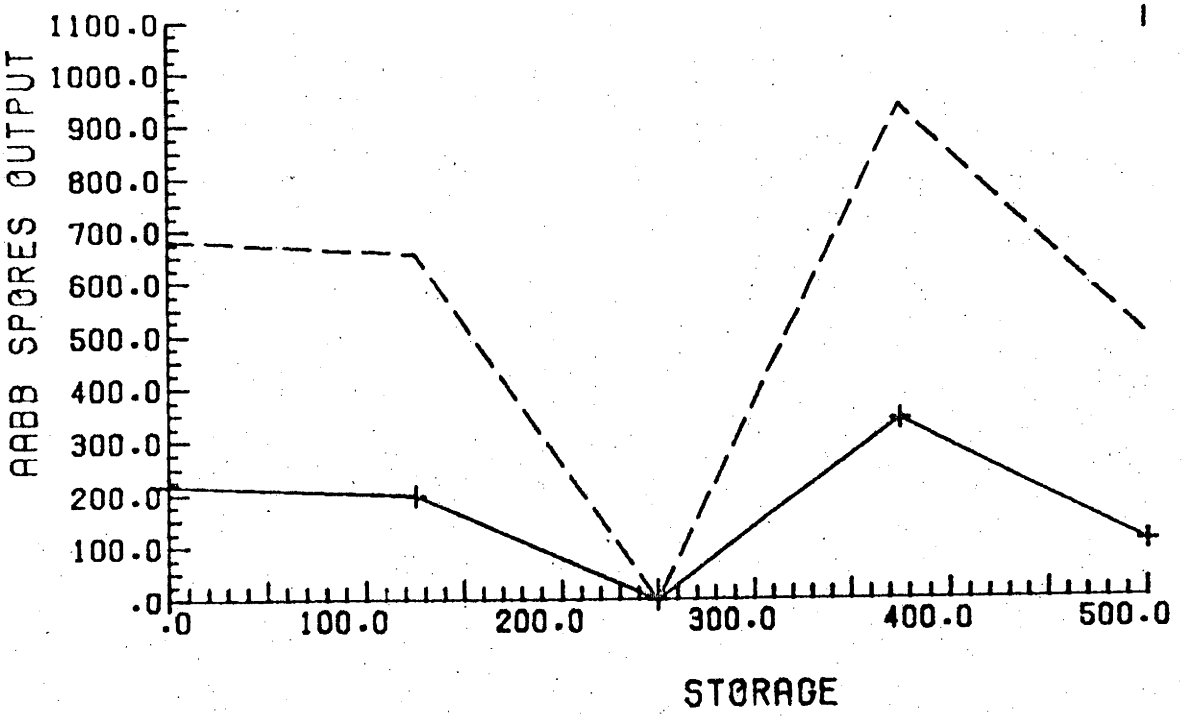

FIG 6.4.96B: AGRIBUSINESS WITH GROWING PLANT
AABB OUTPUT VS AABB

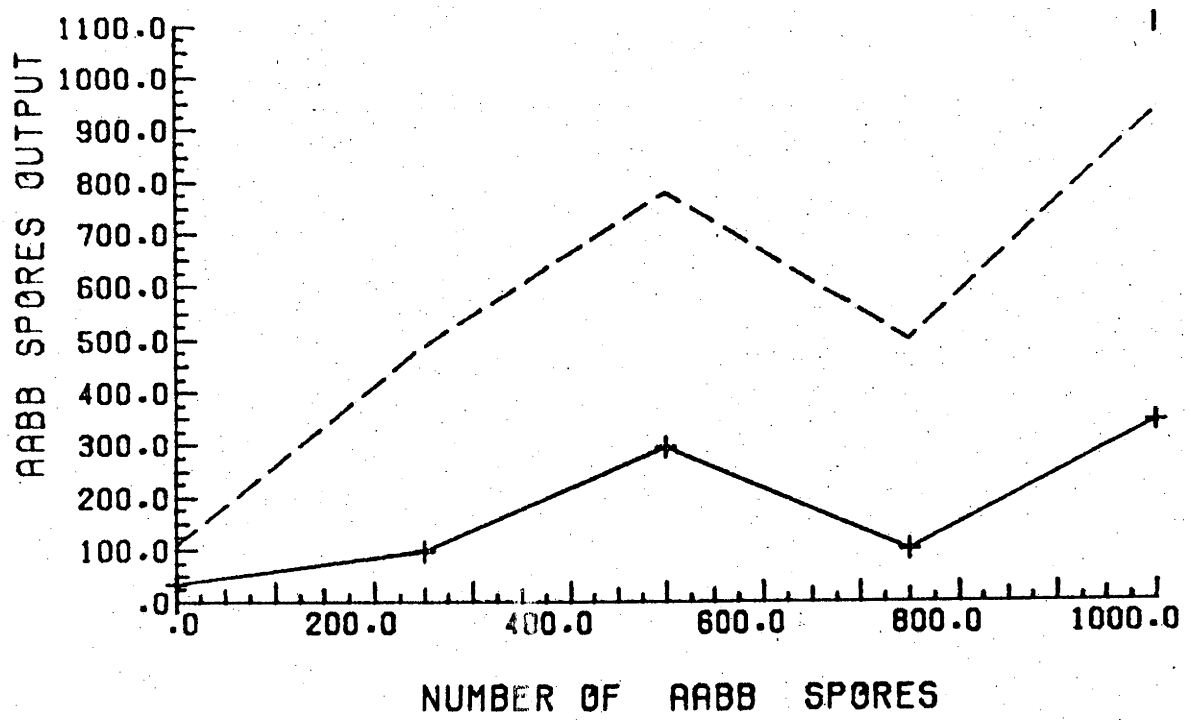

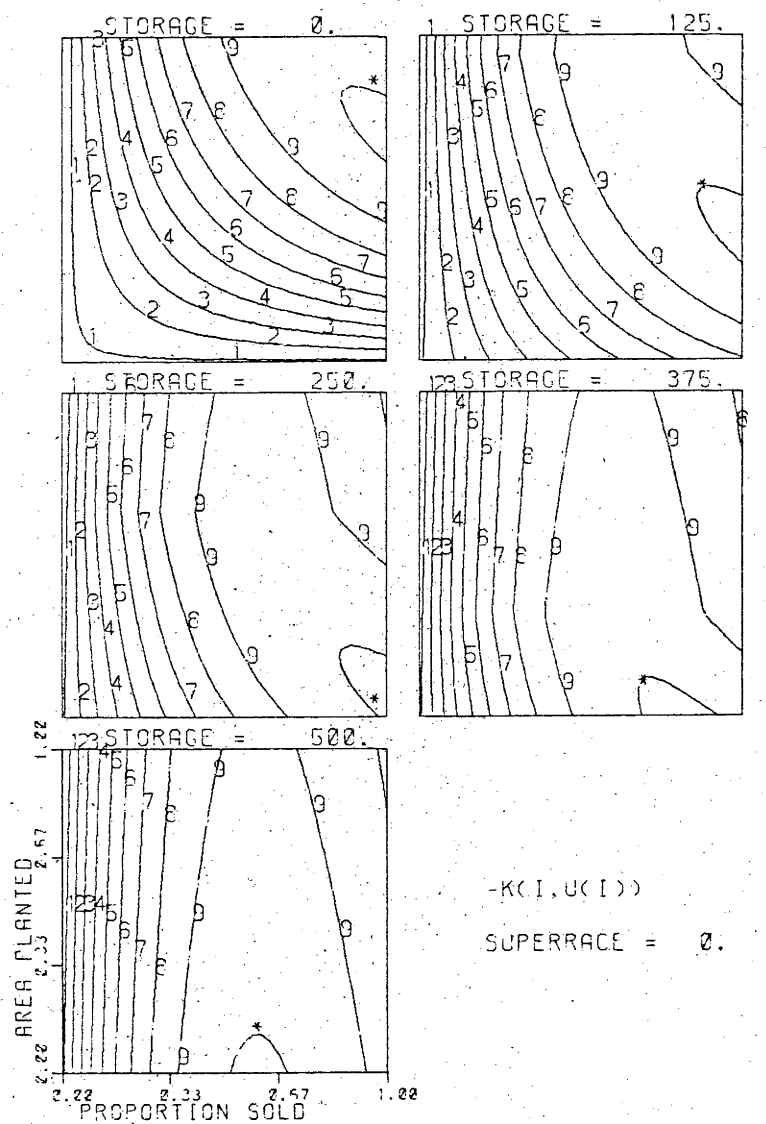

FIG. 5.4.97A: AGRIBUSINESS WITH GROWING PLANT

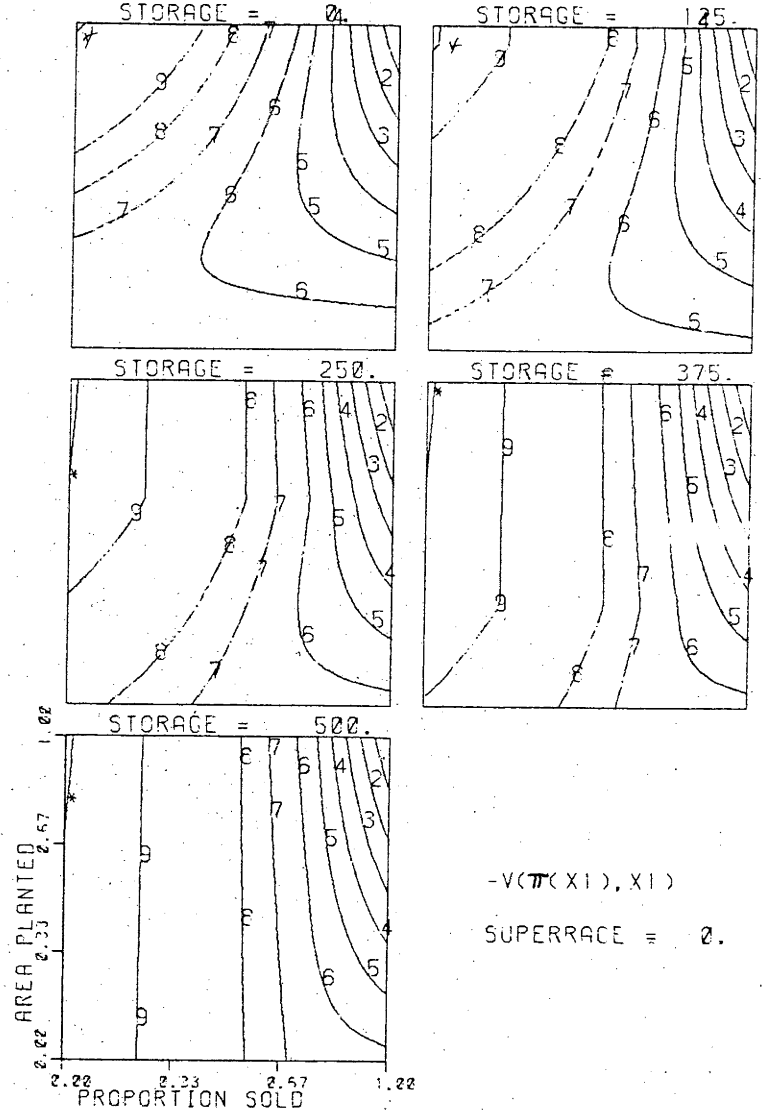

FIG. 5.4.97B: AGRIBUSINESS WITH GROWING PLANT

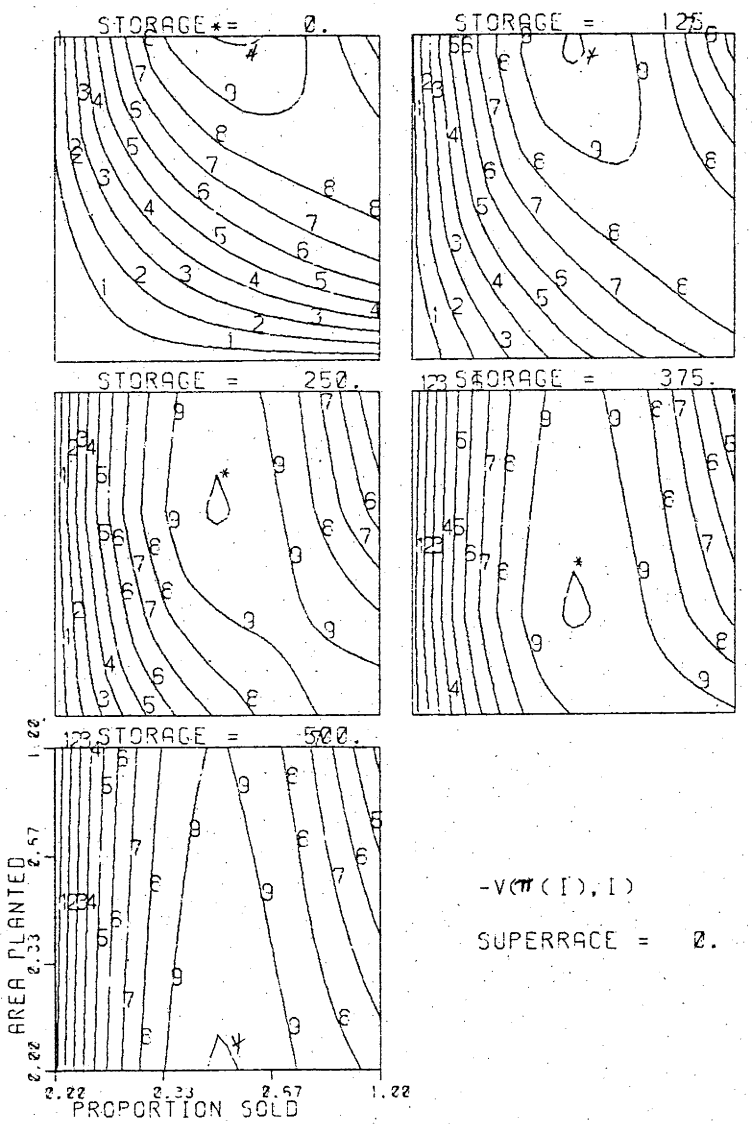

FIG. 5.4.97C: AGRIBUSINESS WITH GROWING PLANT

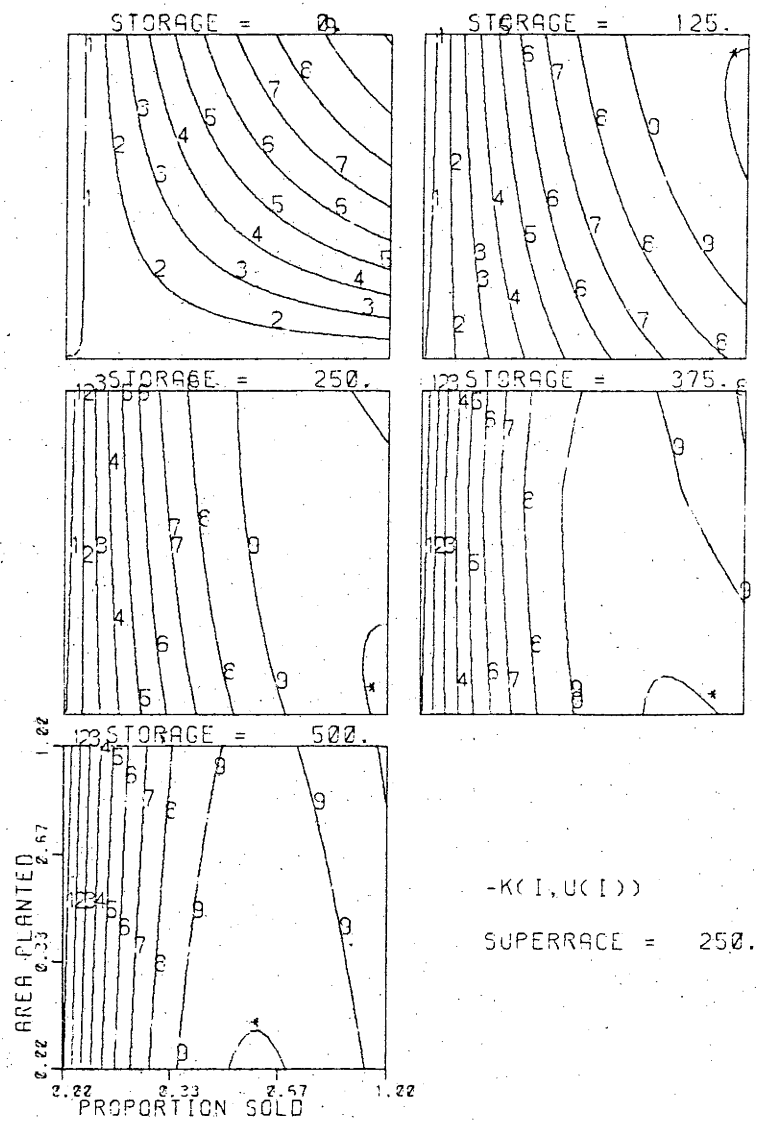

FIG. 5.4.98A: AGRIBUSINESS WITH GROWING PLANT

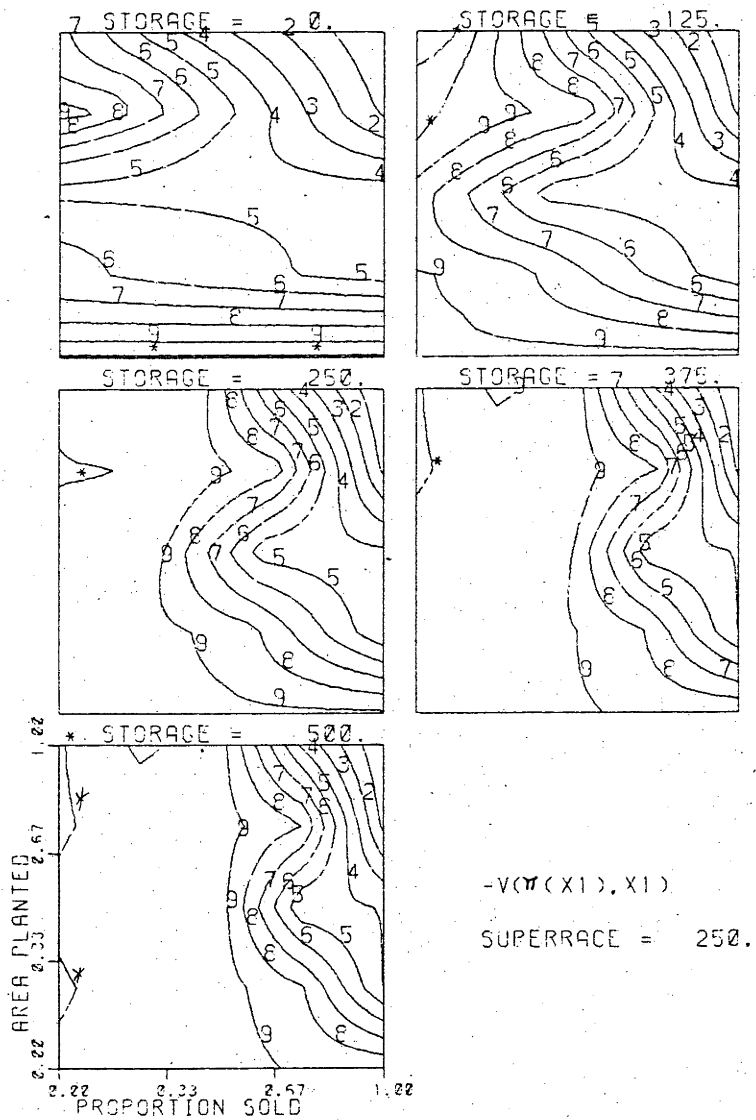

FIG. 5.4.98B: AGRIBUSINESS WITH GROWING PLANT

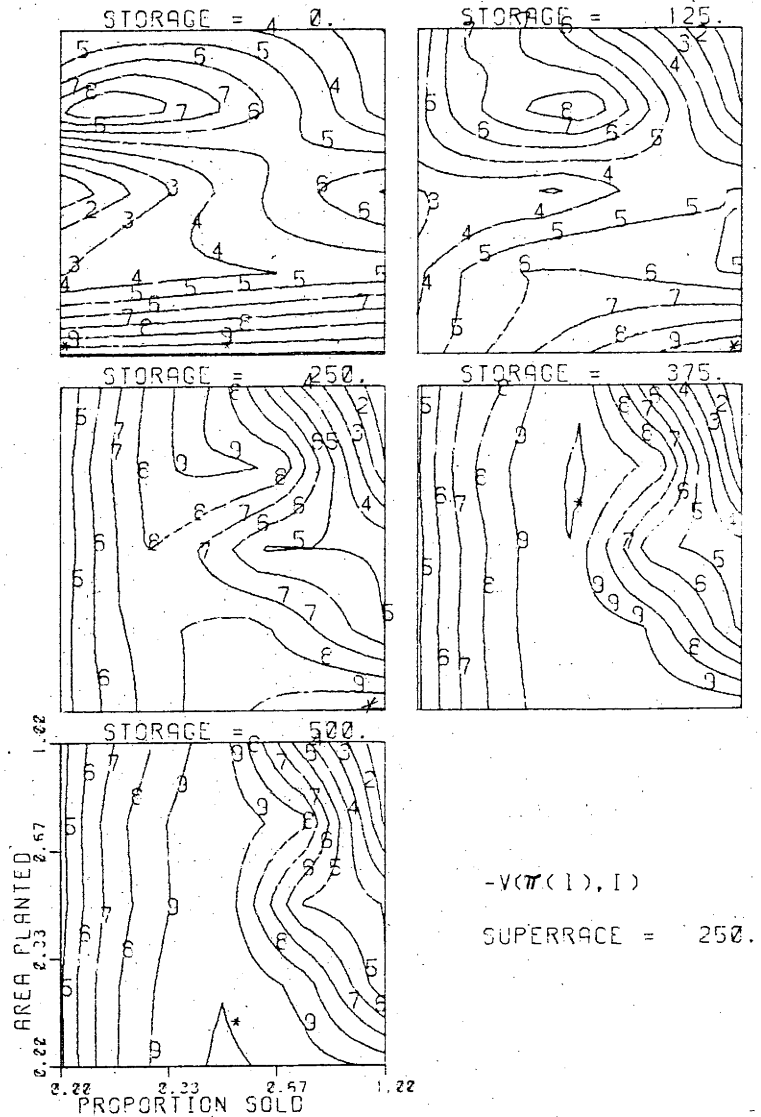

FIG. 5.4.98C: AGRIBUSINESS WITH GROWING PLANT

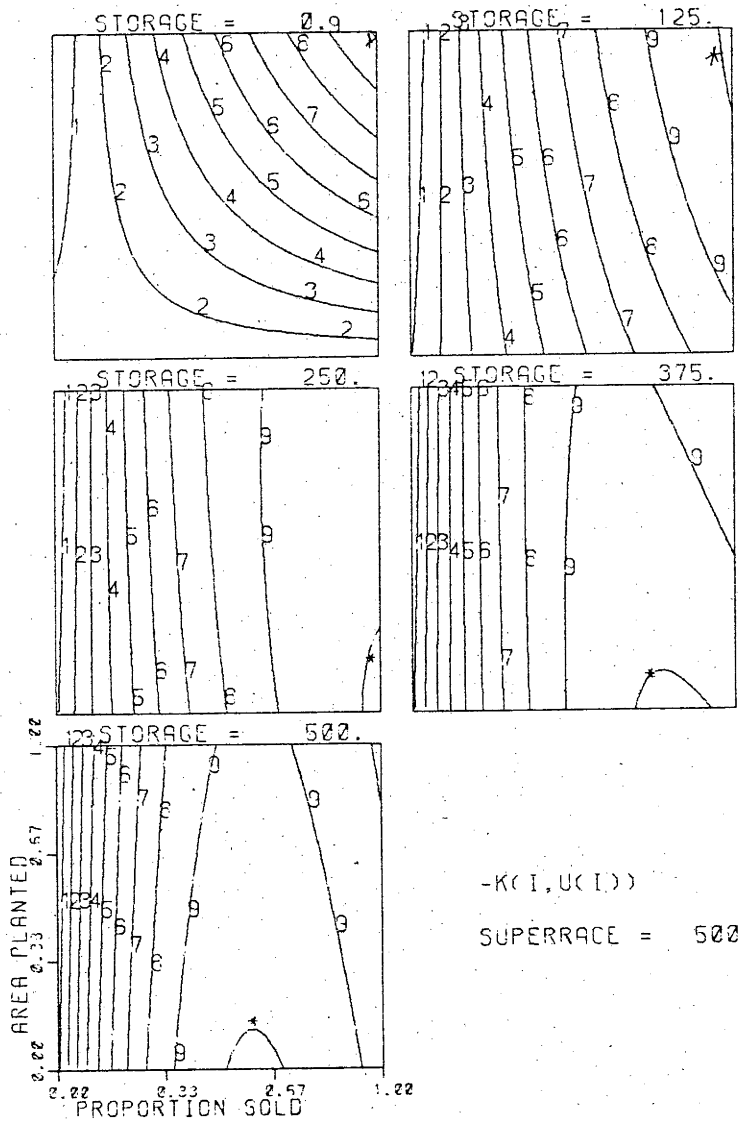

FIG. 6.4.99A: AGRIBUSINESS WITH GROWING PLANT

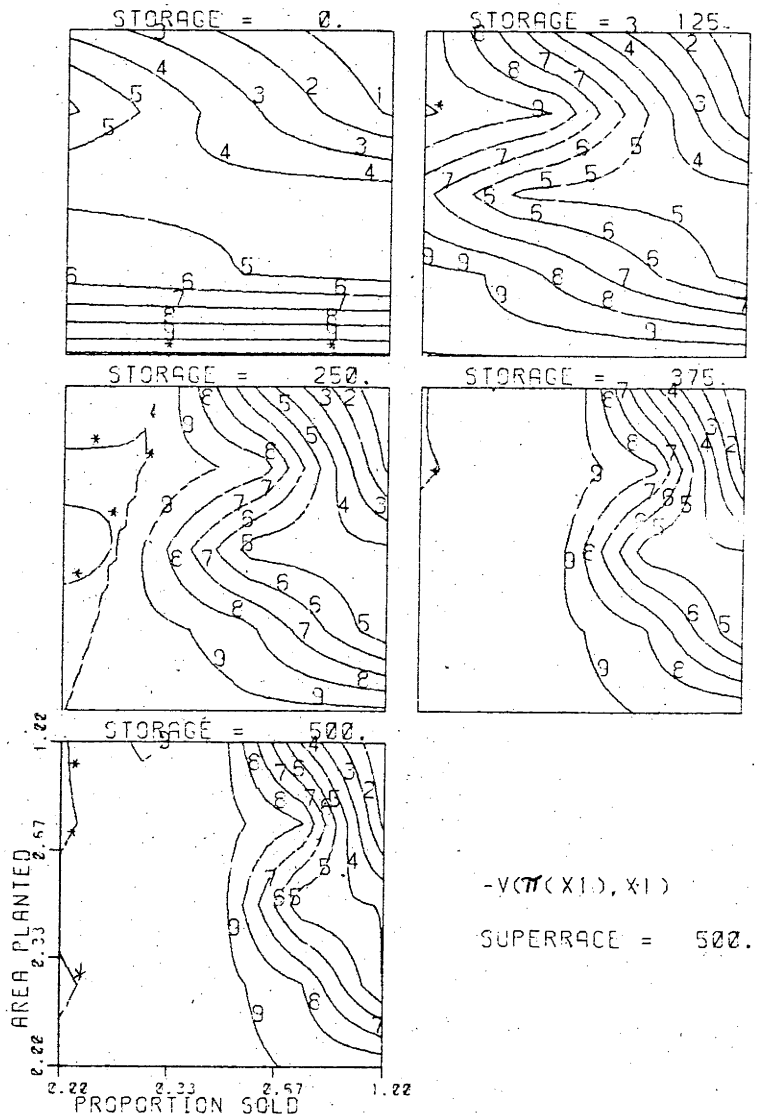

FIG. 6.4.99B: AGRIBUSINESS WITH GROWING PLANT

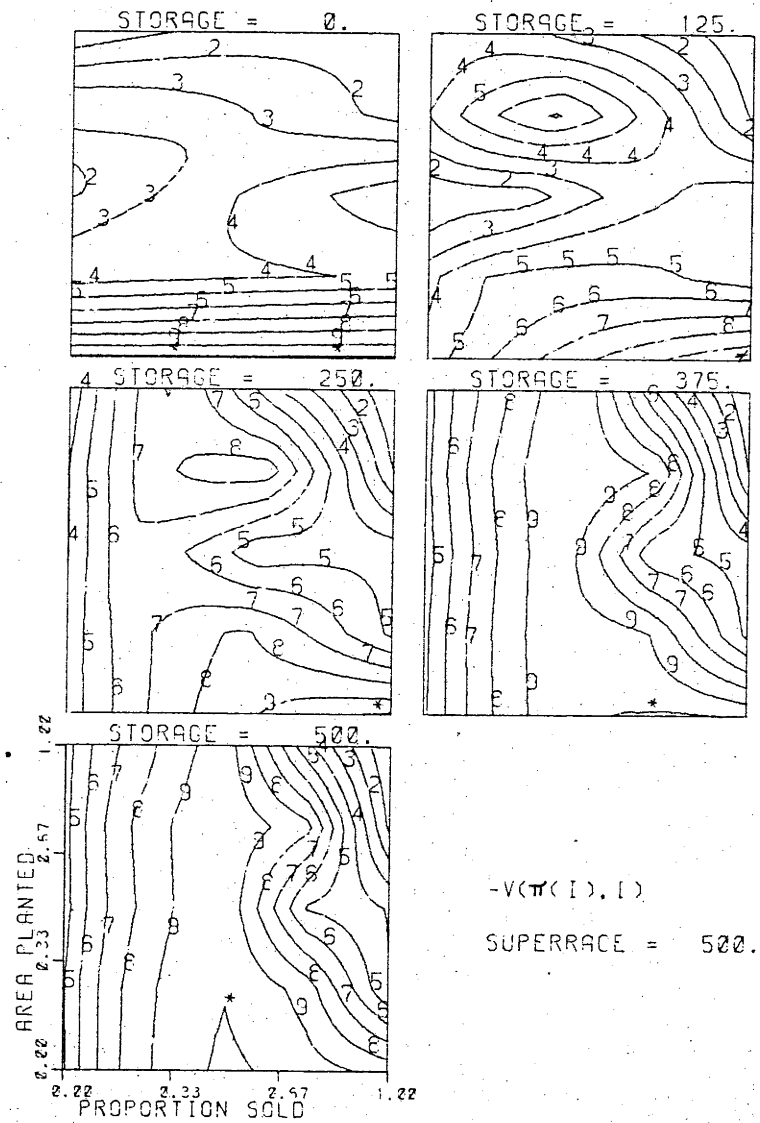

-V(π(I),I)

SUPERRACE = 500.

FIG. 5.4.99C: AGRIBUSINESS WITH GROWING PLANT

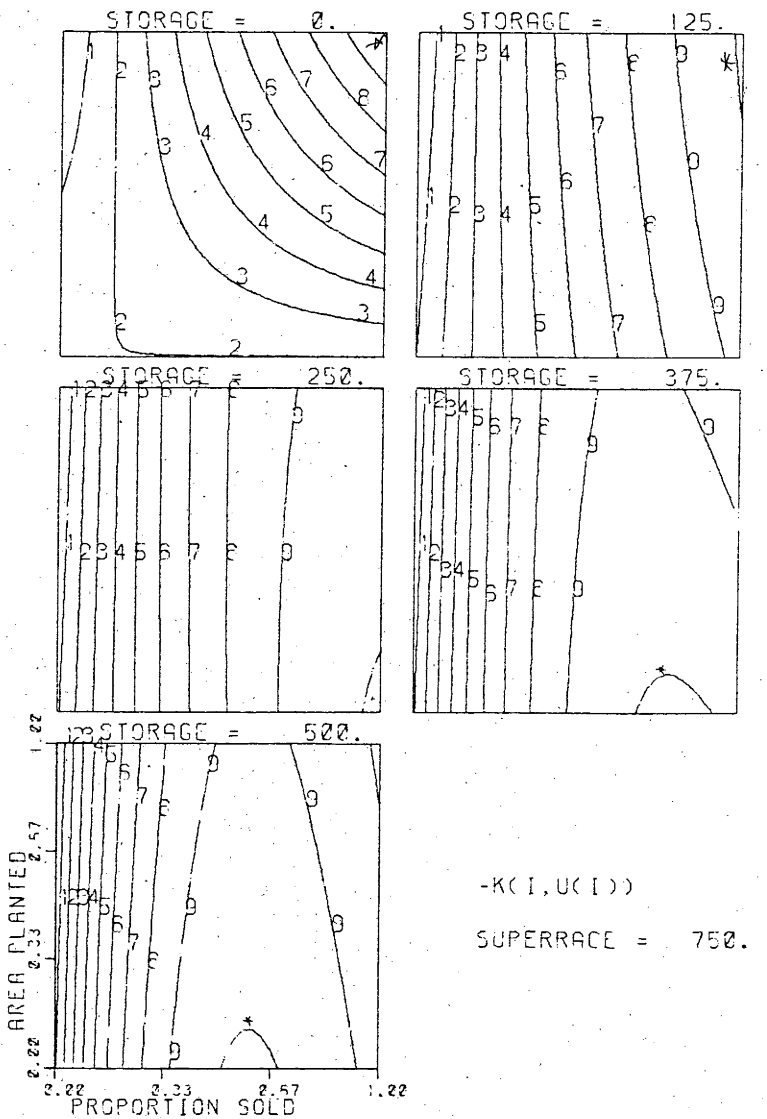

-K(I,U(I))

SUPERRACE = 750.

FIG. 5.4.100A: AGRIBUSINESS WITH GROWING PLANT

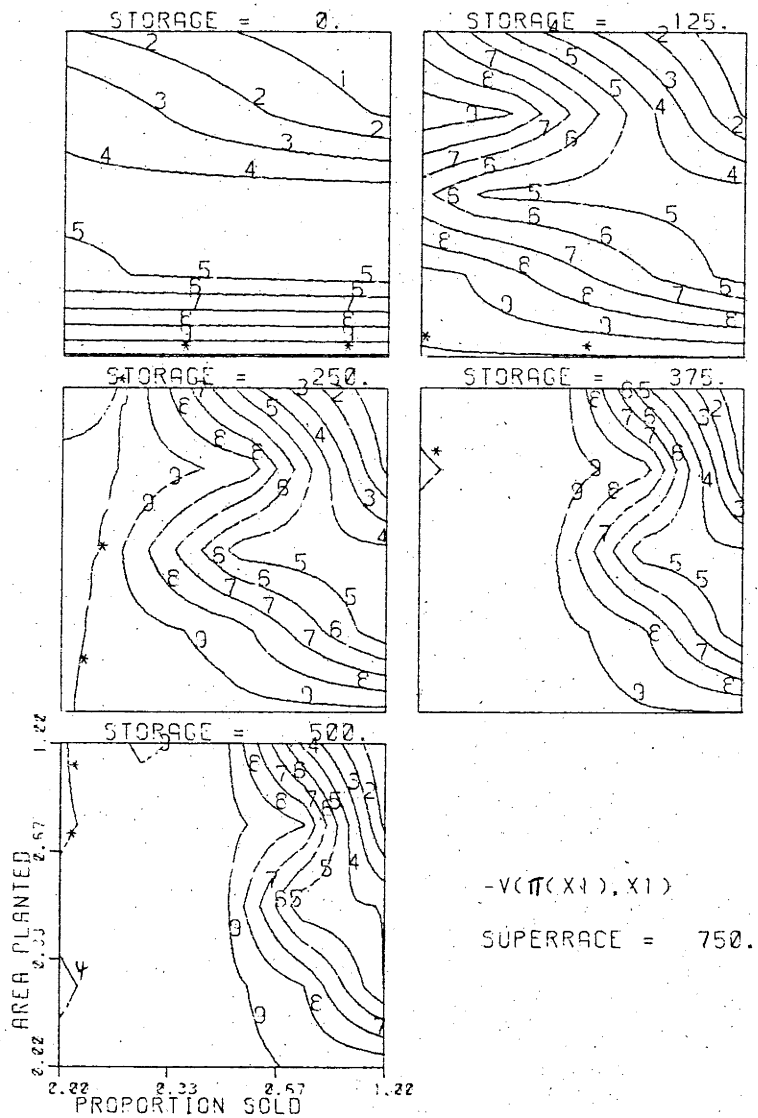

FIG. 6.4.100B: AGRIBUSINESS WITH GROWING PLANT

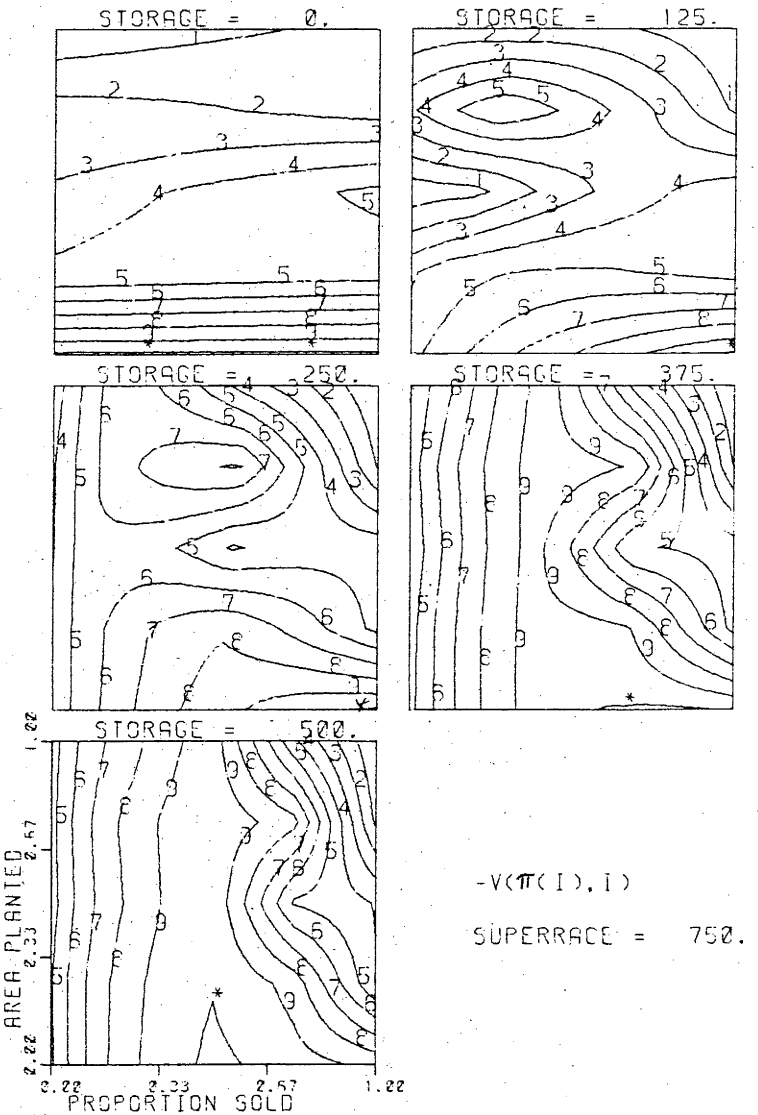

FIG. 6.4.100C: AGRIBUSINESS WITH GROWING PLANT

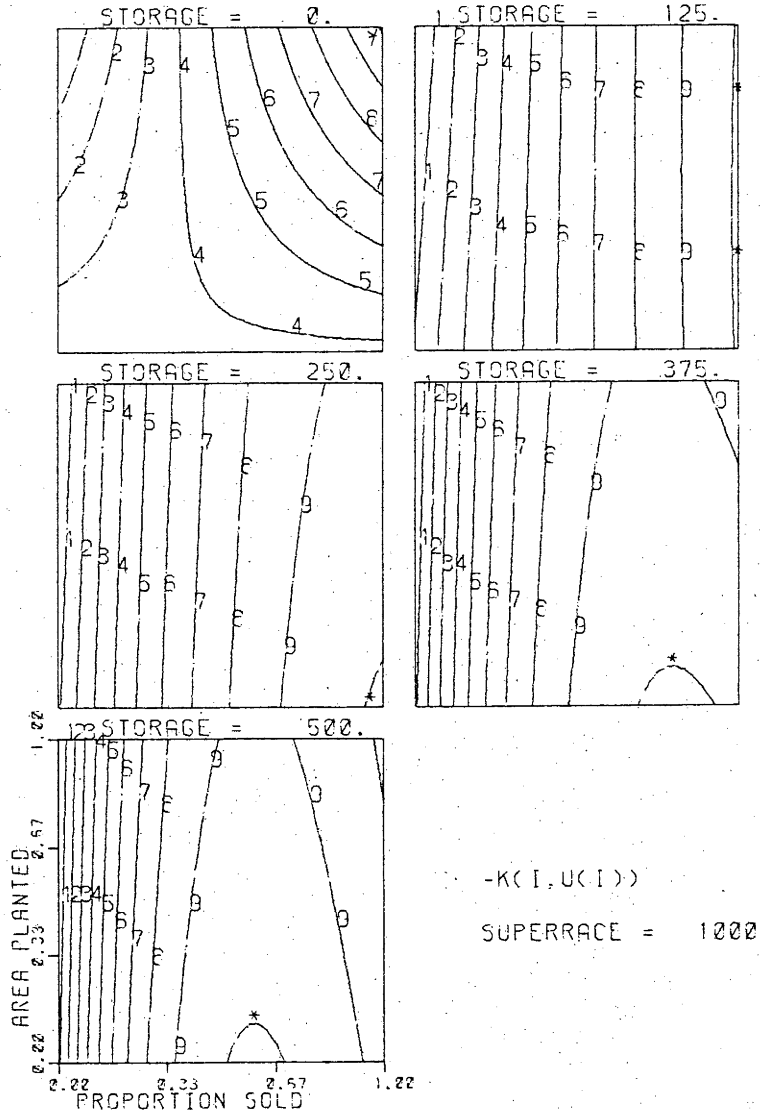

FIG. 5.4.101A: AGRIBUSINESS WITH GROWING PLANT

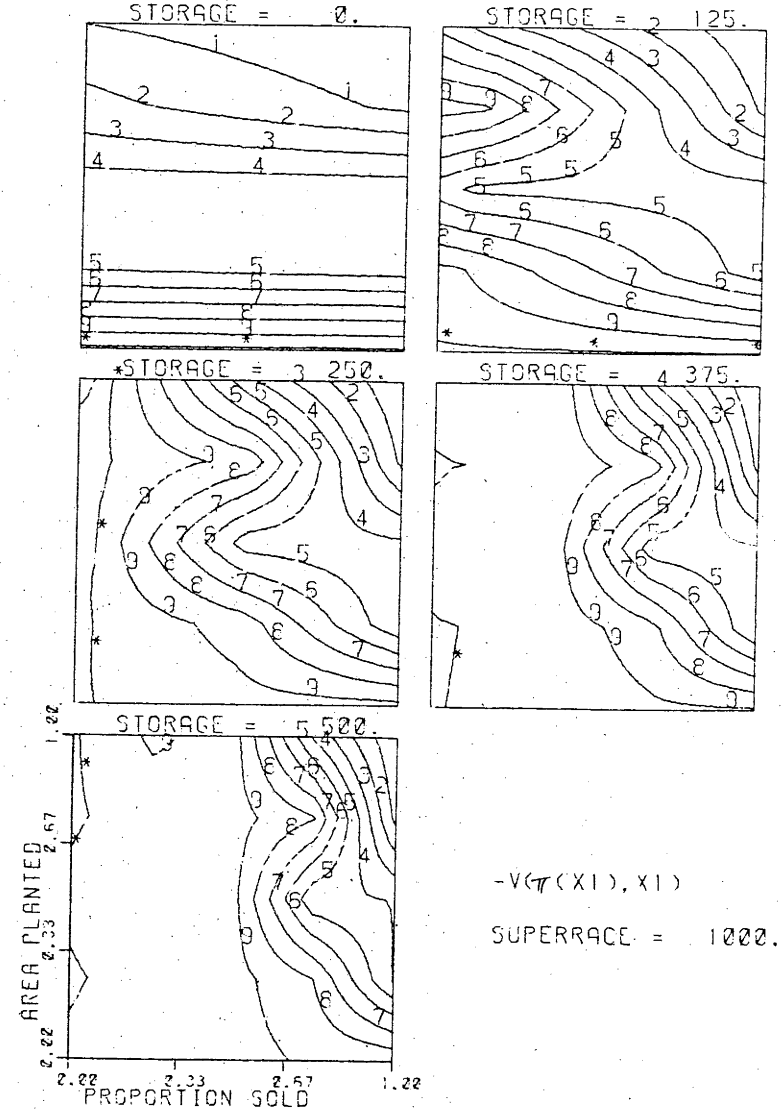

FIG. 5.4.101B: AGRIBUSINESS WITH GROWING PLANT

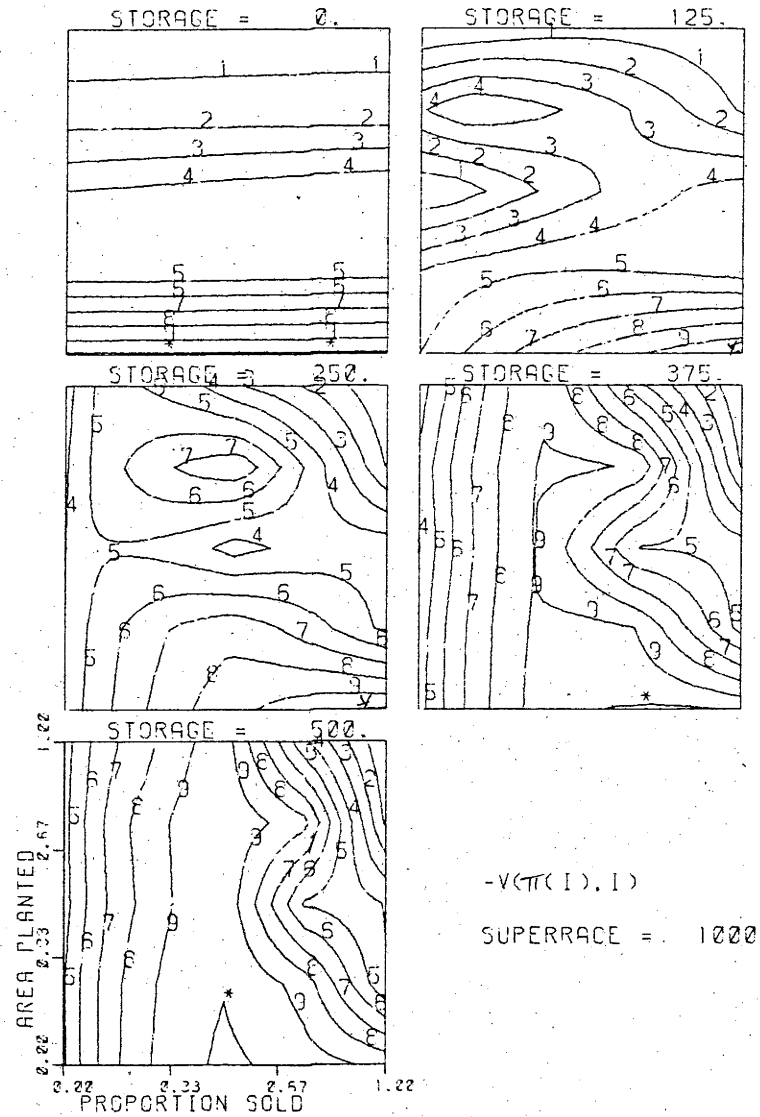

FIG. 6.4.101C: AGRIBUSINESS WITH GROWING PLANT

FIG 6.4.102: AGRIBUSINESS WITH GROWING PLANT FRACTION SOLD THROUGH TIME

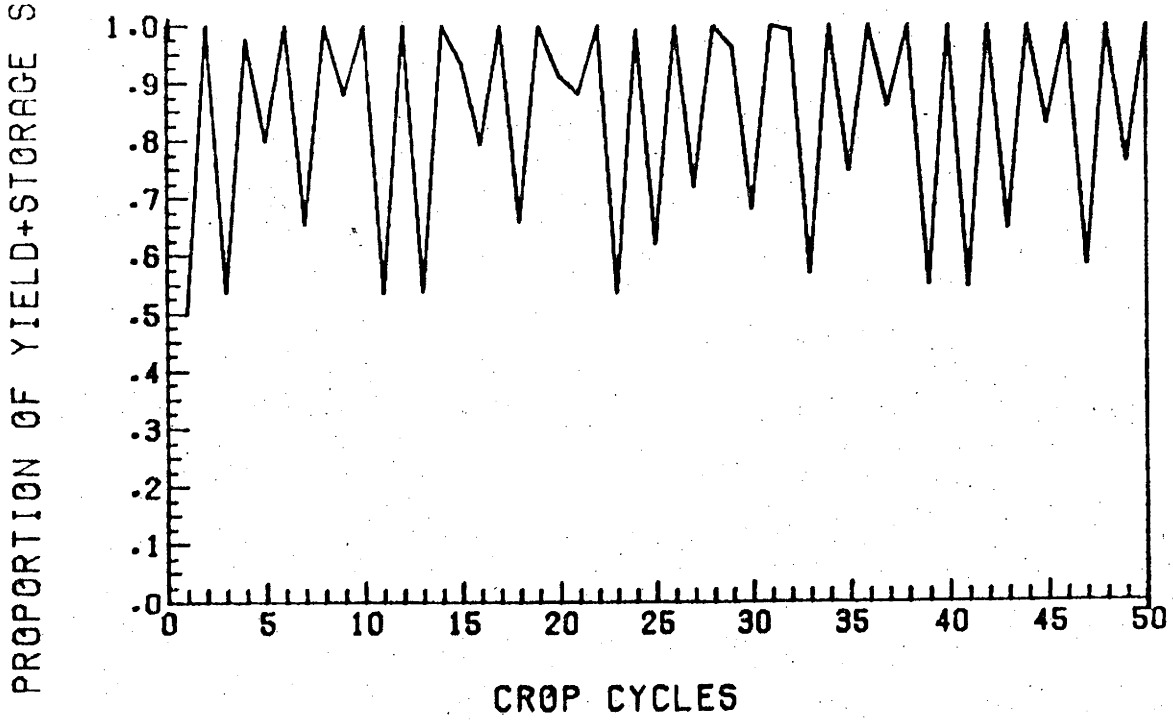

FIG 6.4.103: AGRIBUSINESS WITH GROWING PLANT AREA PLANTED THROUGH TIME

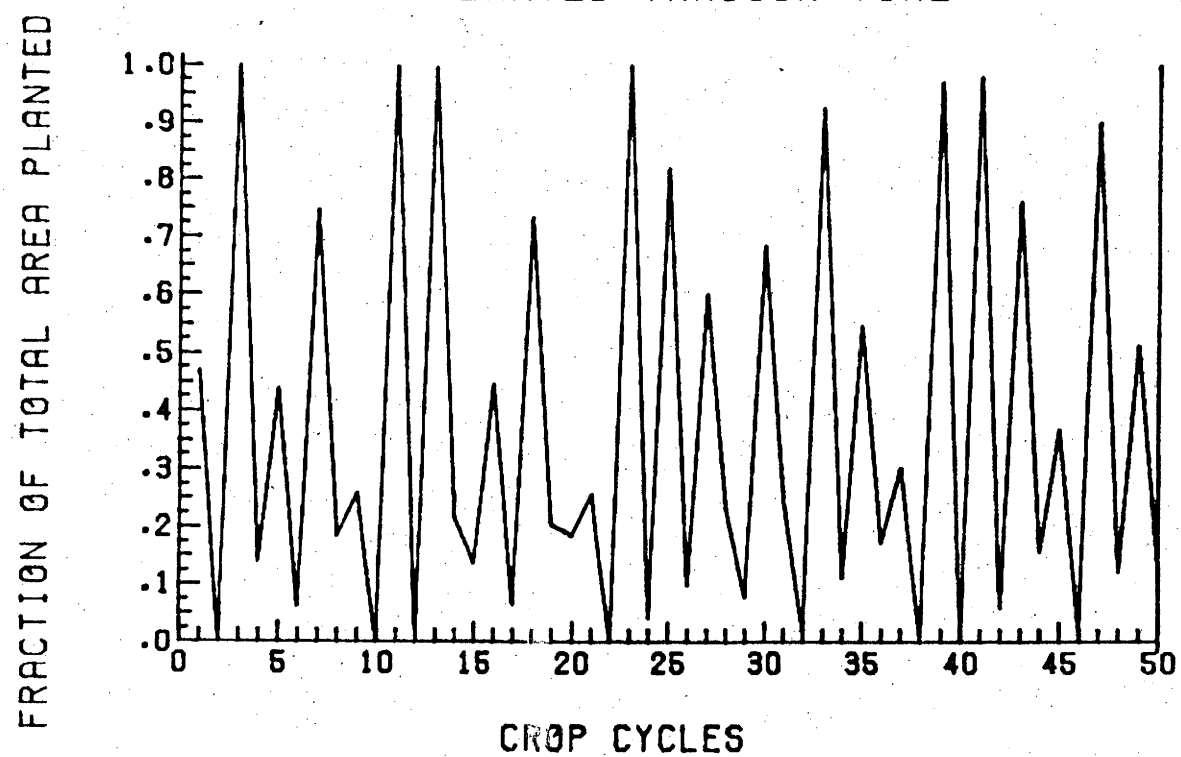

FIG 6.4.104: AGRIBUSINESS WITH GROWING PLANT RETURN THROUGH TIME

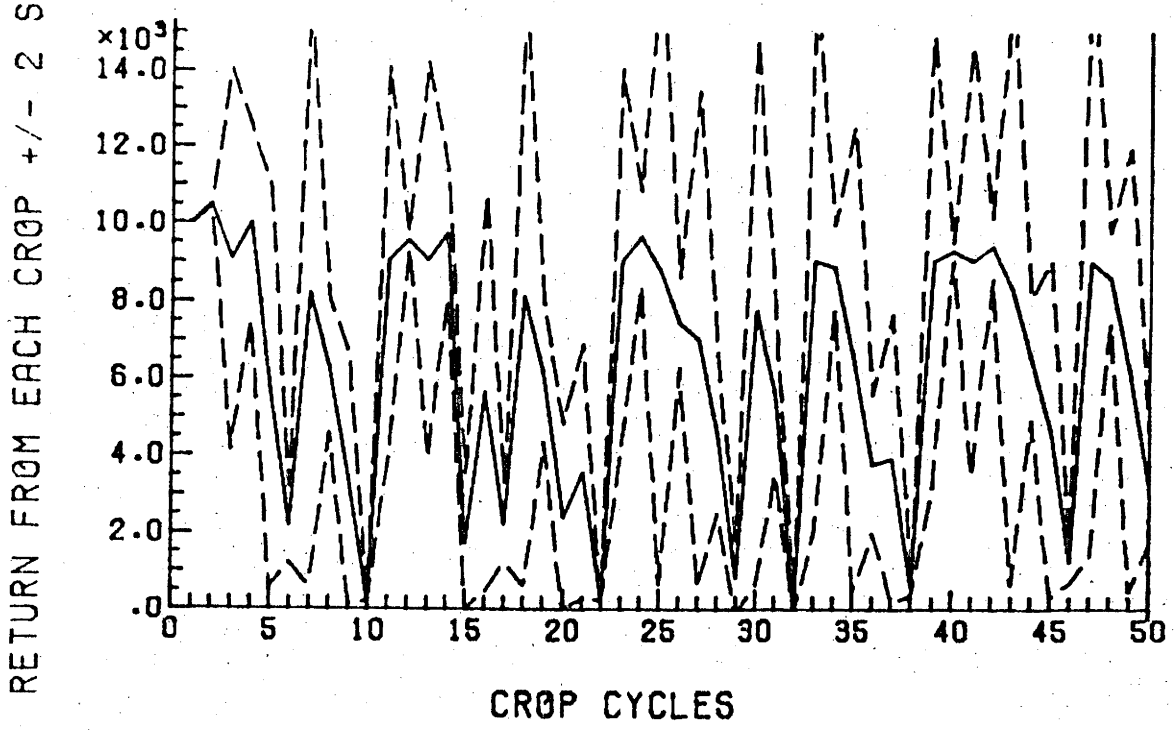

FIG 6.4.105: AGRIBUSINESS WITH GROWING PLANT YIELD THROUGH TIME

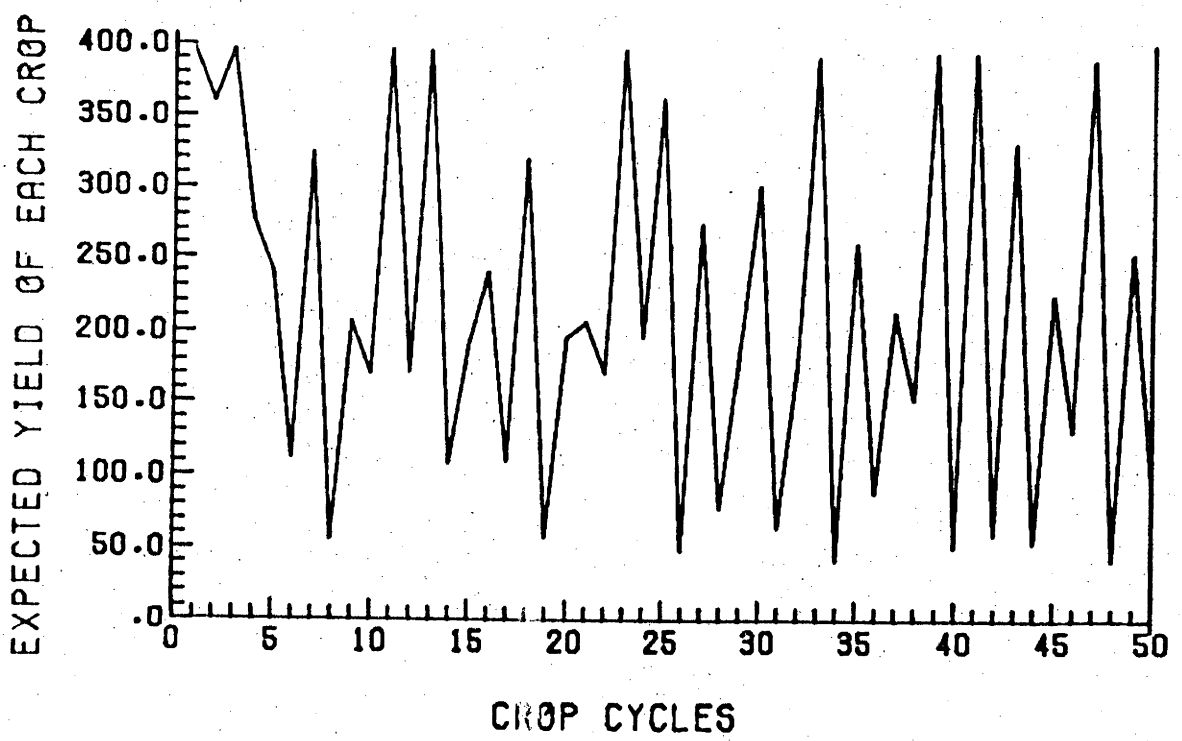

FIG 6.4.106: AGRIBUSINESS WITH GROWING PLANT STORAGE THROUGH TIME

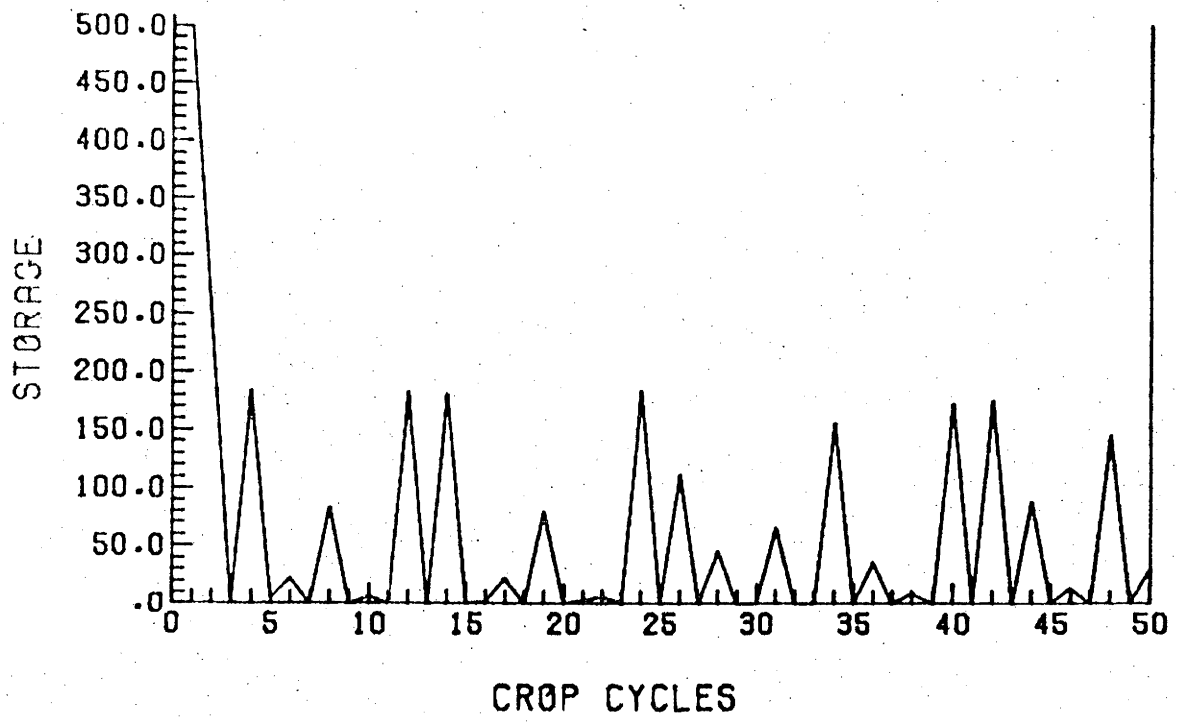

FIG 6.4.107: AGRIBUSINESS WITH GROWING PLANT SUPERRACE THROUGH TIME

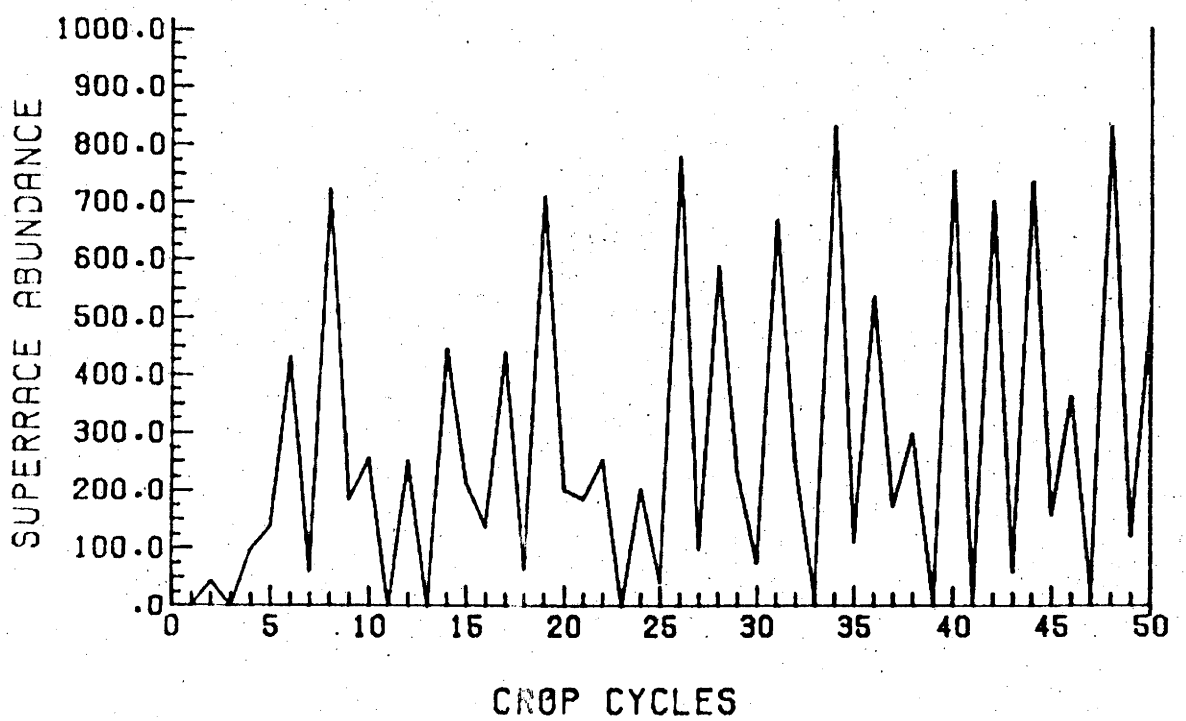

FIG 6.4.108A: AGRIBUSINESS GAME WITH GROWTH (1 - P(EXT)) VS STORAGE

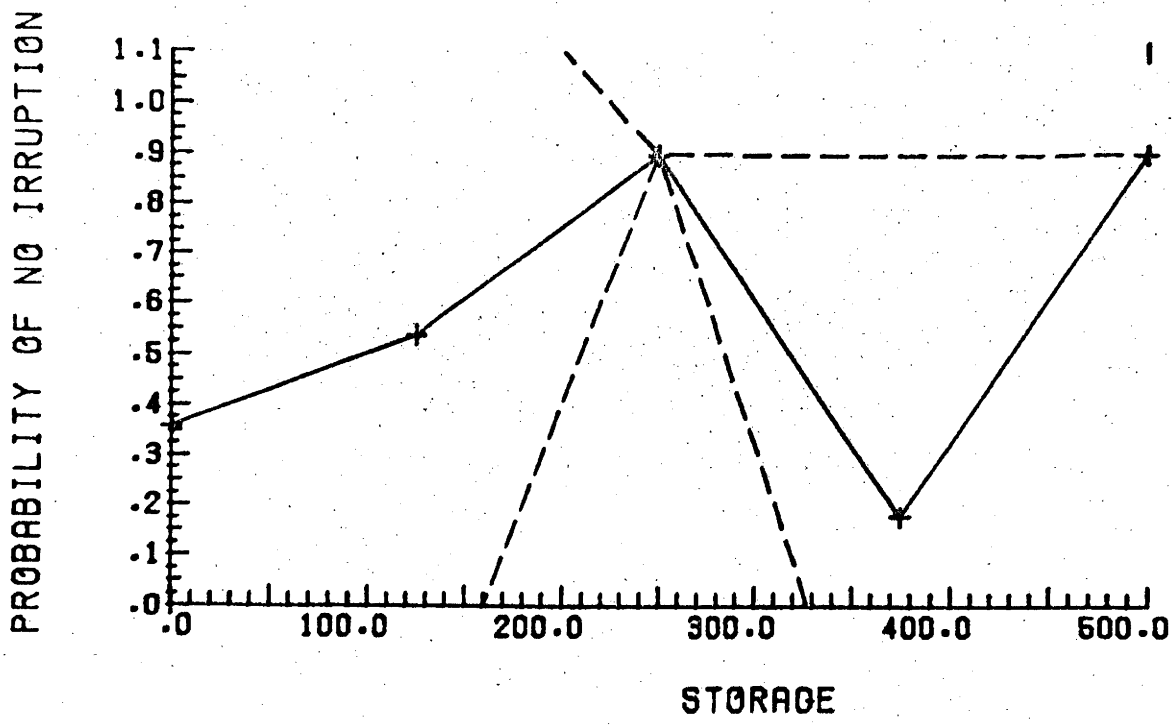

FIG 6.4.108B: AGRIBUSINESS GAME WITH GROWTH (1 - P(EXT)) VS AABB

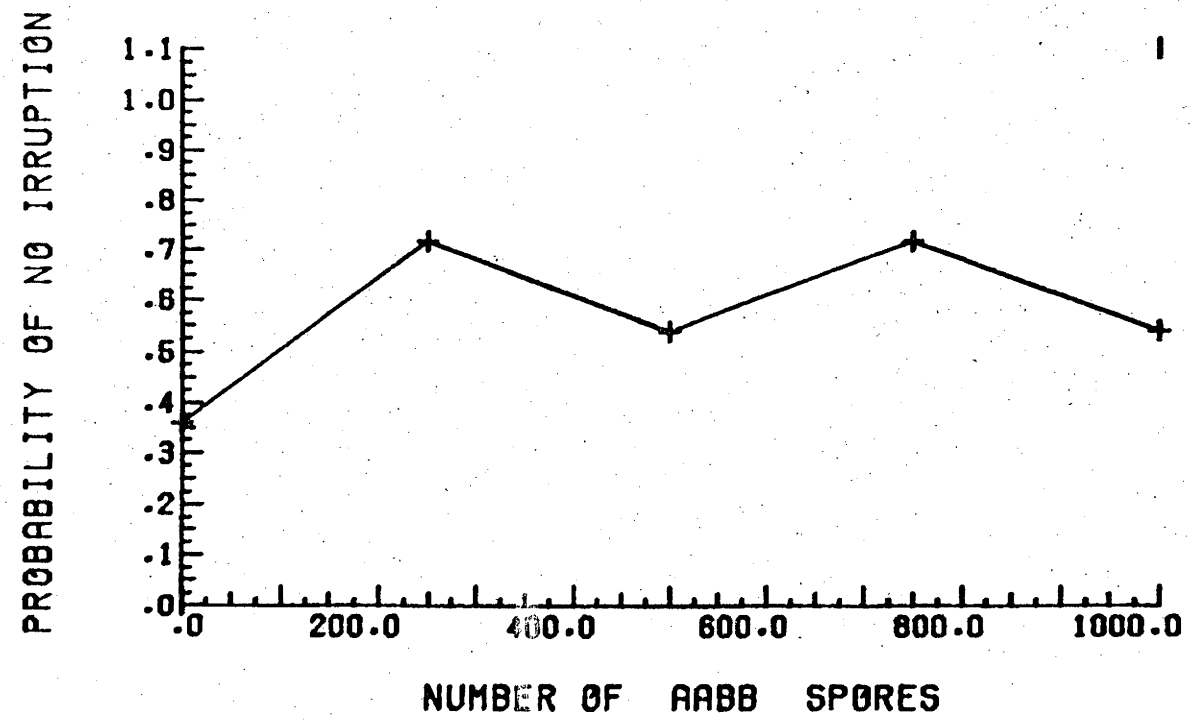

FIG 6.4.109A: AGRIBUSINESS GAME WITH GROWTH MUTATION RATE VS STORAGE

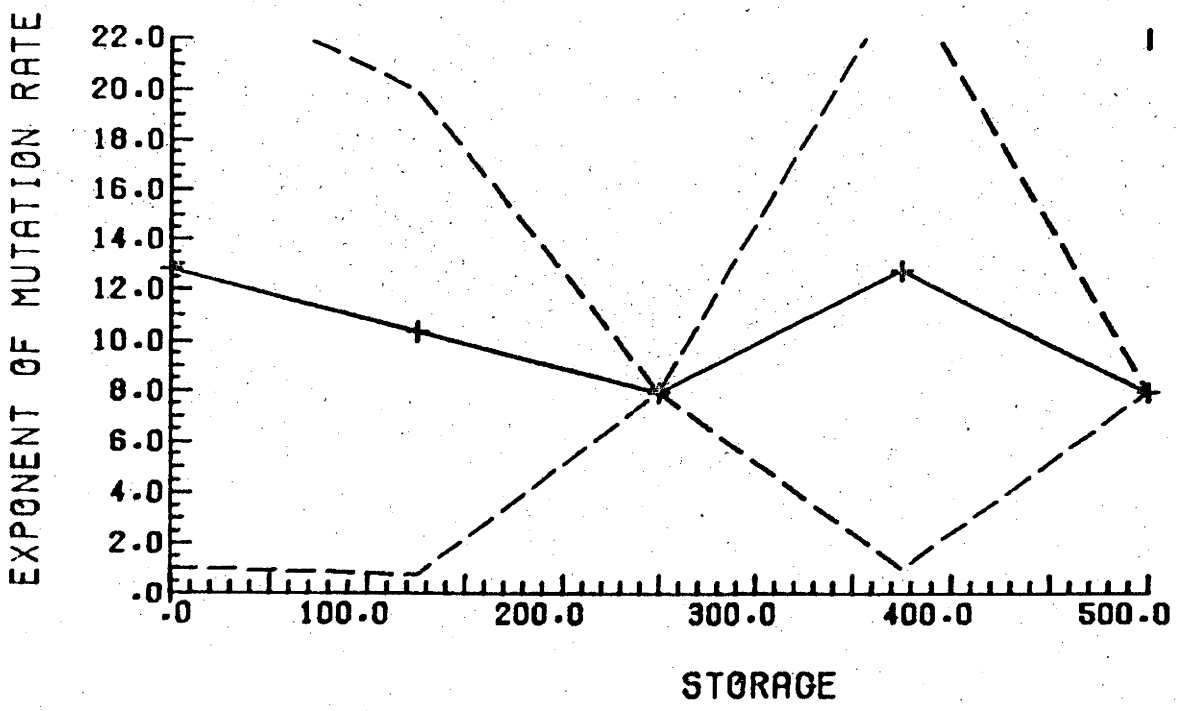

FIG 6.4.109B: AGRIBUSINESS GAME WITH GROWTH MUTATION RATE VS AABB

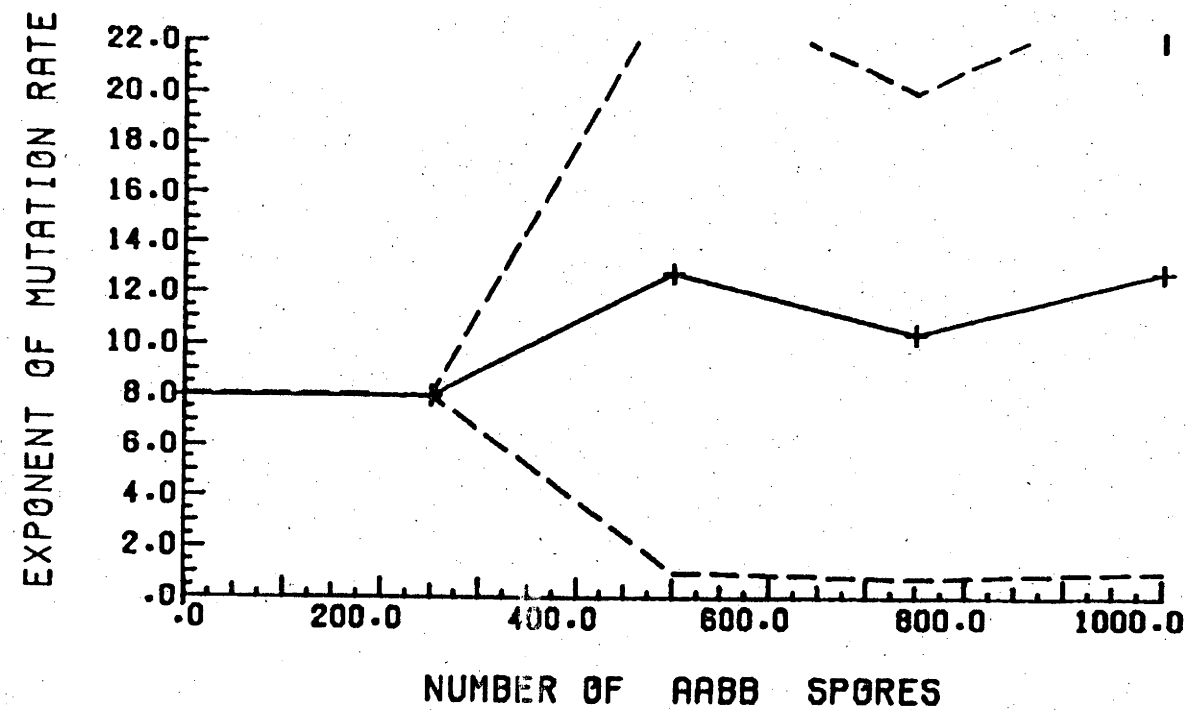

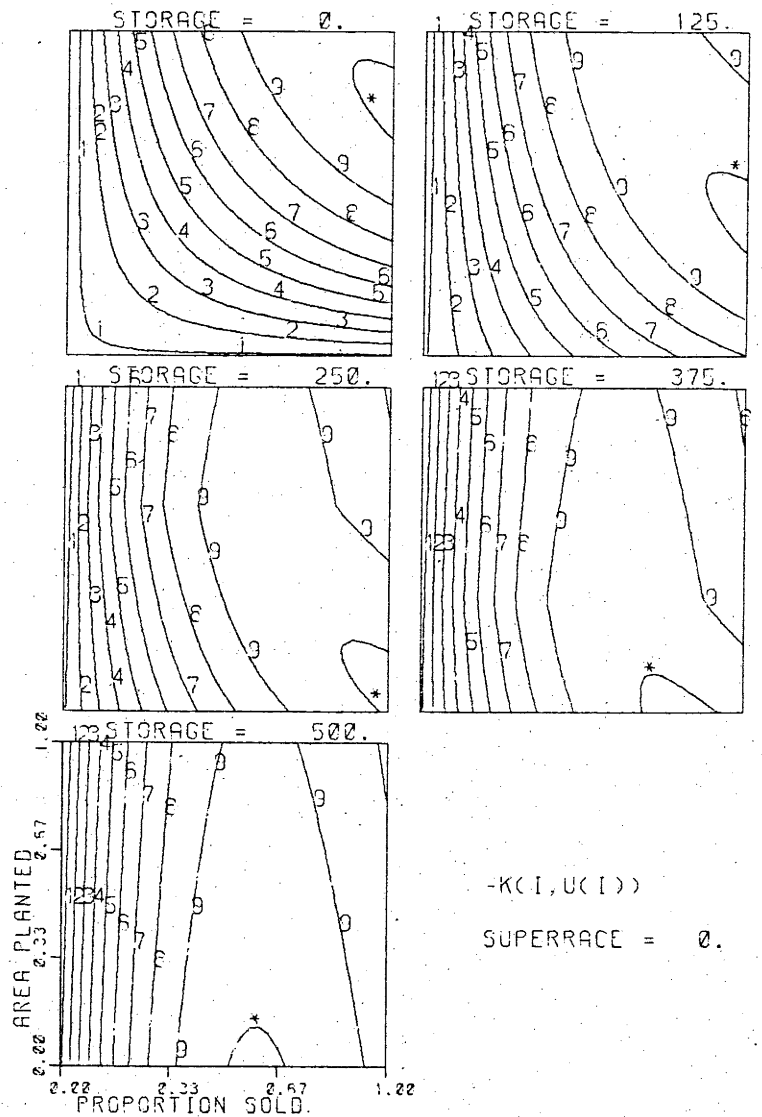

FIG. 6.4.110A: AGRIBUSINESS GAME WITH GROWTH

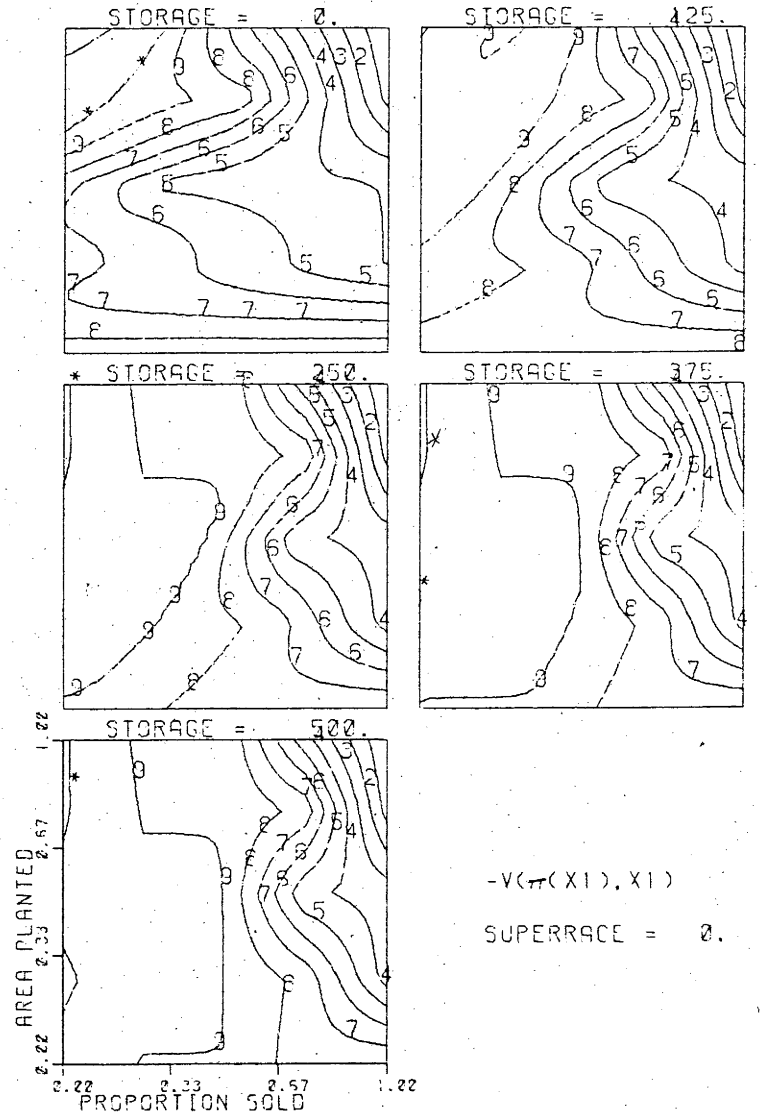

FIG. 6.4.110B: AGRIBUSINESS GAME WITH GROWTH

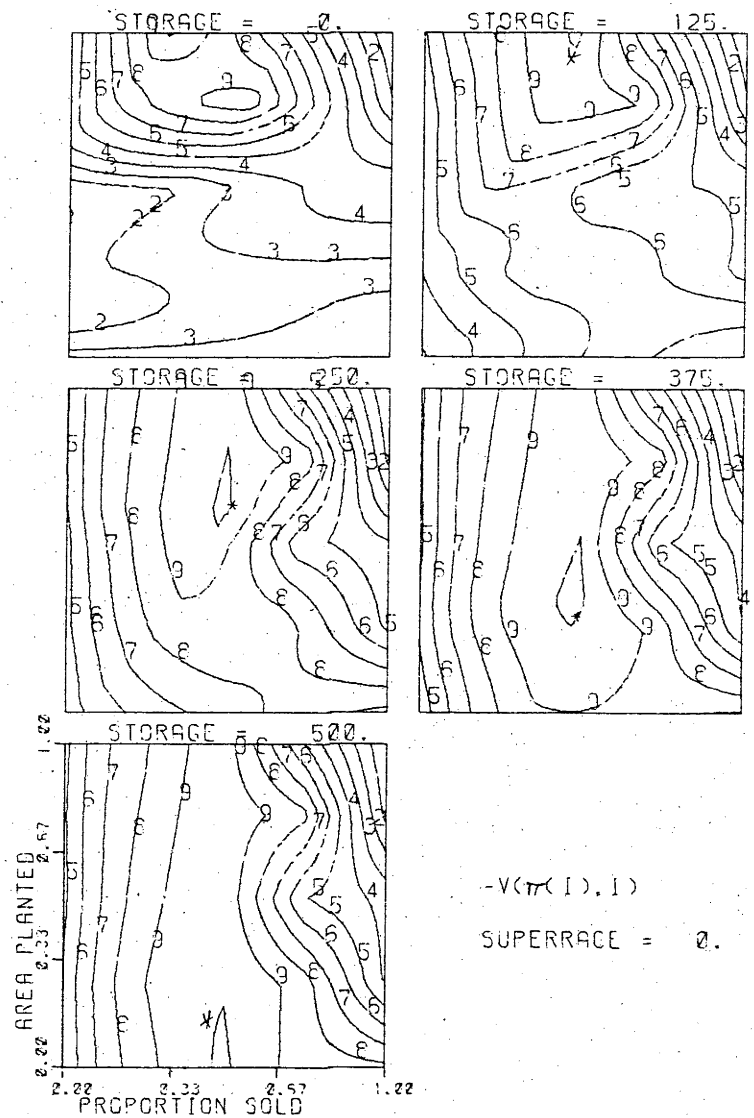

FIG. 6.4.110C: AGRIBUSINESS GAME WITH GROWTH

FIG 6.4.111: AGRIBUSINESS GAME WITH GROWTH FRACTION SOLD THROUGH TIME

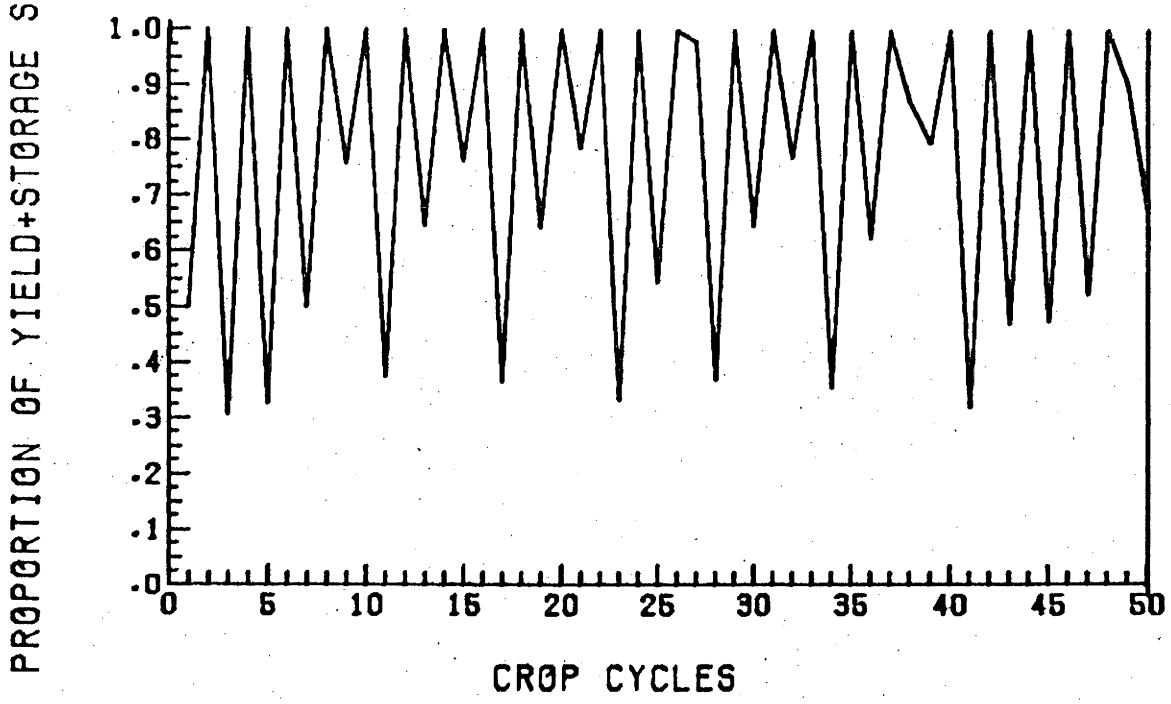

FIG 6.4.112: AGRIBUSINESS GAME WITH GROWTH AREA PLANTED THROUGH TIME

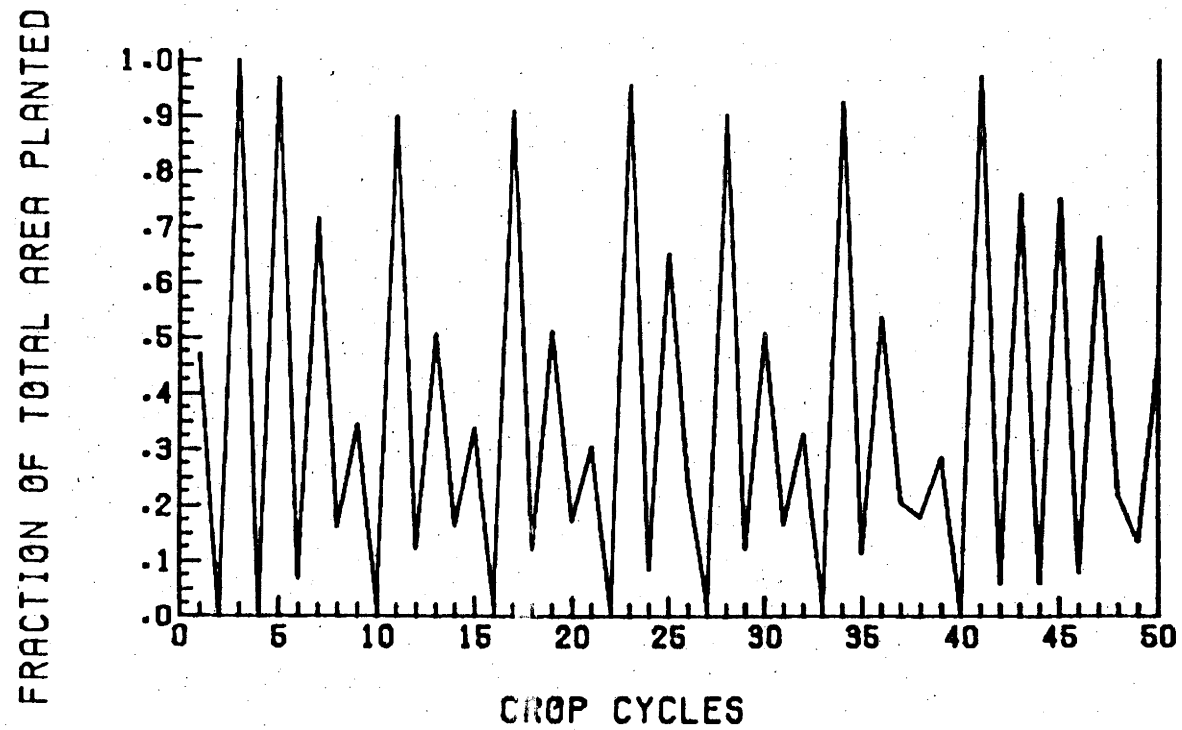

FIG 6.4.115: AGRIBUSINESS GAME WITH GROWTH (1-P(EXT)) THROUGH TIME

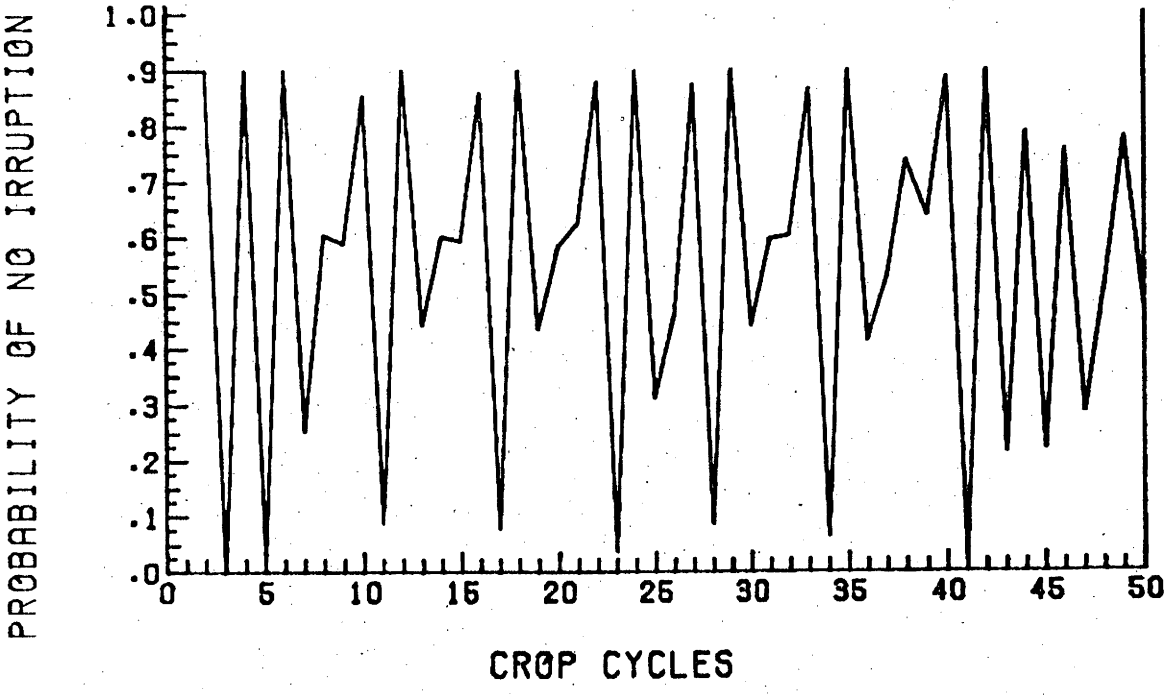

FIG 6.4.116: AGRIBUSINESS GAME WITH GROWTH MUTATION RATE THROUGH TIME

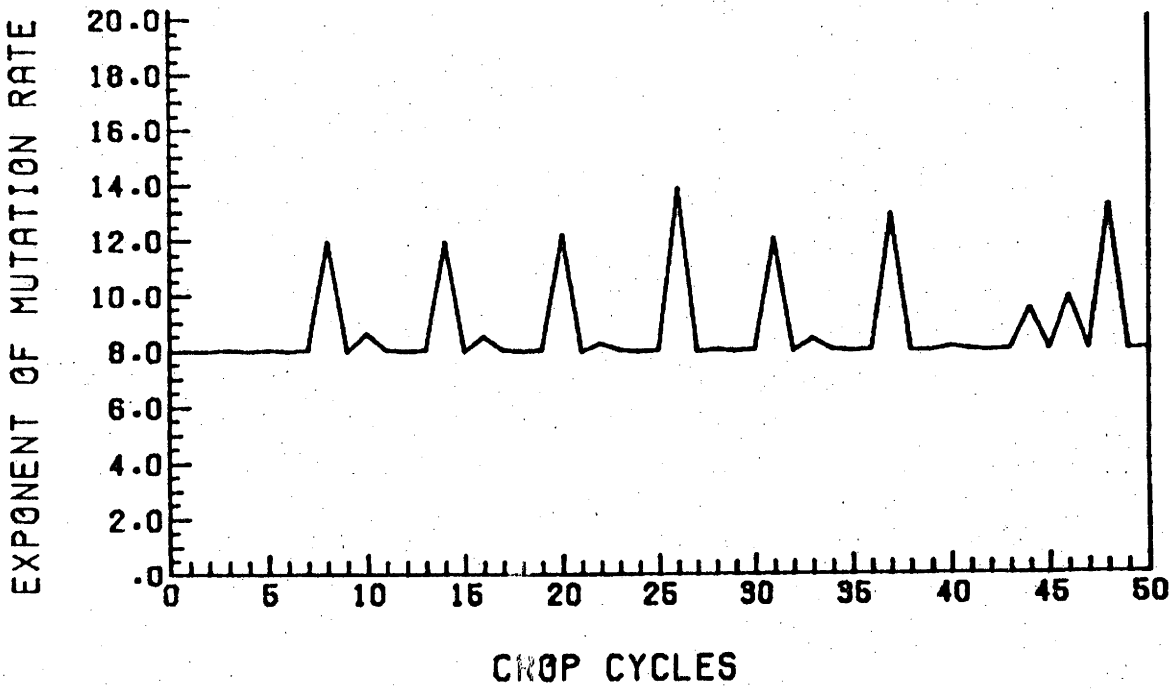

FIG. 6.5.1: DISEASE FREE SUBSISTENCE
DAYS WORKED VS STORAGE

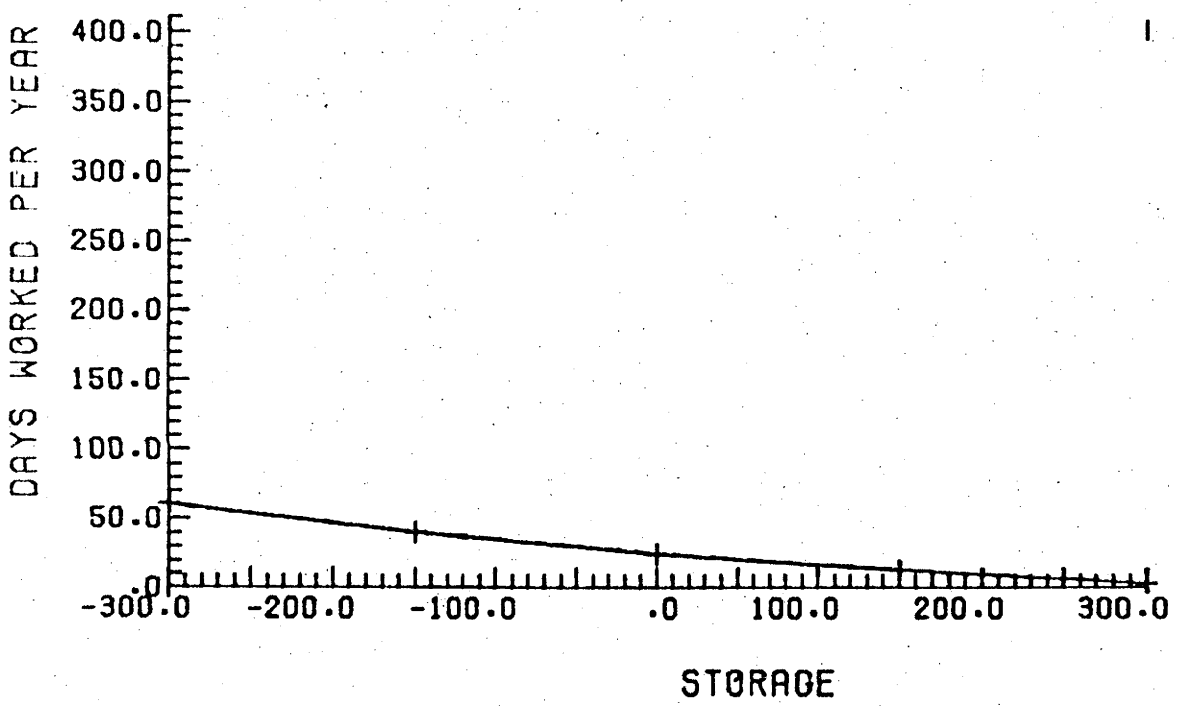

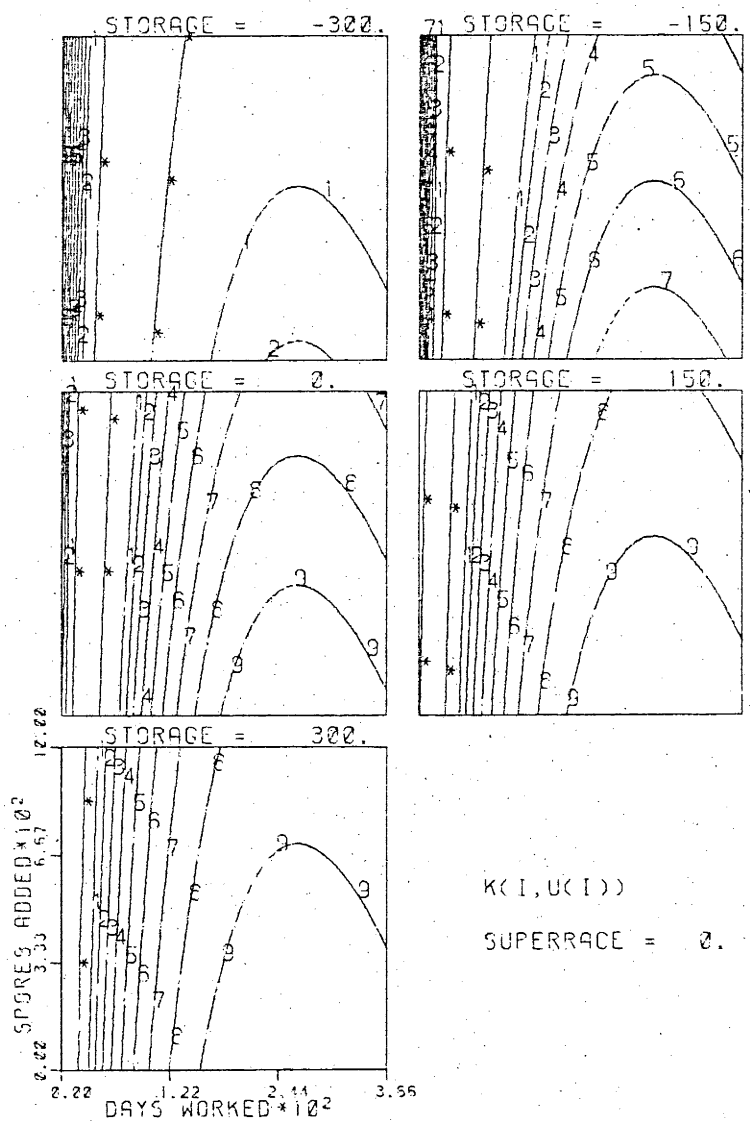

FIG. 6 5.2A: DISEASE-FREE SUBSISTENCE

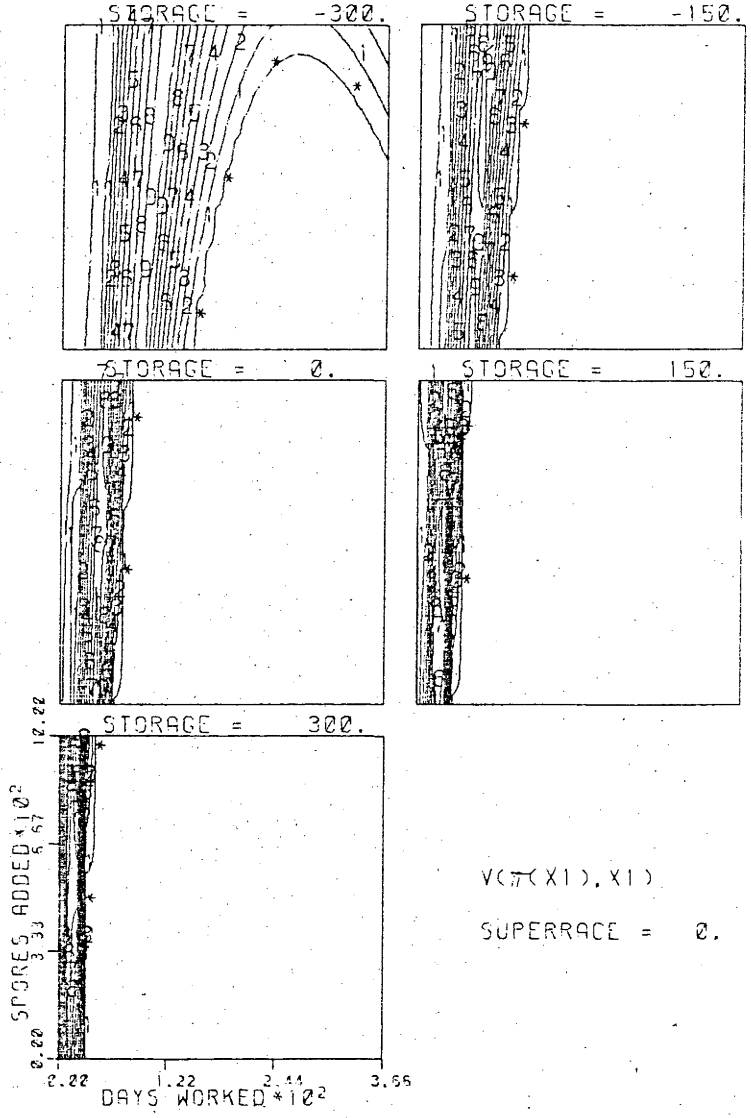

FIG. 6 5.2B: DISEASE-FREE SUBSISTENCE

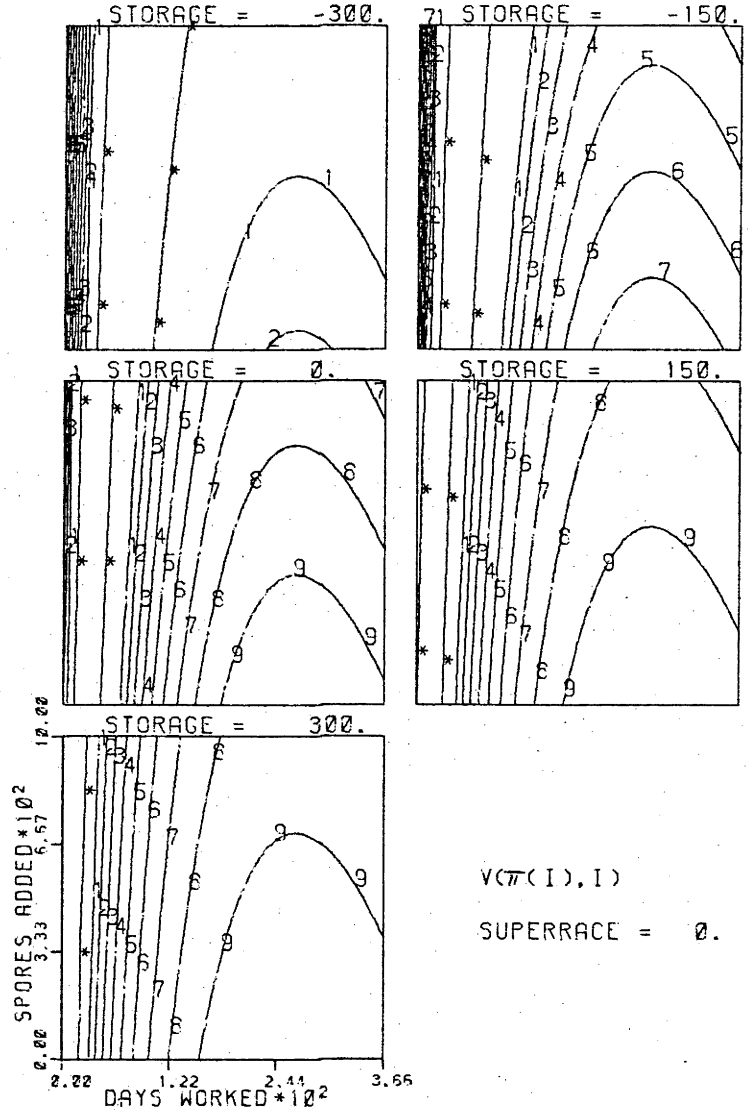

FIG. 6.5.2C: DISEASE-FREE SUBSISTENCE

FIG. 6.5.3: DISEASE-FREE WORK DONE THROUGH TIME

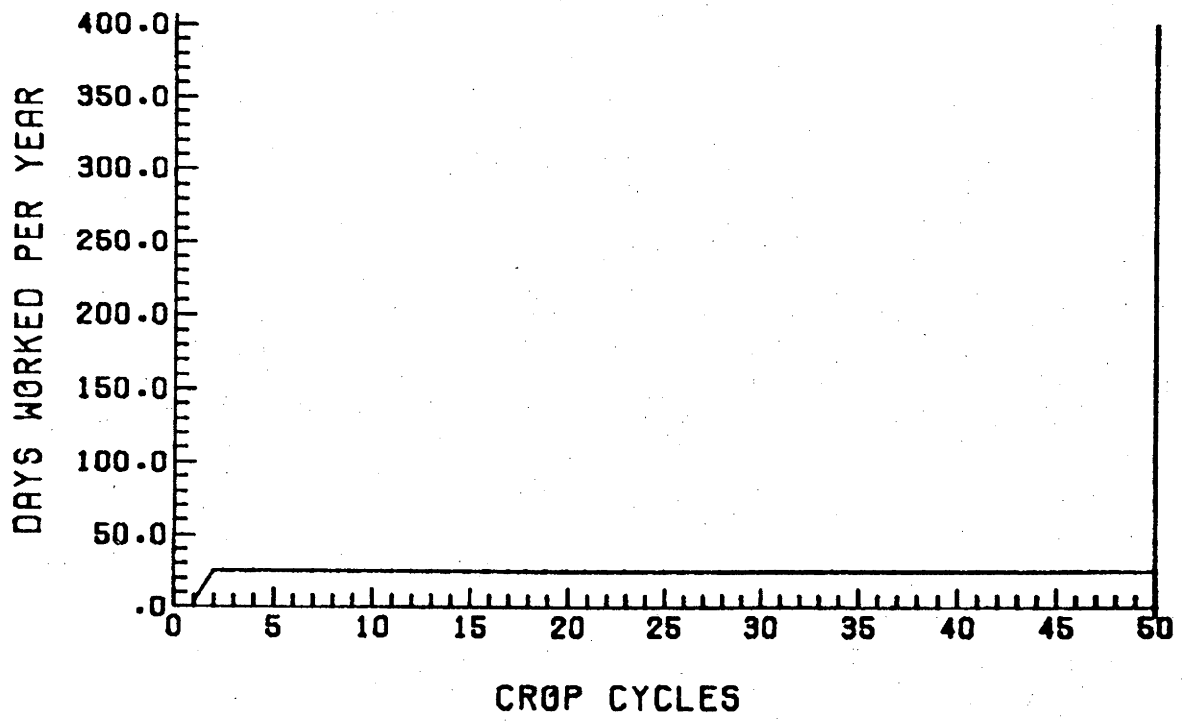

FIG. 6.5.4: DISEASE-FREE STORAGE THROUGH TIME

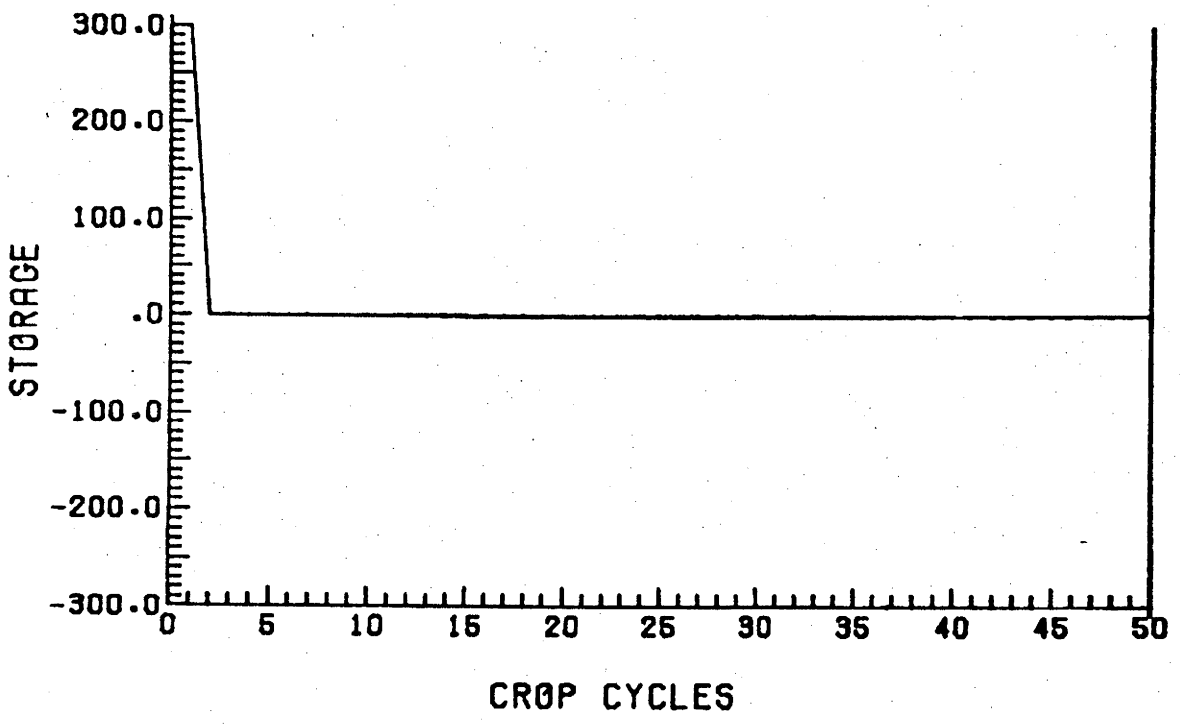

FIG. 6.5.5A: SUBSISTENCE BASELINE
DAYS WORKED VS STORAGE

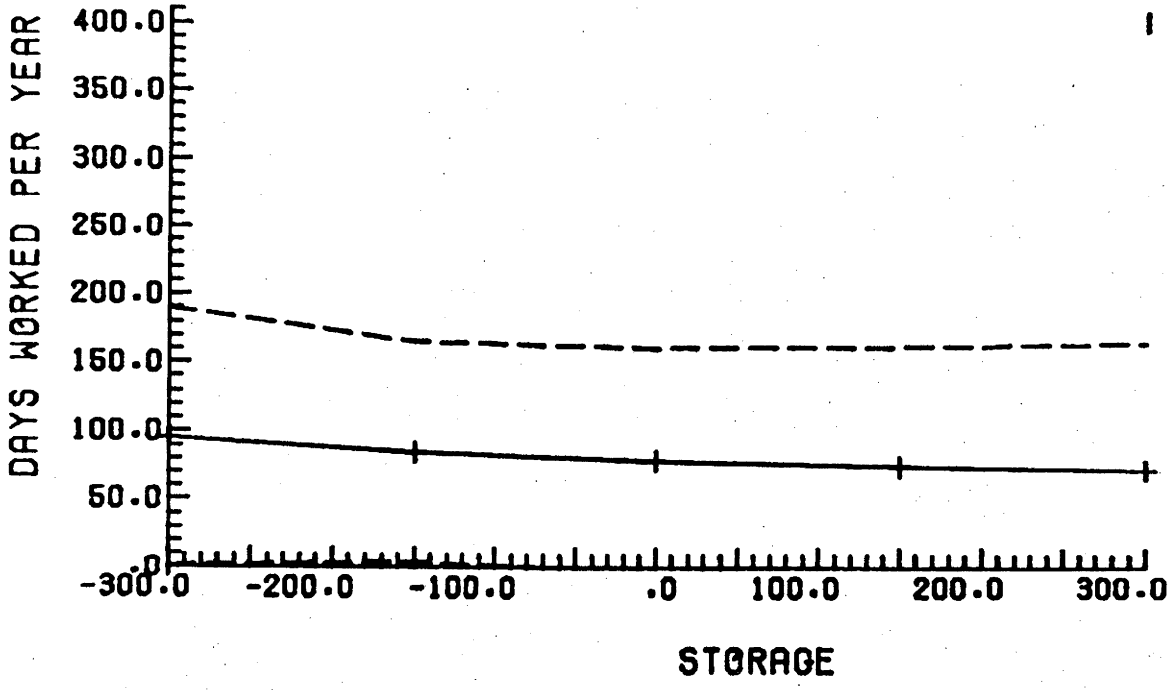

FIG. 6.5.5B: SUBSISTENCE BASELINE
DAYS WORKED VS AABB

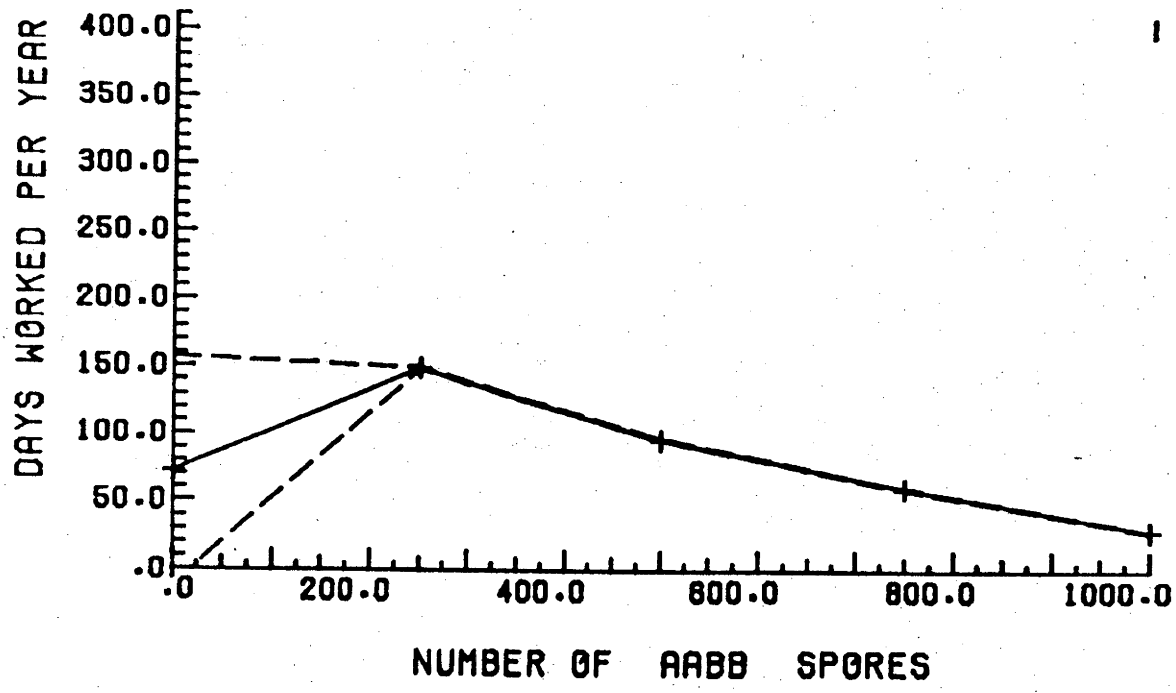

FIG. 6.5.6A: SUBSISTENCE BASELINE
E(COST) VS STORAGE

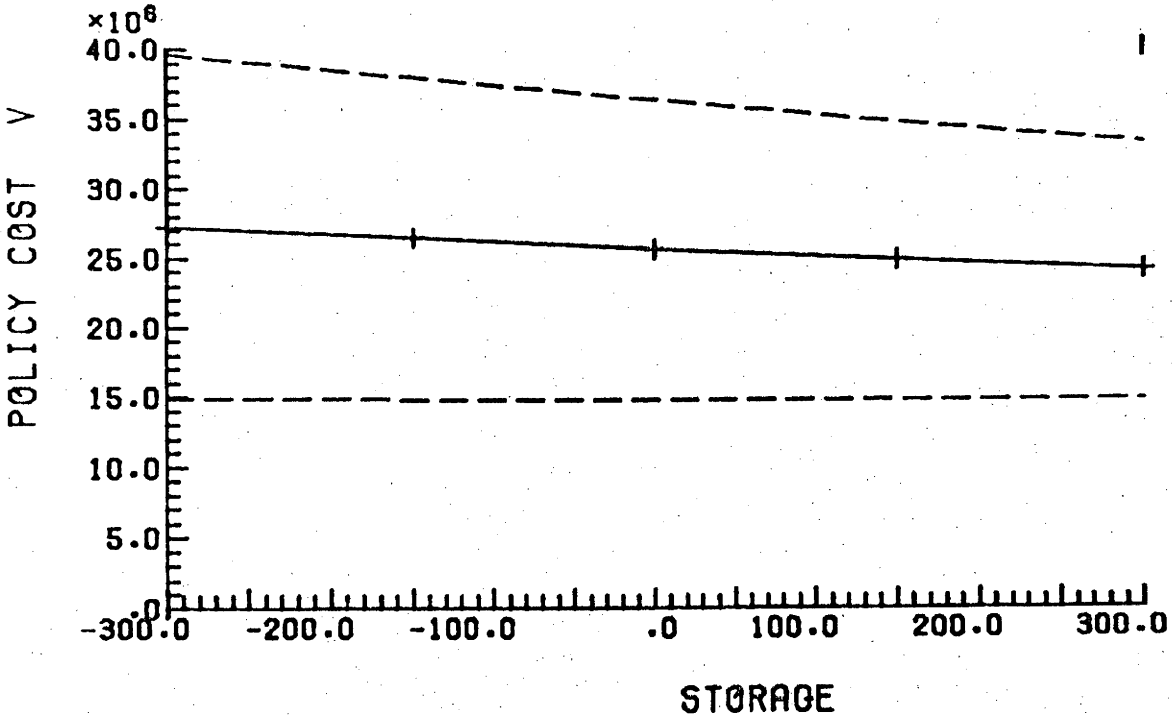

FIG. 6.5.6B: SUBSISTENCE BASELINE
E(COST) VS AABB

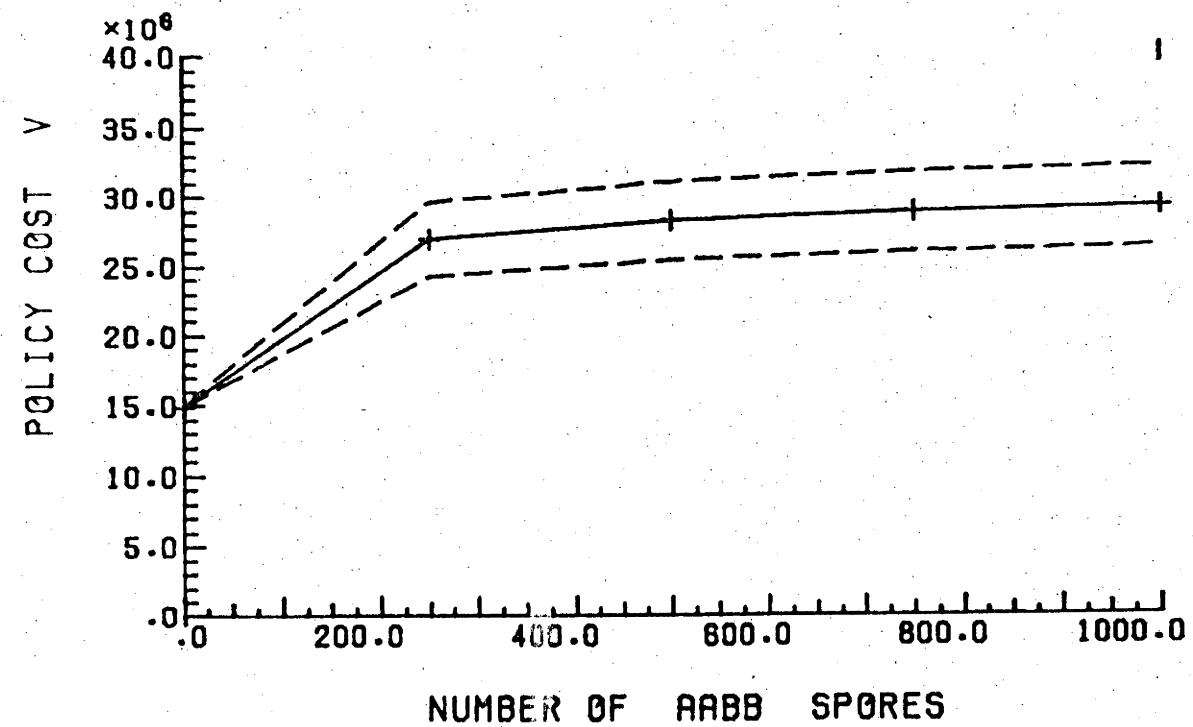

FIG. 6.5.7A: SUBSISTENCE BASELINE
STORAGE INPUT/OUTPUT

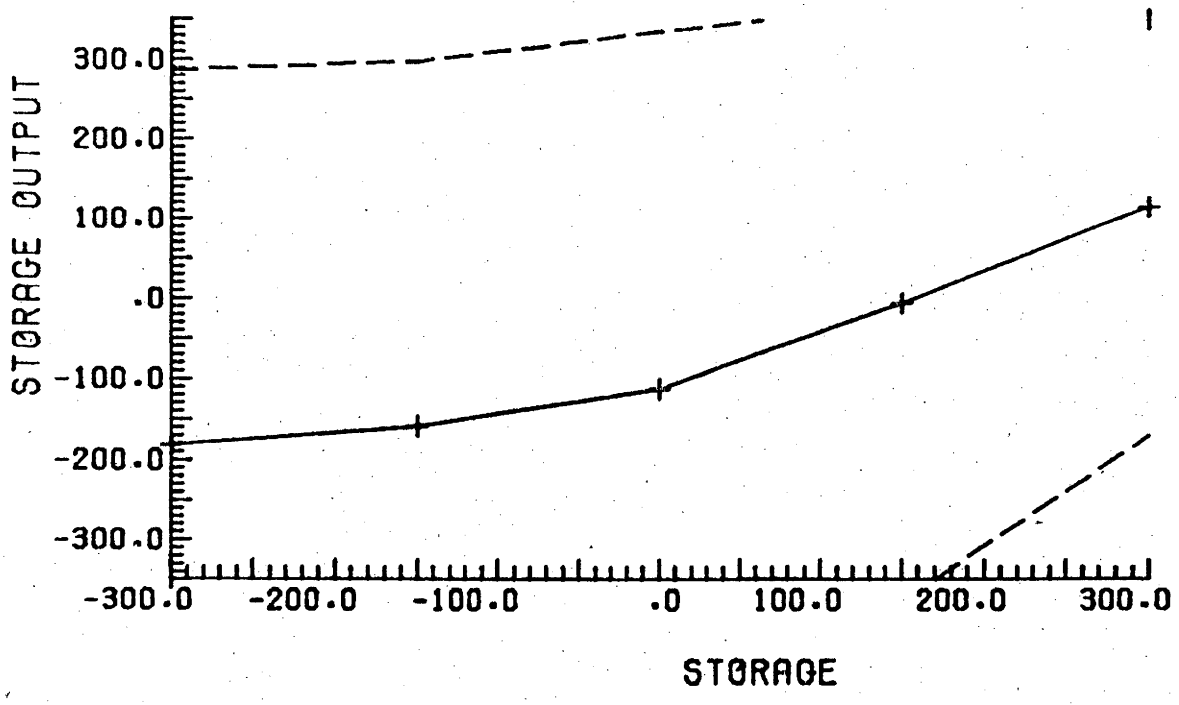

FIG. 6.5.7B: SUBSISTENCE BASELINE
STORAGE OUTPUT VS AABB

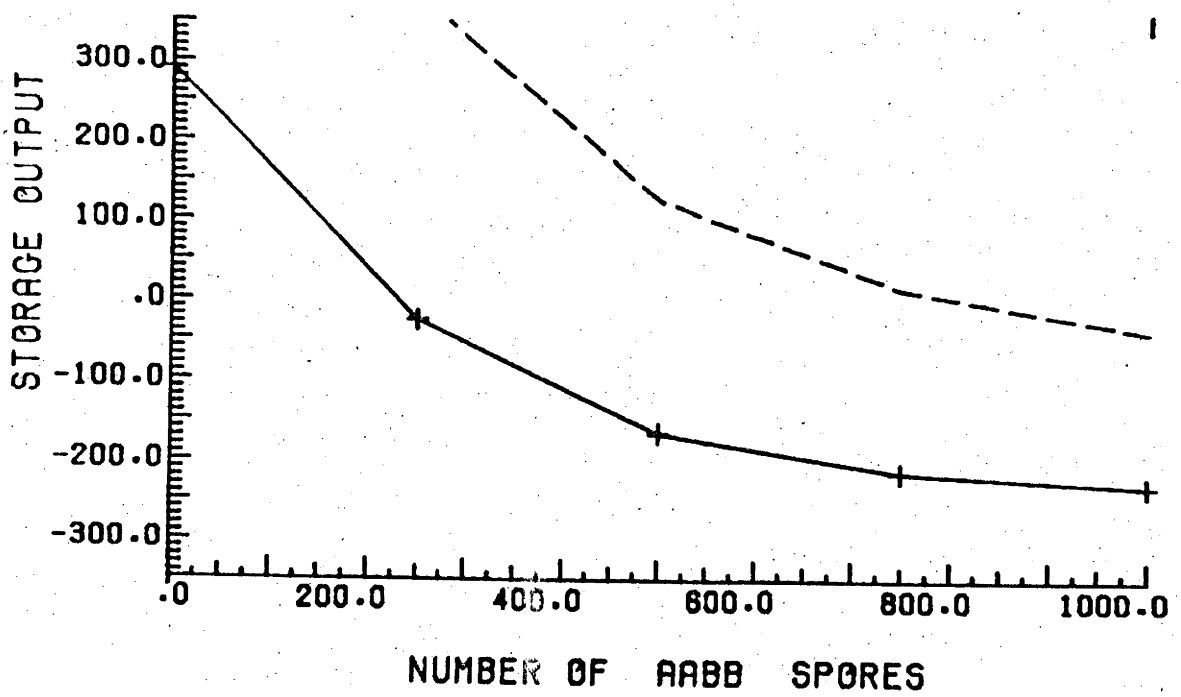

FIG. 6.5.8: SUBSISTENCE BASELINE
SUPERRACE INPUT/0UTPUT

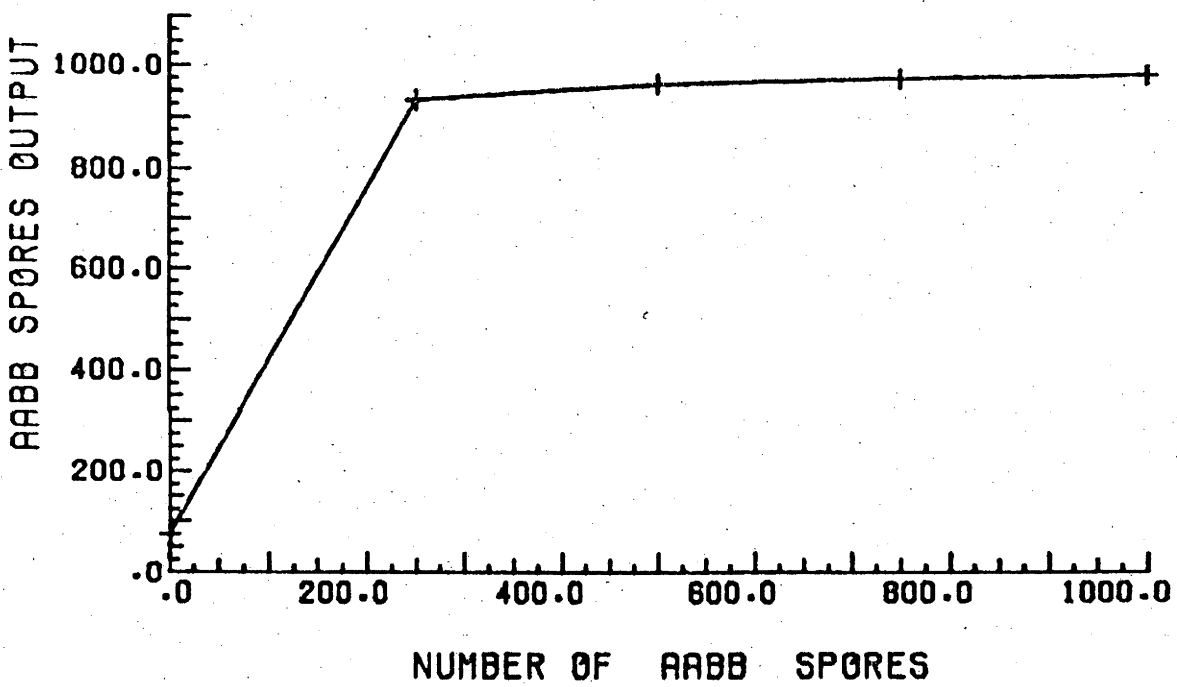

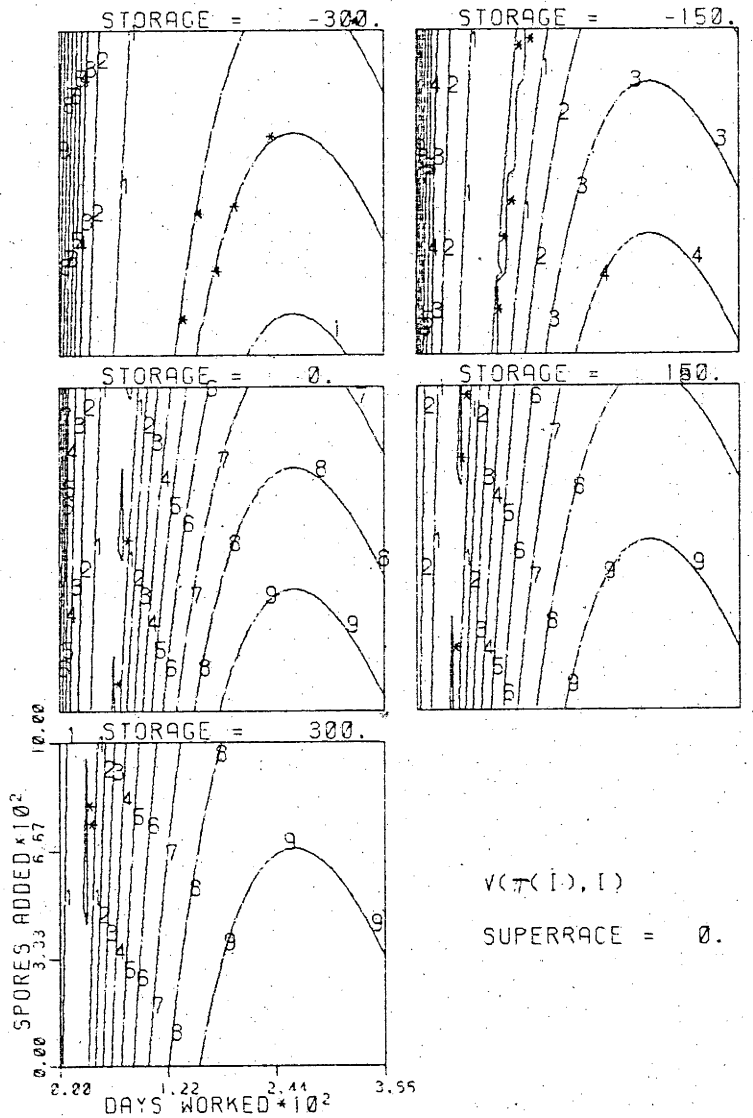

FIG. 6.5.9C· SUBSISTENCE BASELINE COST

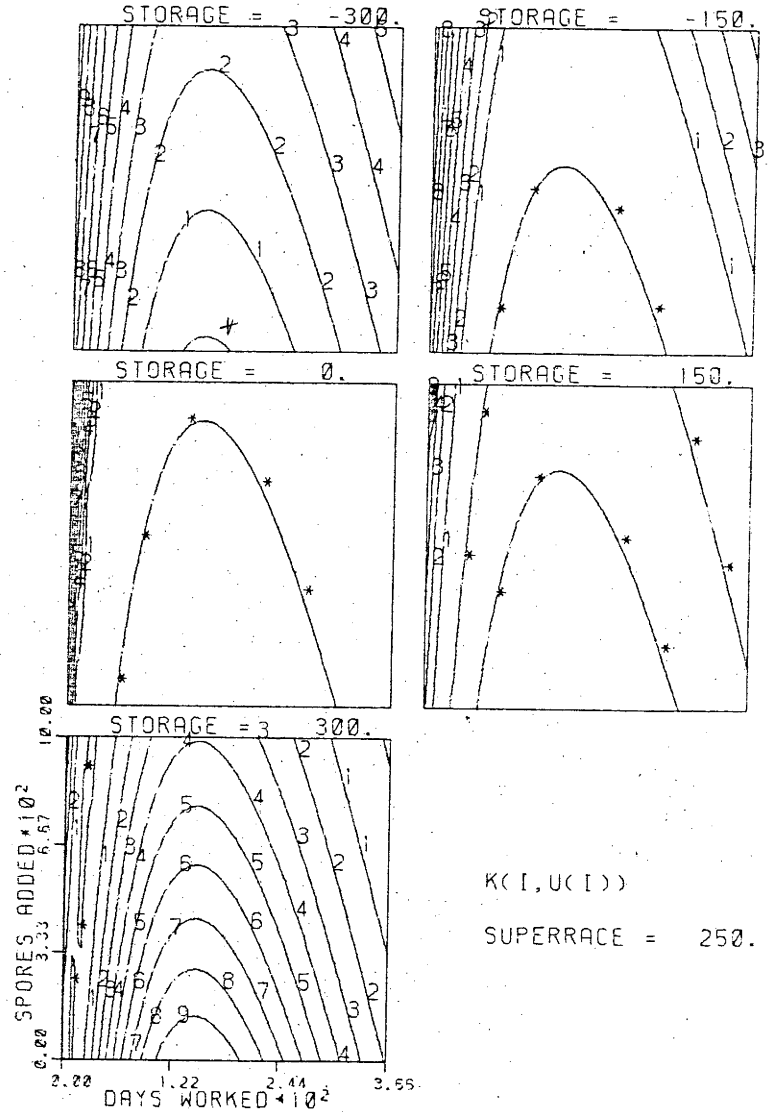

FIG. 6.5.10A· SUBSISTENCE BASELINE COST

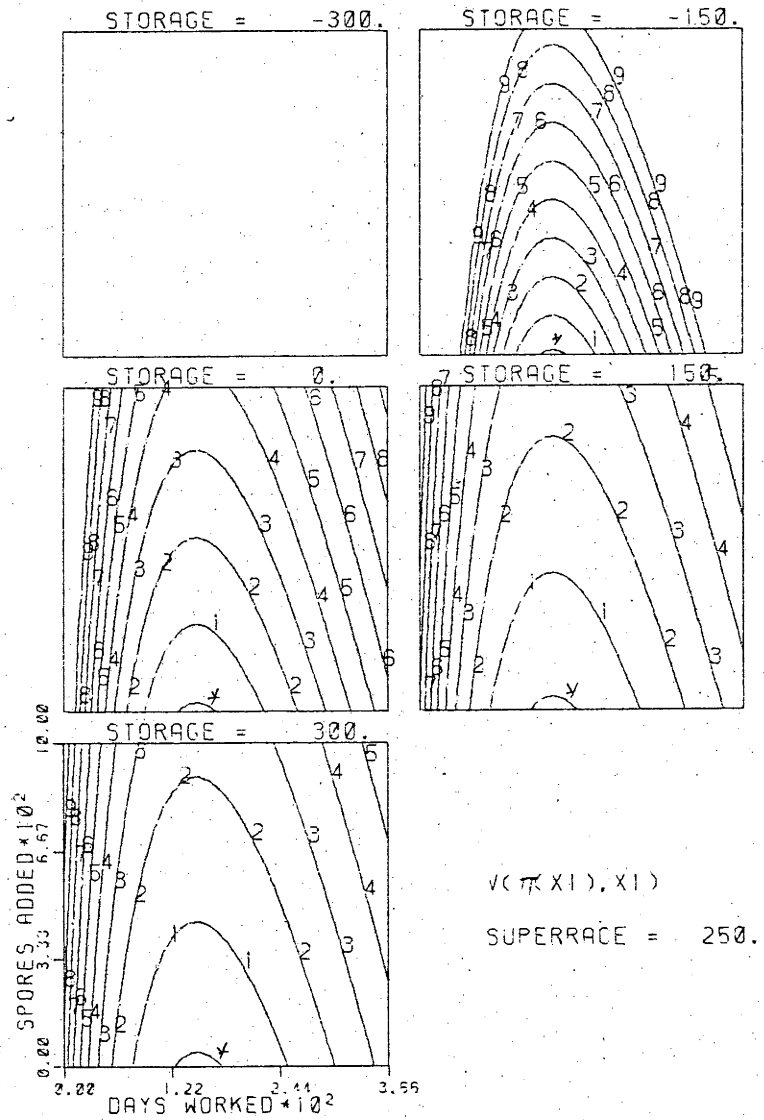

FIG. 6.5.10B· SUBSISTENCE BASELINE COST

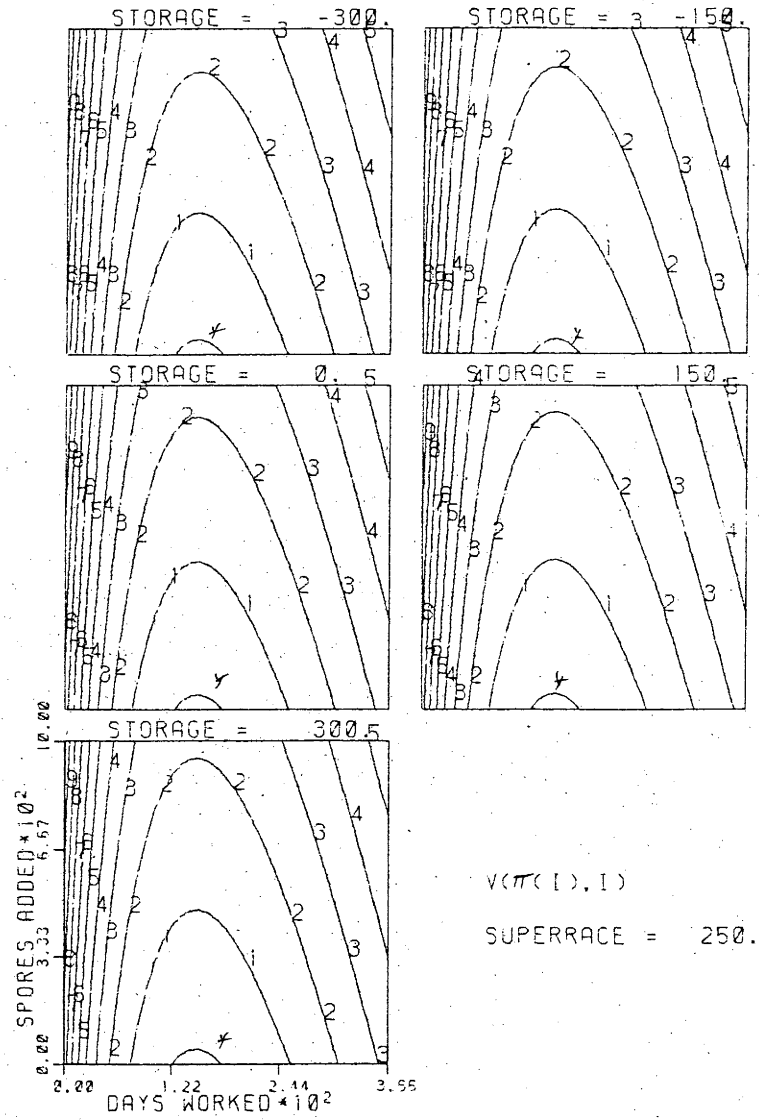

FIG. 6.5.10C· SUBSISTENCE BASELINE COST

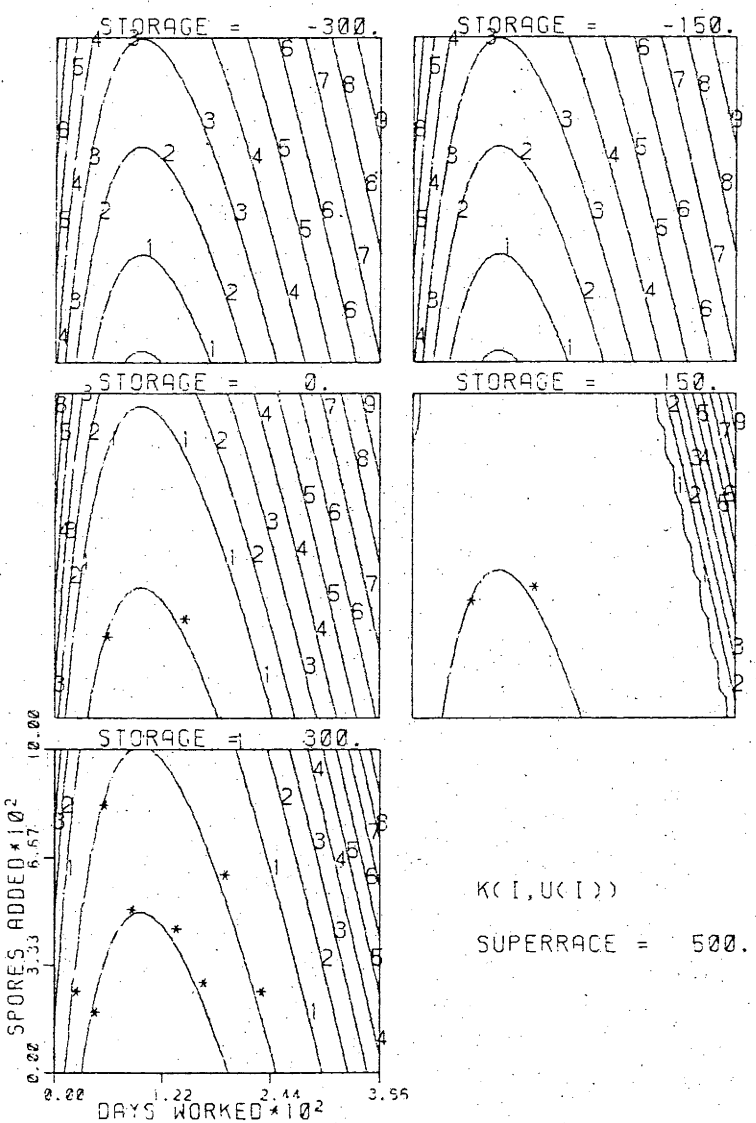

FIG. 6.5.11A: SUBSISTENCE BASELINE COST

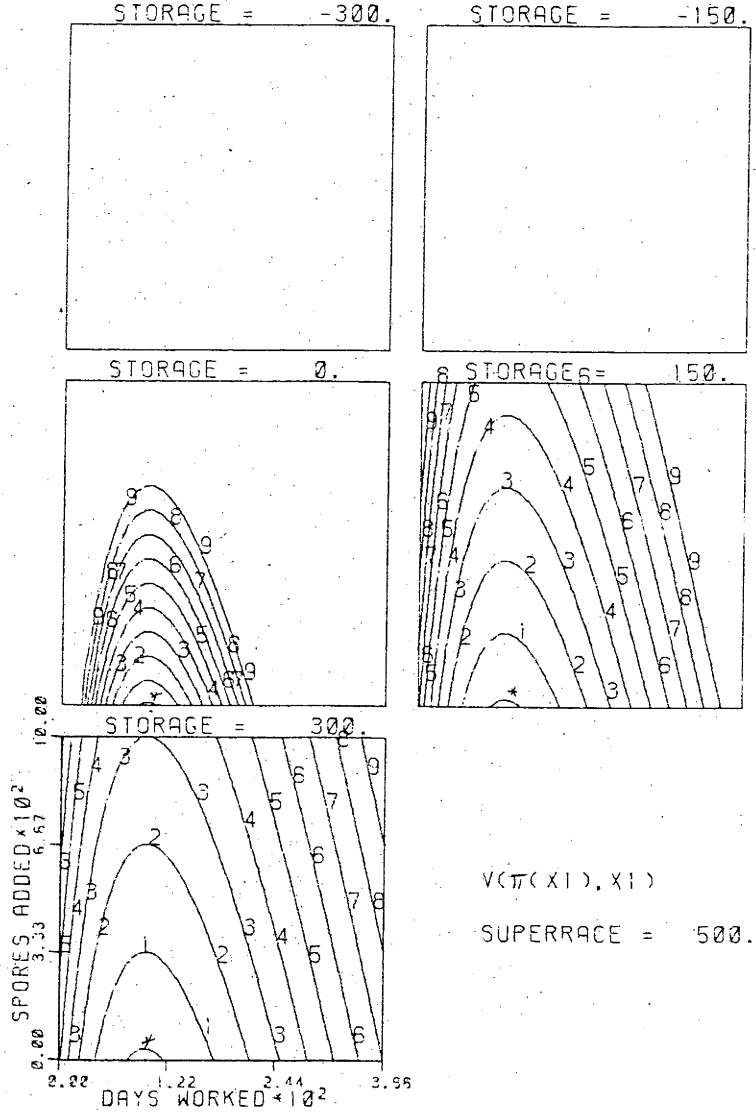

FIG. 6.5.11B: SUBSISTENCE BASELINE COST

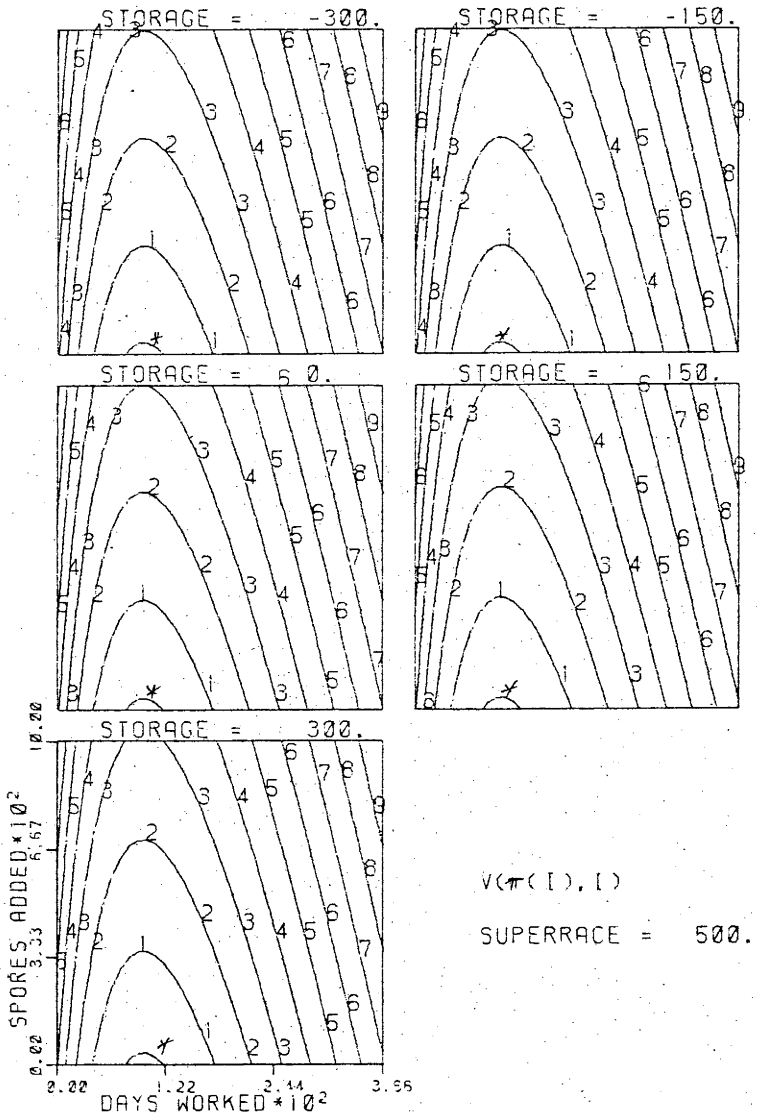

FIG. 6.5.11C: SUBSISTENCE BASELINE COST

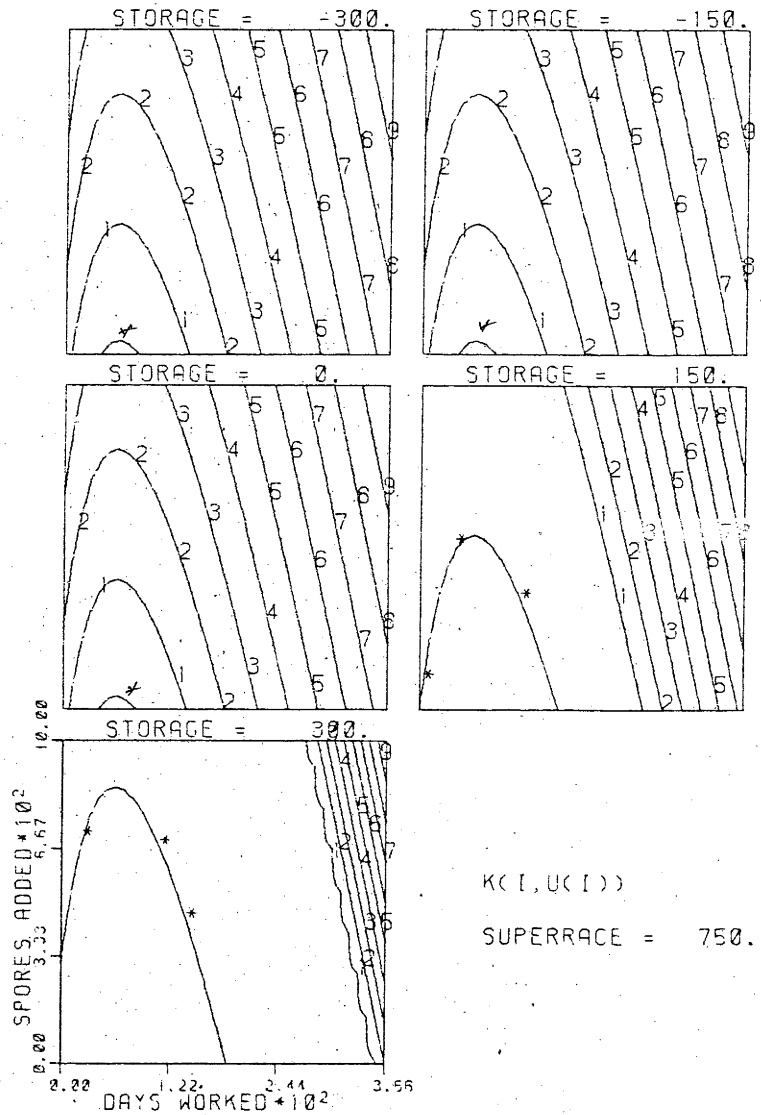

FIG. 6.5.12A: SUBSISTENCE BASELINE COST

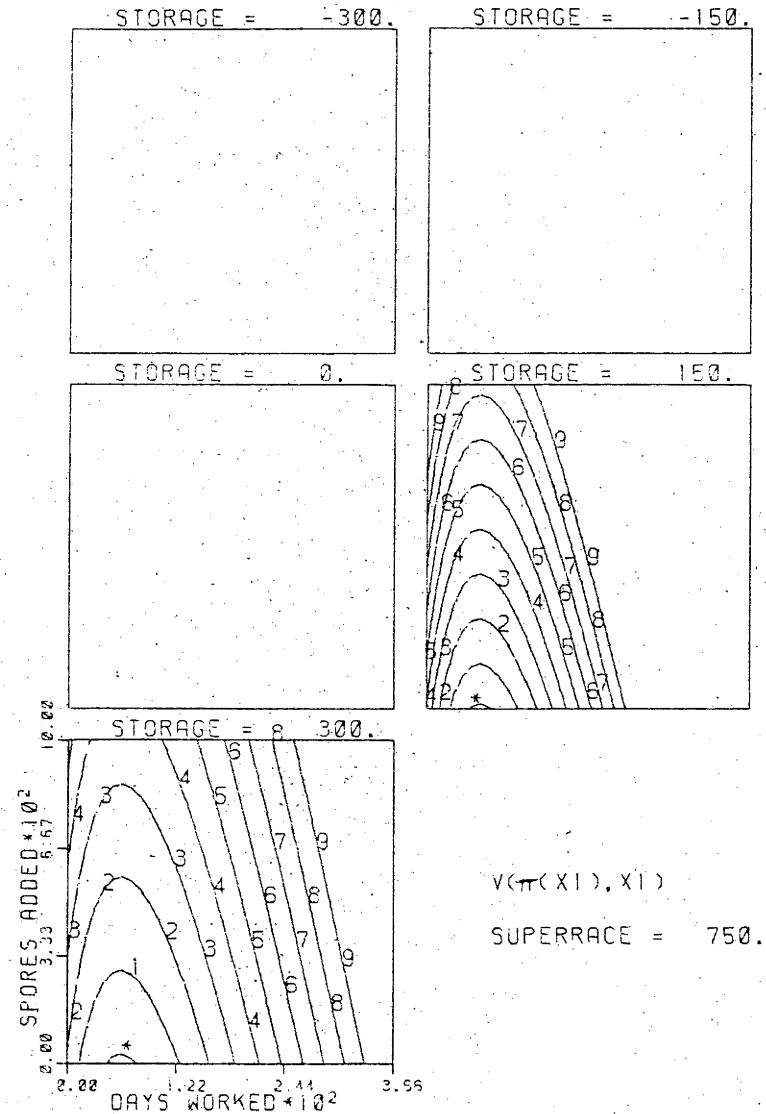

FIG. 6.5.12B: SUBSISTENCE BASELINE COST

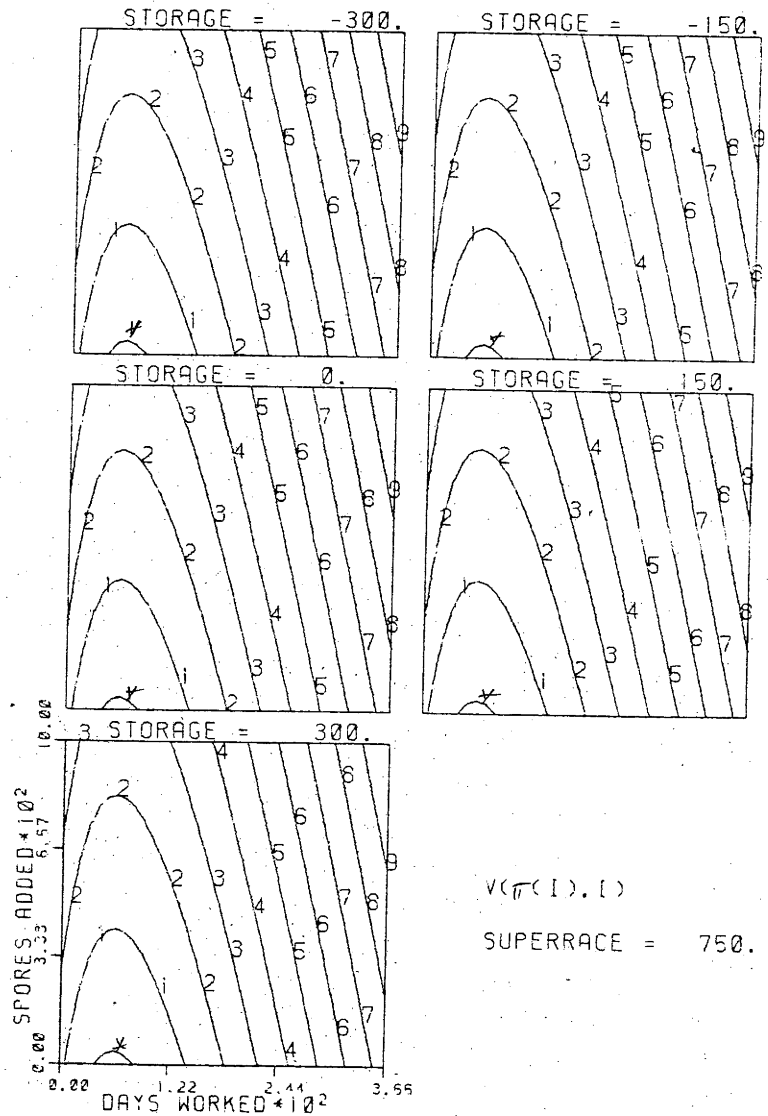

FIG. 6.5.12C: SUBSISTENCE BASELINE COST

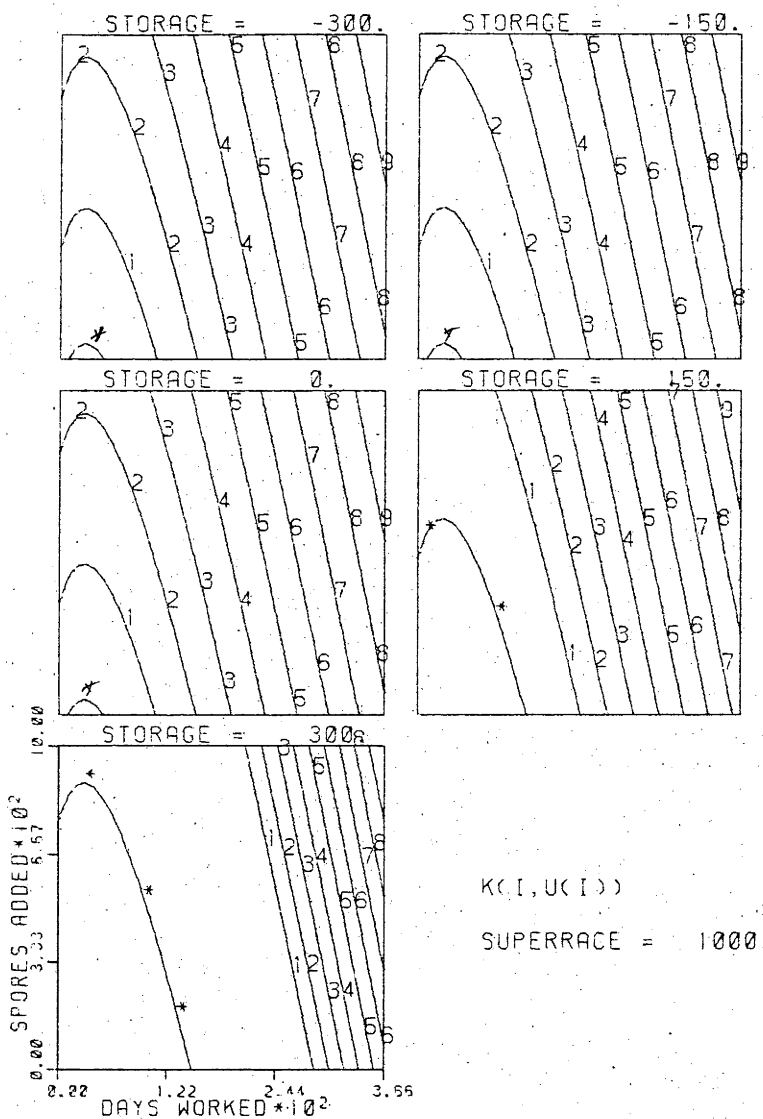

FIG. 6.5.13A: SUBSISTENCE BASELINE COST

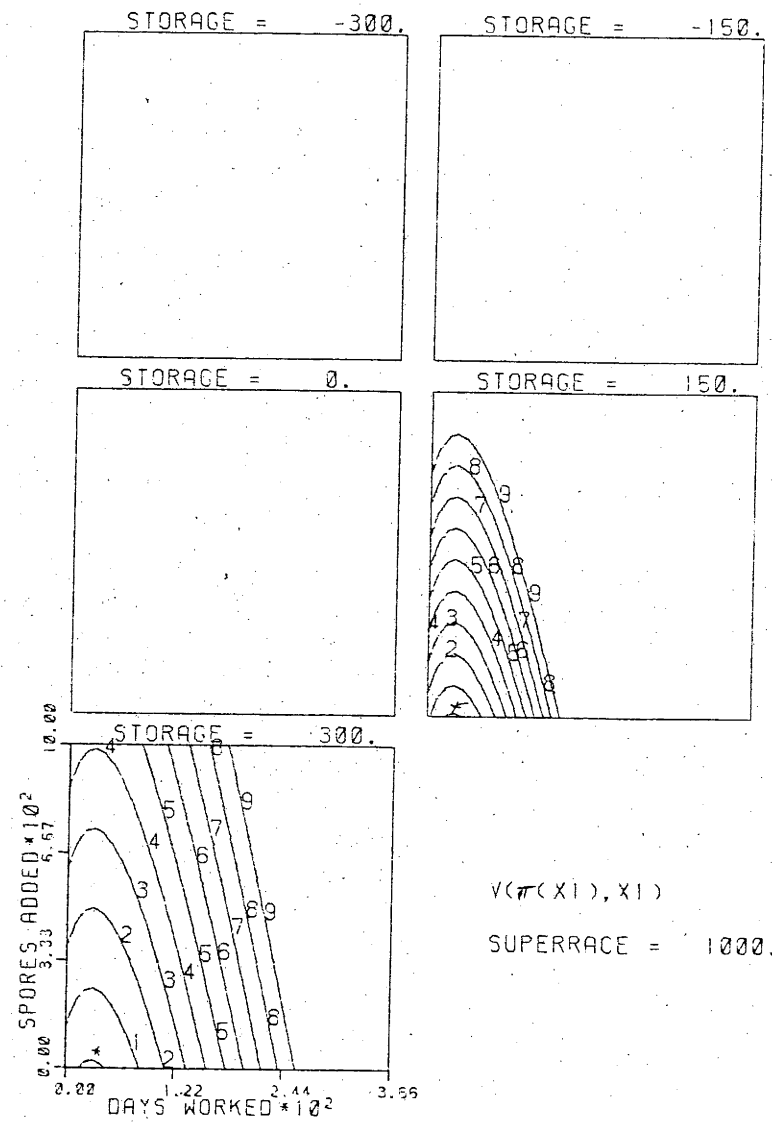

FIG. 6.5.13B: SUBSISTENCE BASELINE COST

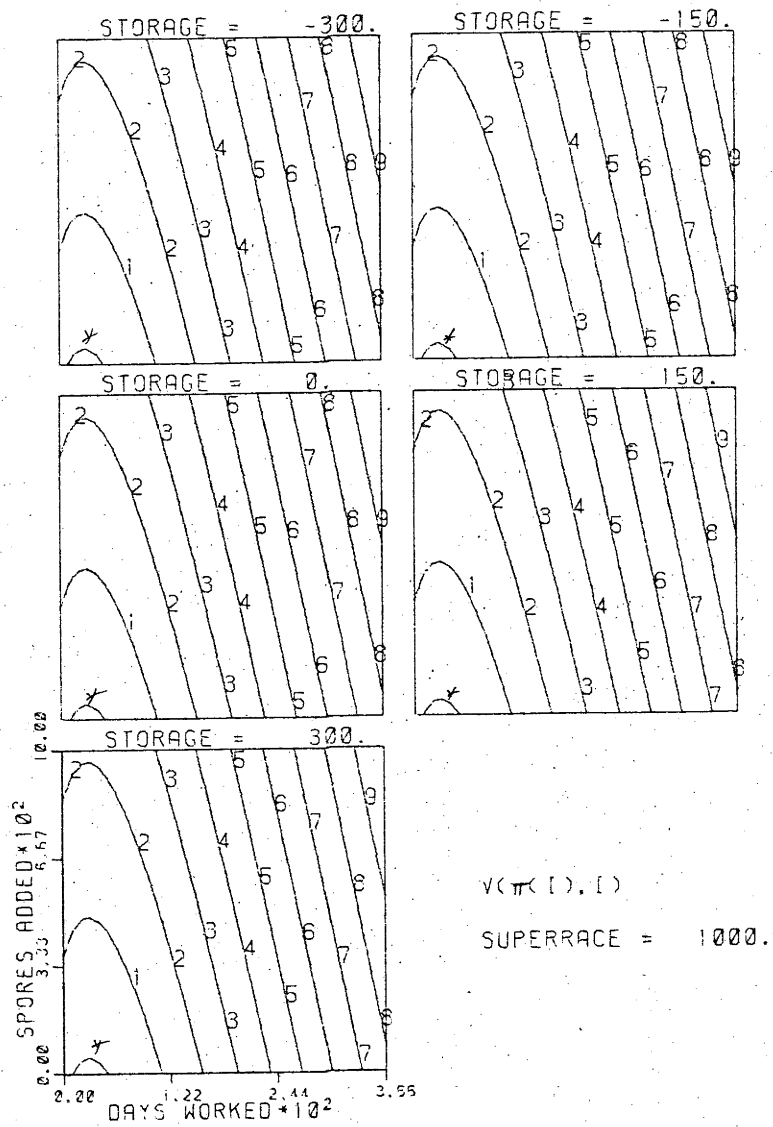

FIG. 6.5.13C: SUBSISTENCE BASELINE COST

FIG. 6.5.14: SUBSISTENCE BASELINE: WORK THROUGH TIME

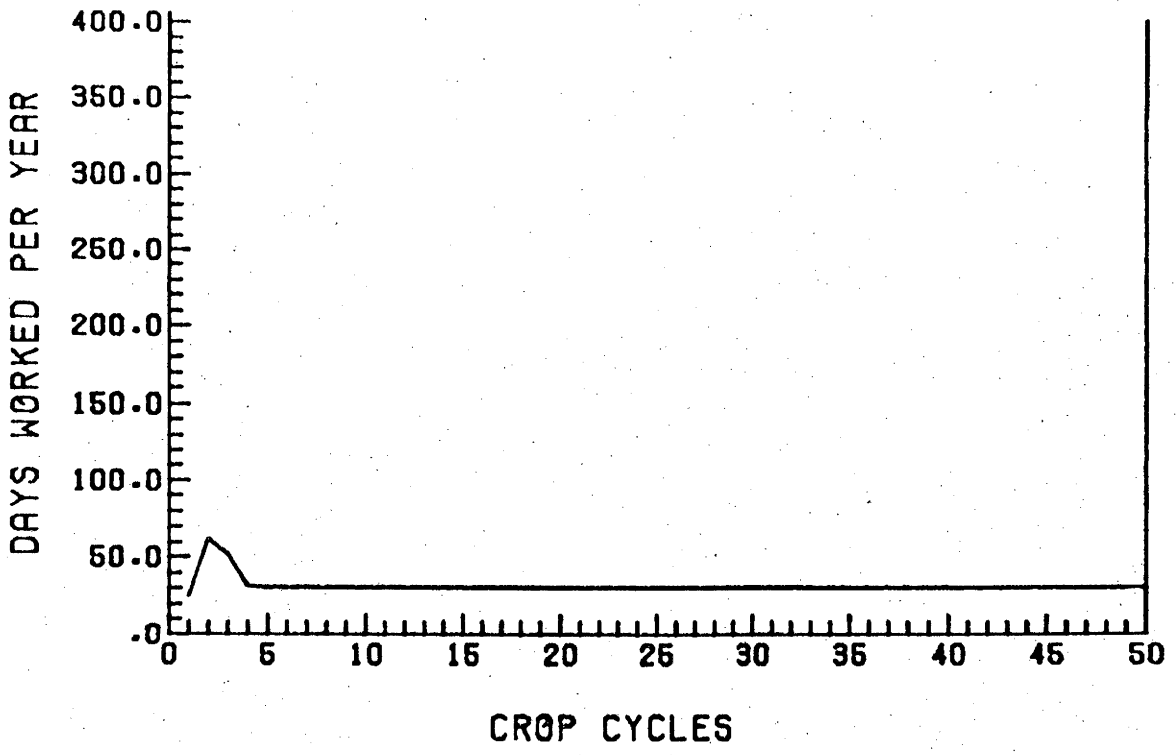

FIG. 6.5.15: SUBSISTENCE BASELINE: COST THROUGH TIME

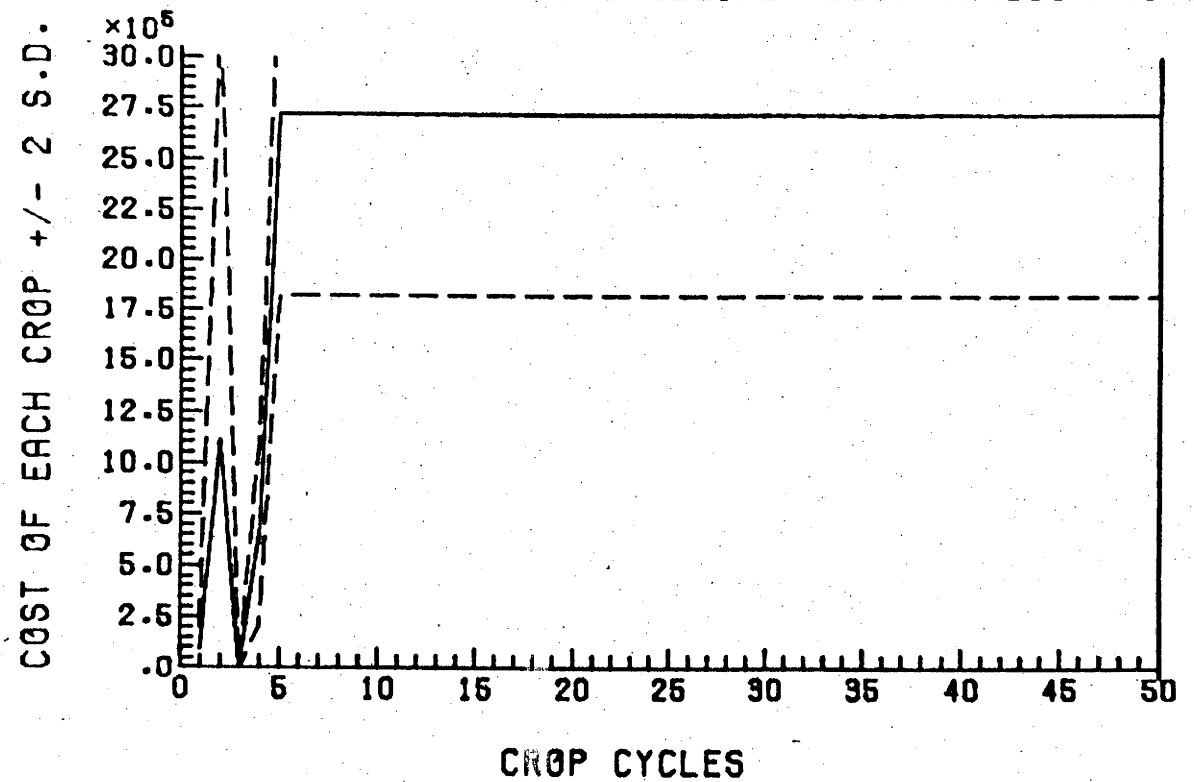

FIG. 6.5.16: SUBSISTENCE BASELINE: YIELD THROUGH TIME

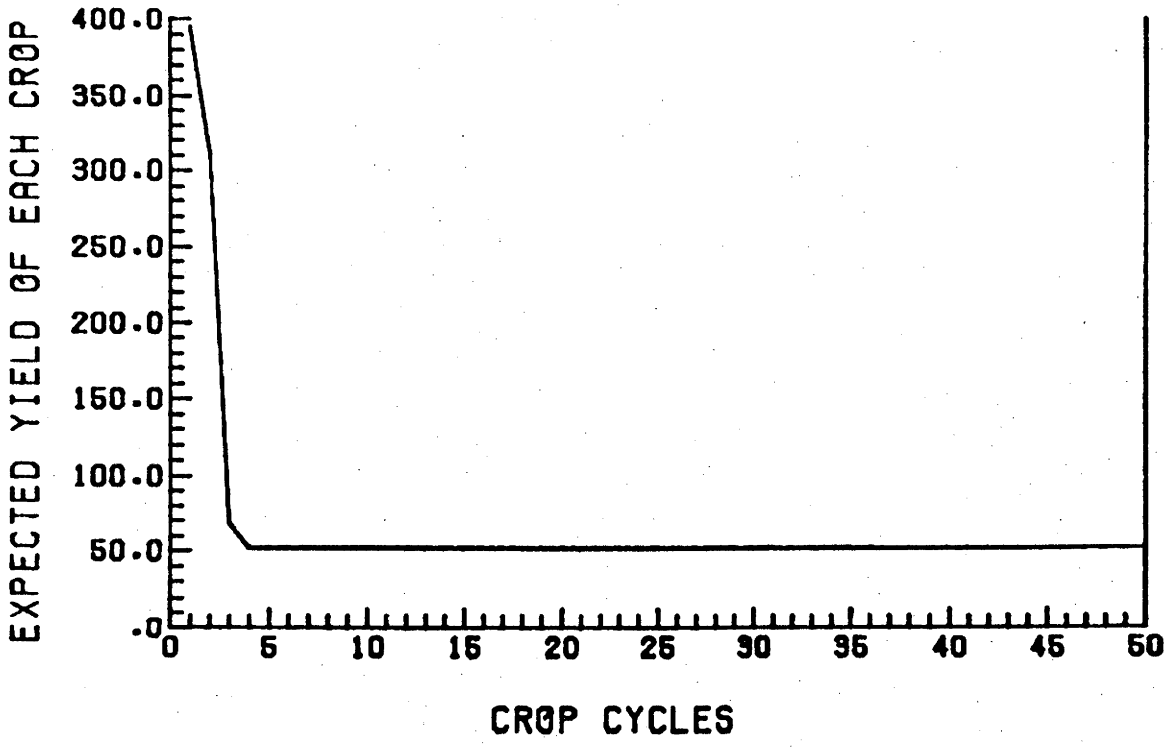

FIG. 6.5.17: SUBSISTENCE BASELINE: STORAGE THROUGH TIME

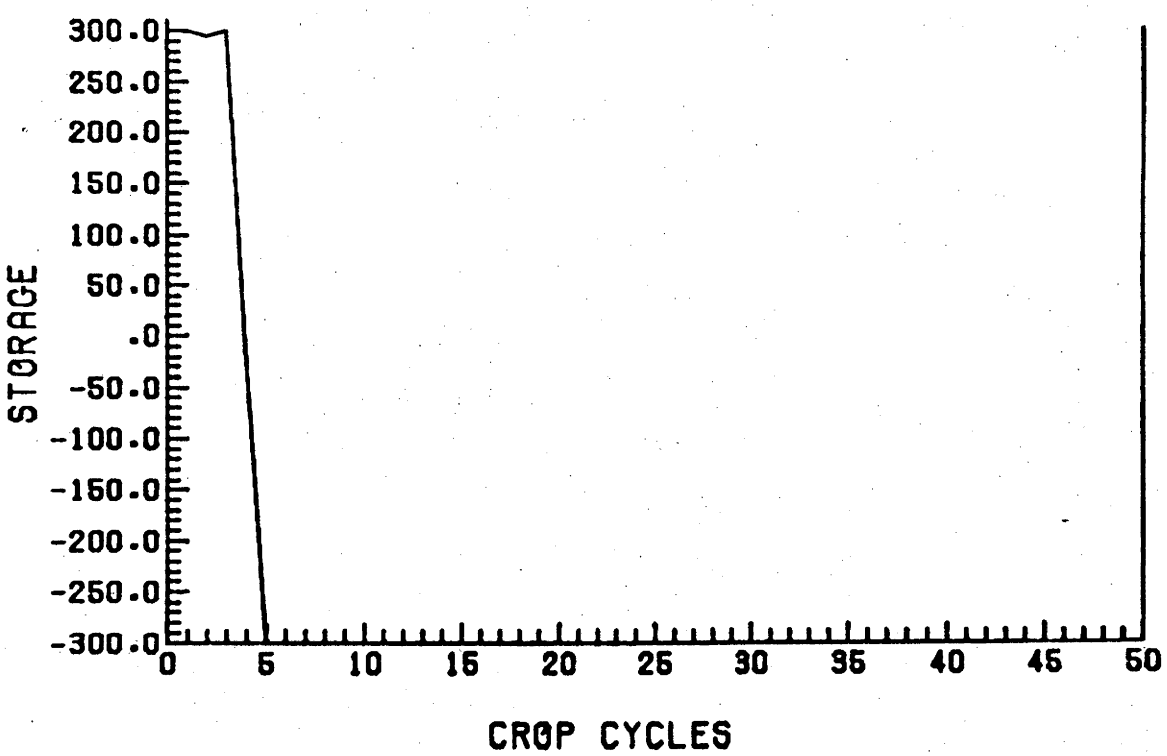

FIG. 6.5.18: SUBSISTENCE BASELINE: SUPERRACE THROUGH TIME

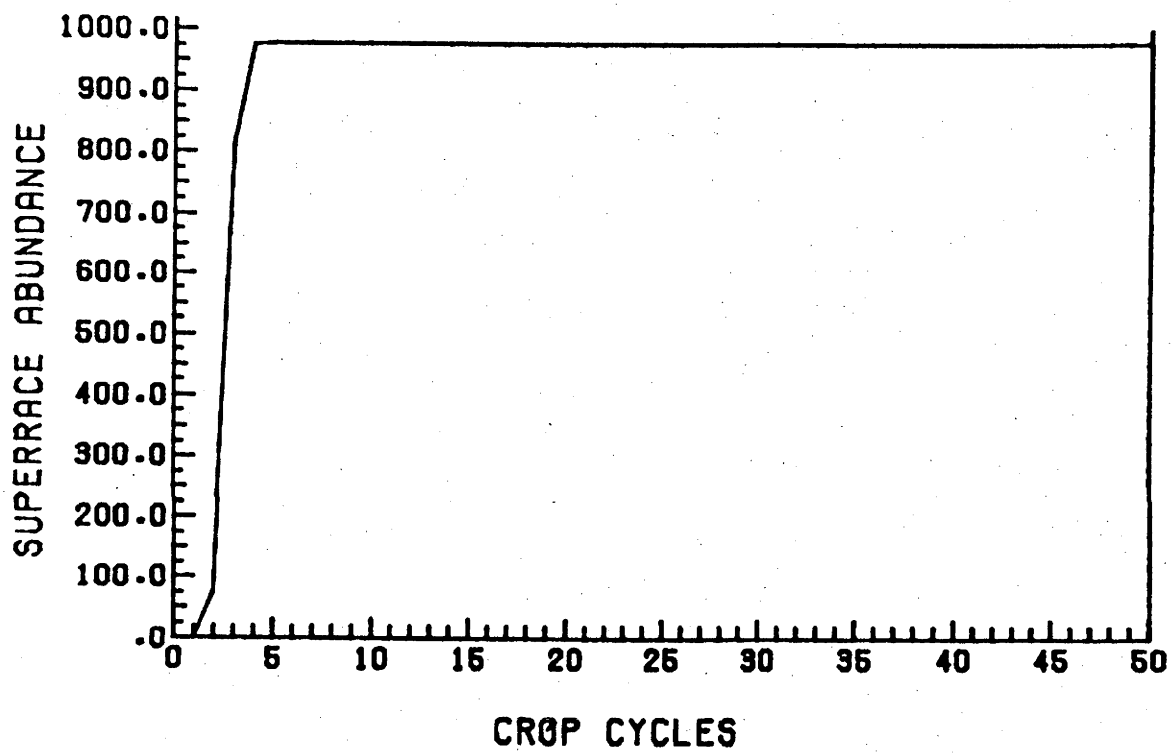

FIG. 6.5.19A: SUBSISTENCE + GENERAL RESISTANCE
DAYS WORKED VS STORAGE

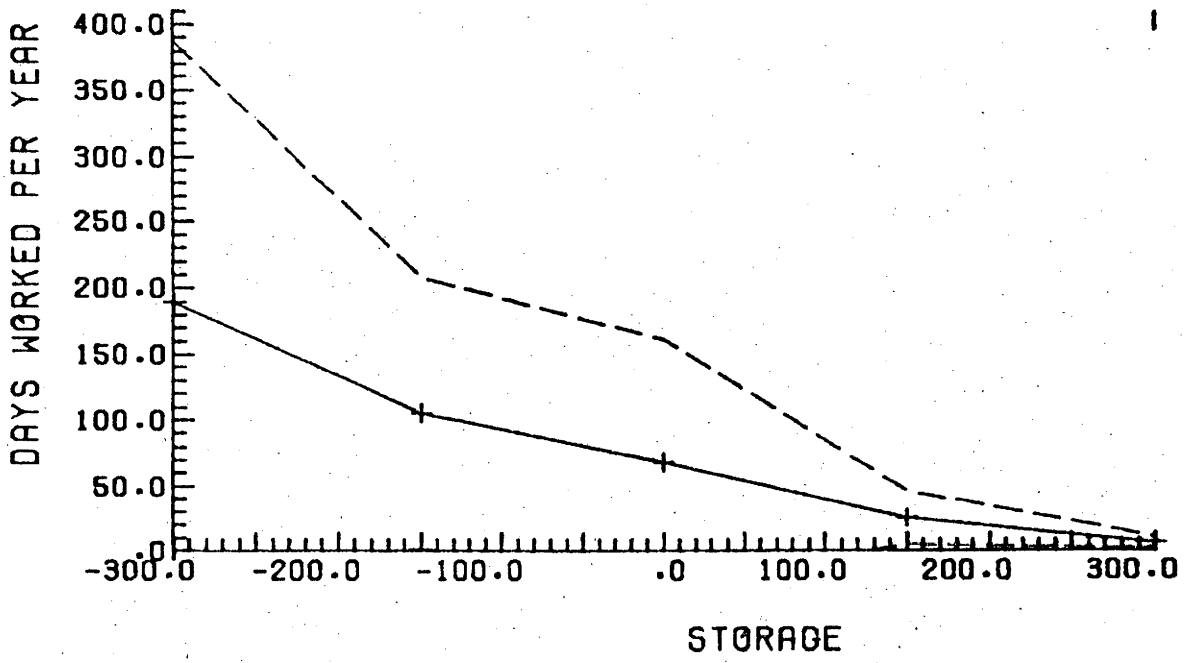

FIG. 6.5.19B: SUBSISTENCE + GENERAL RESISTANCE
DAYS WORKED VS AABB

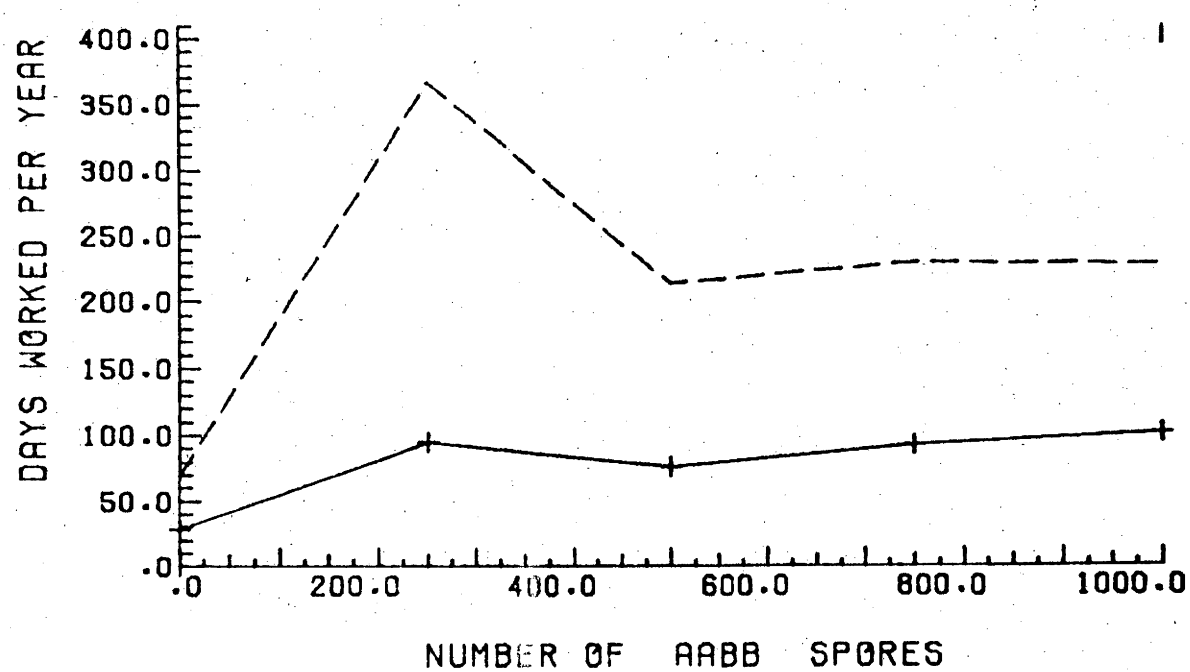

FIG. 6.5.21A: SUBSISTENCE + GENERAL RESISTANCE STORAGE INPUT/OUTPUT

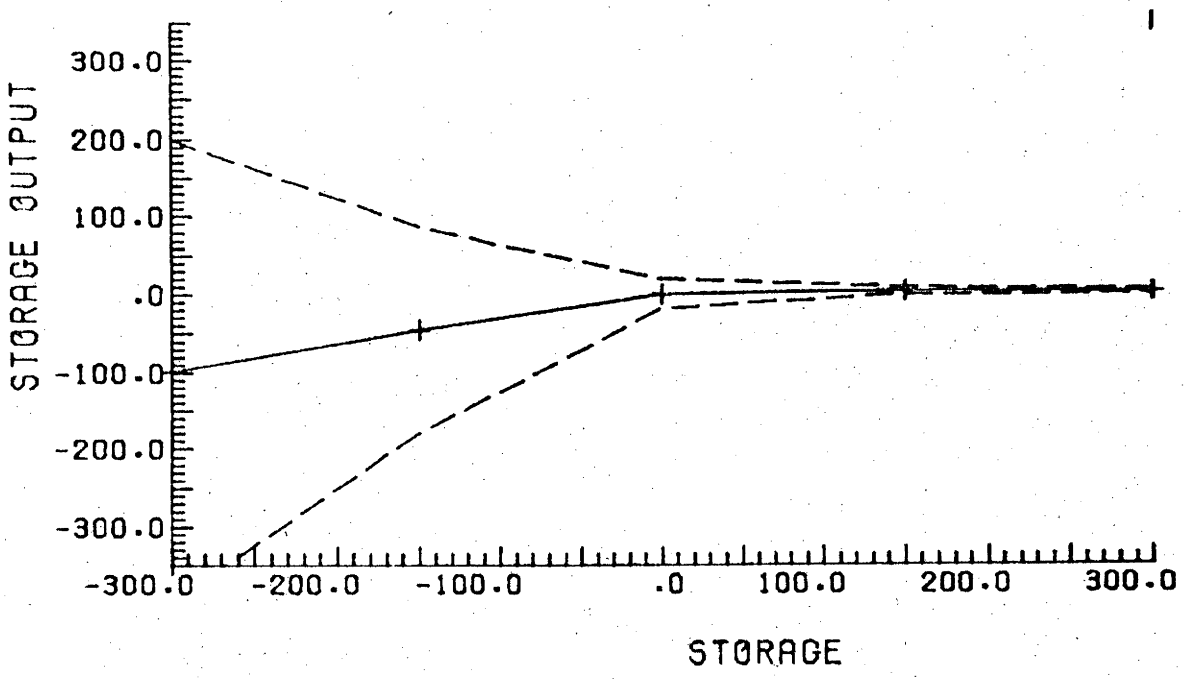

FIG. 6.5.21B: SUBSISTENCE + GENERAL RESISTANCE STORAGE OUTPUT VS AABB

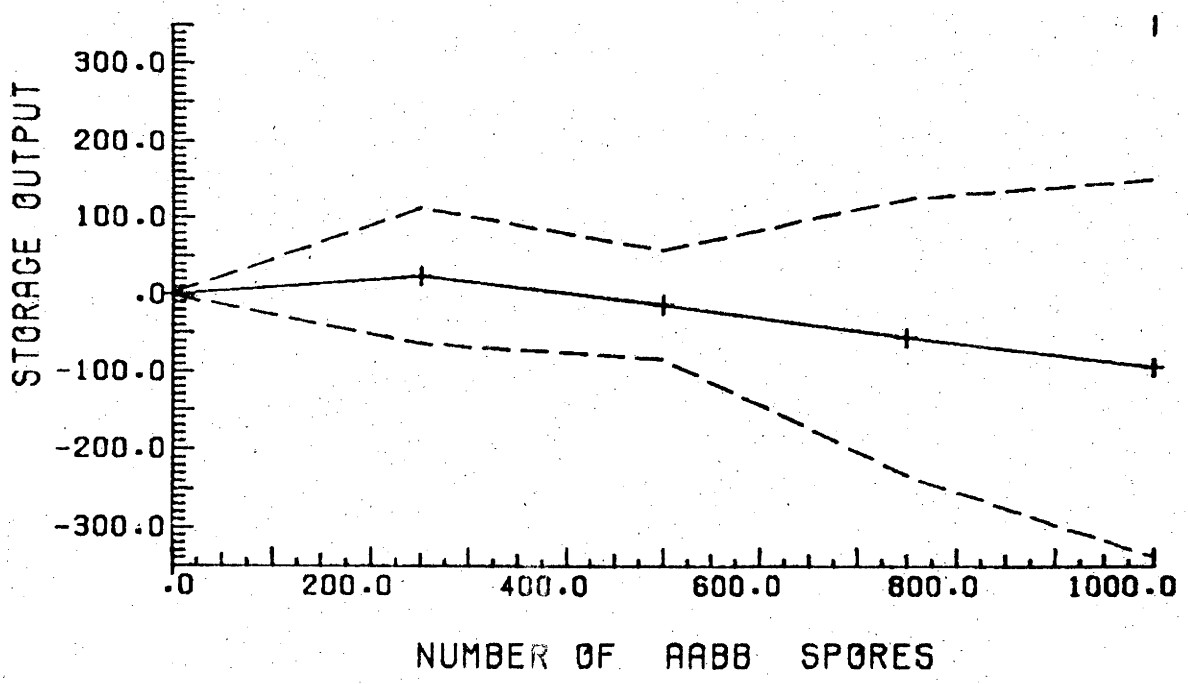

FIG. 6.5.22: SUBSISTENCE + GENERAL RESISTANCE SUPERRACE INPUT/OUTPUT

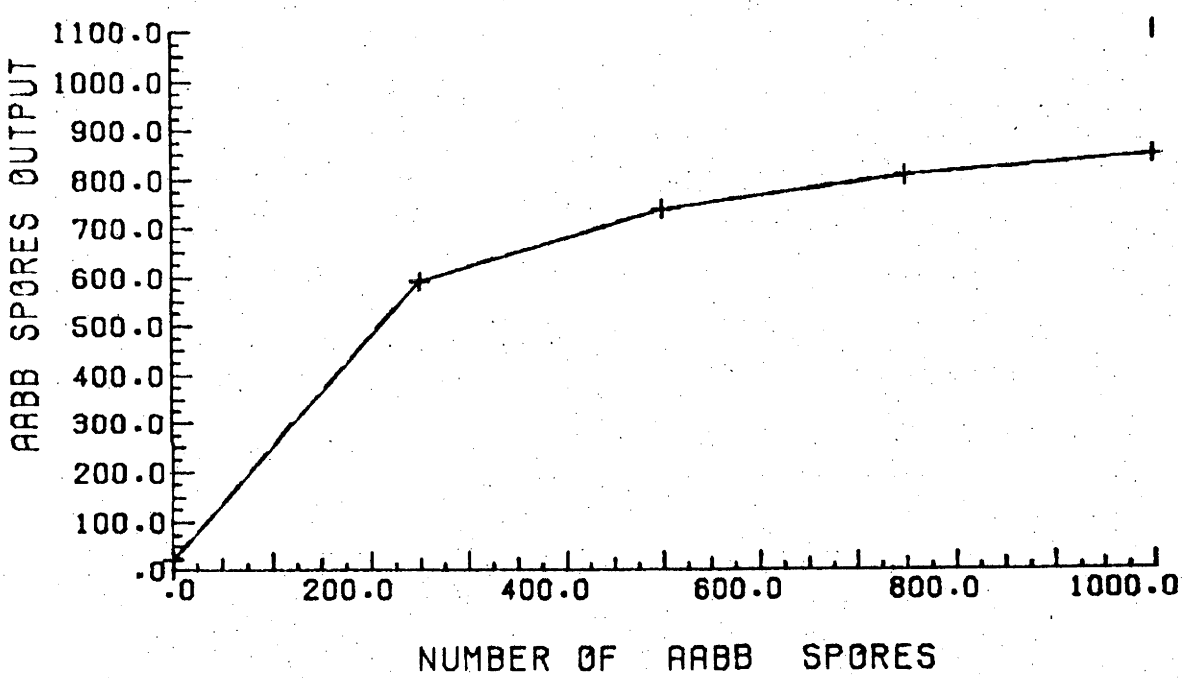

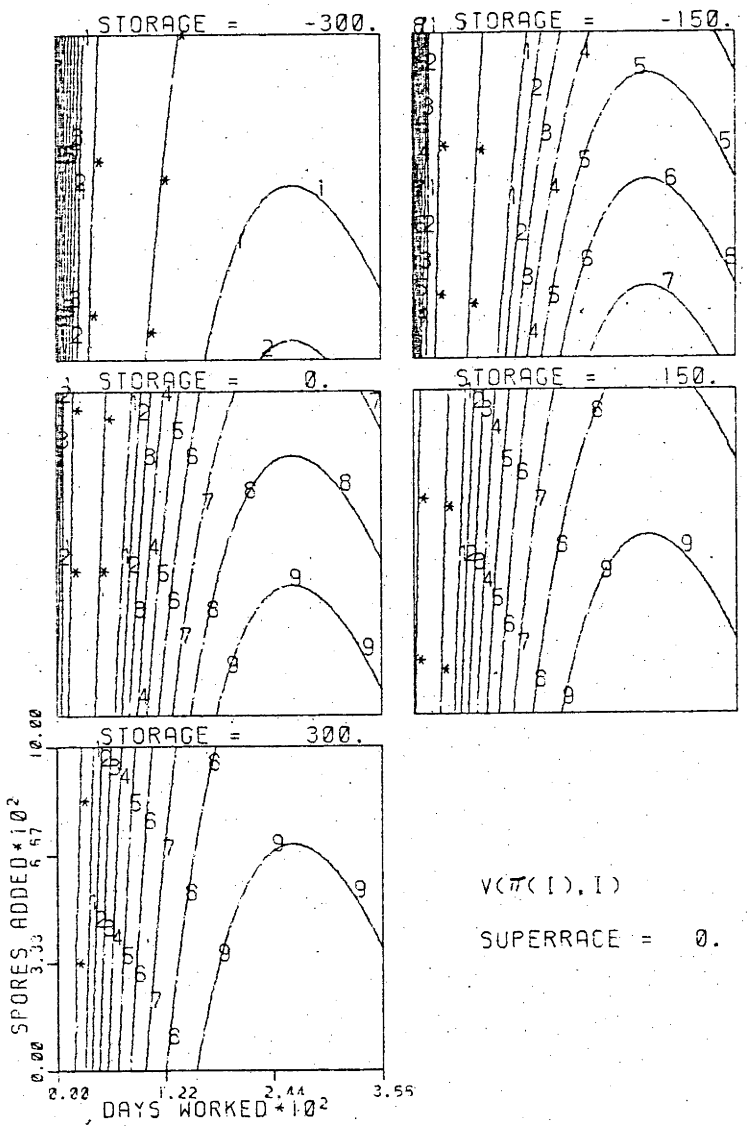

FIG. 6.5.23· SUBSISTENCE COST WITH GENERAL RESISTANCE

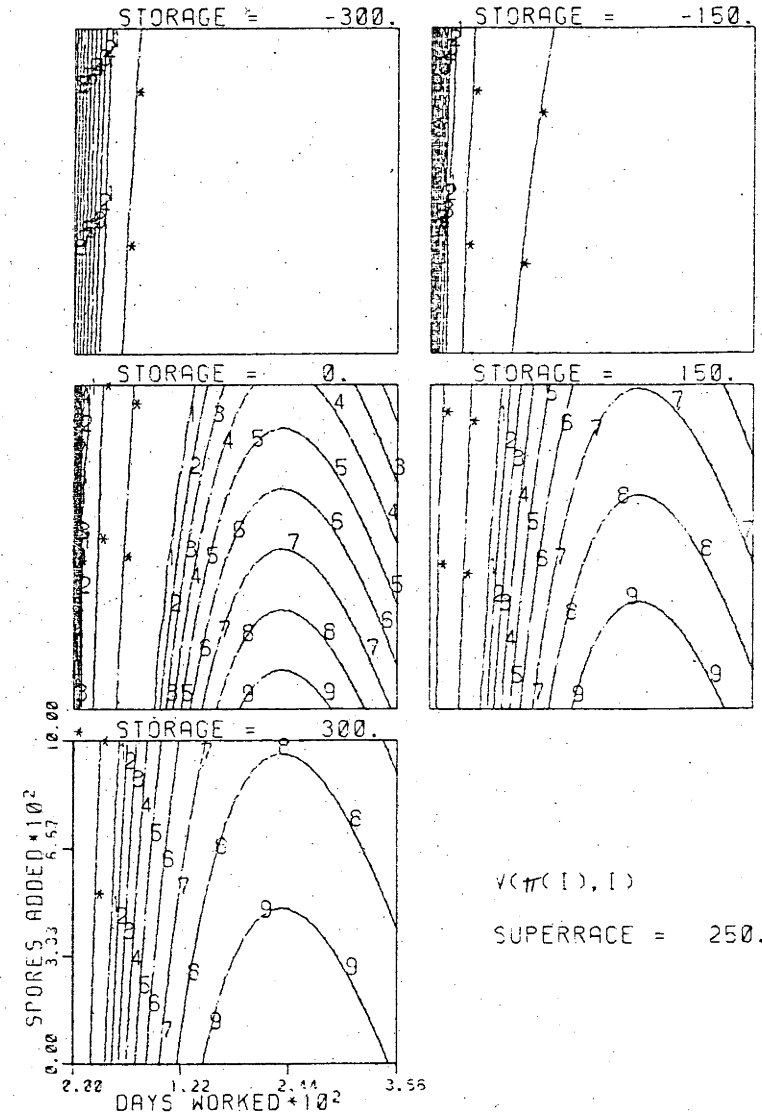

FIG. 6.5.24· SUBSISTENCE COST WITH GENERAL RESISTANCE

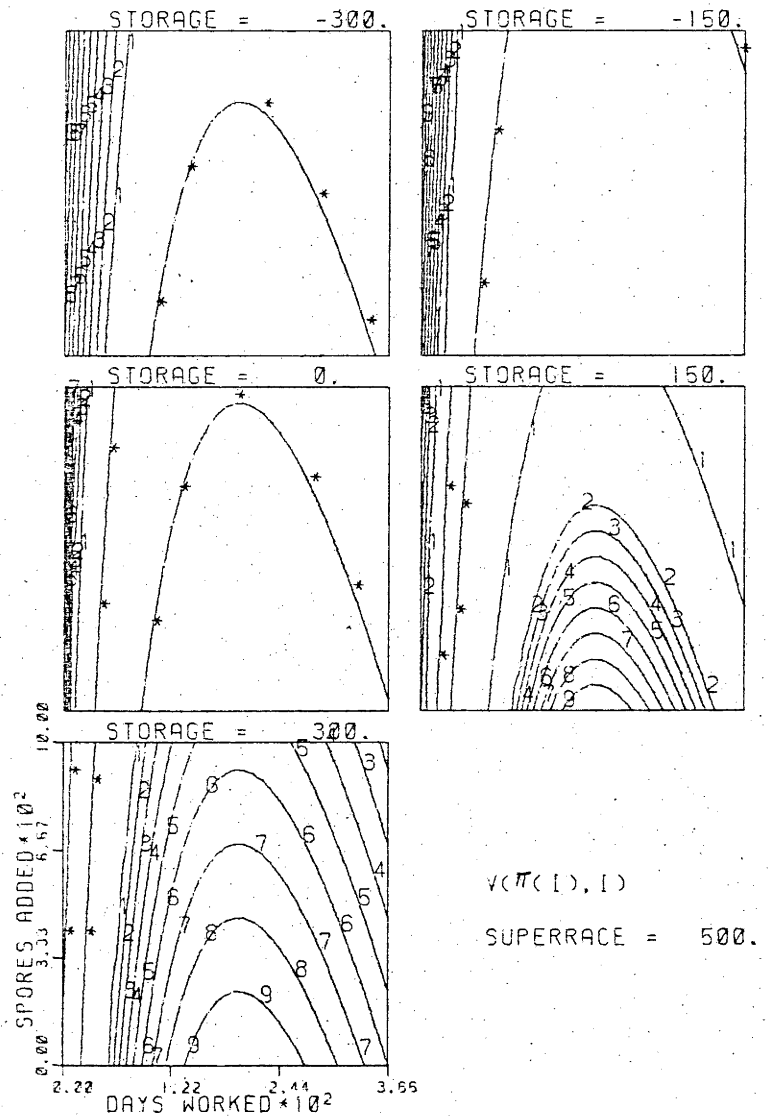

FIG. 6 5.25· SUBSISTENCE COST WITH GENERAL RESISTANCE

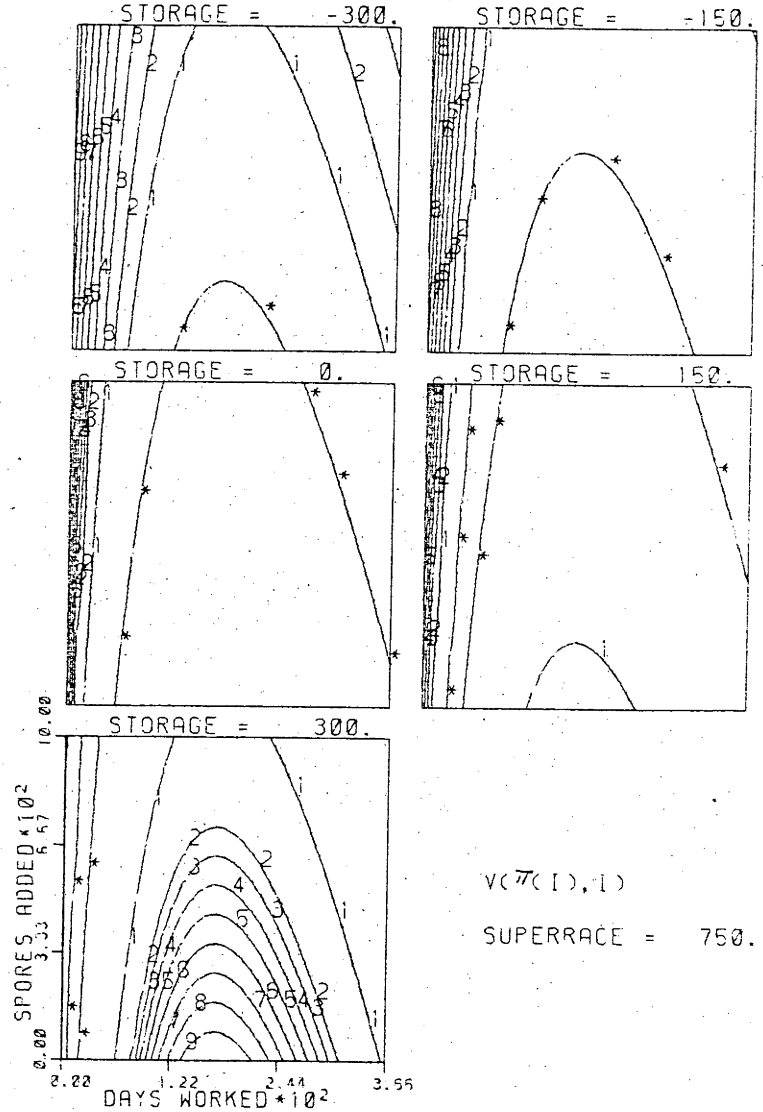

FIG. 6.5.26· SUBSISTENCE COST WITH GENERAL RESISTANCE

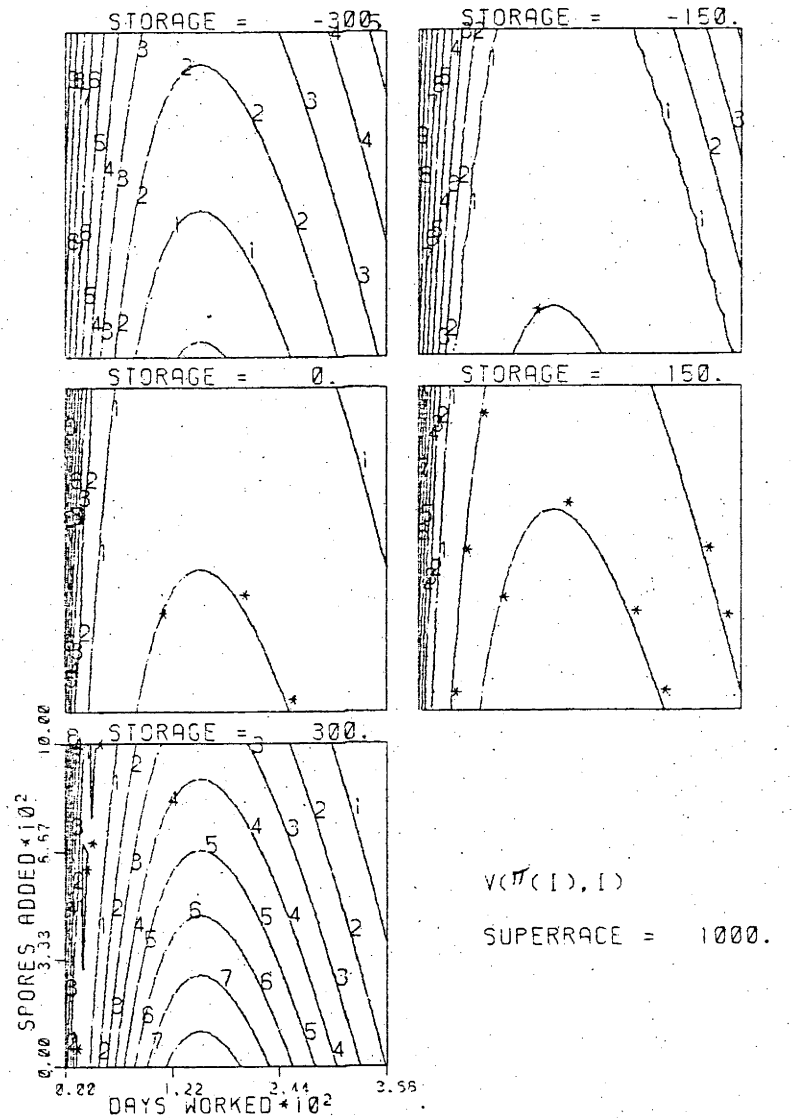

FIG. 6.5.27: SUBSISTENCE COST WITH GENERAL RESISTANCE

FIG. 6.5.28: SUBSISTENCE + GENERAL RESISTANCE: WORK THROUGH TIME

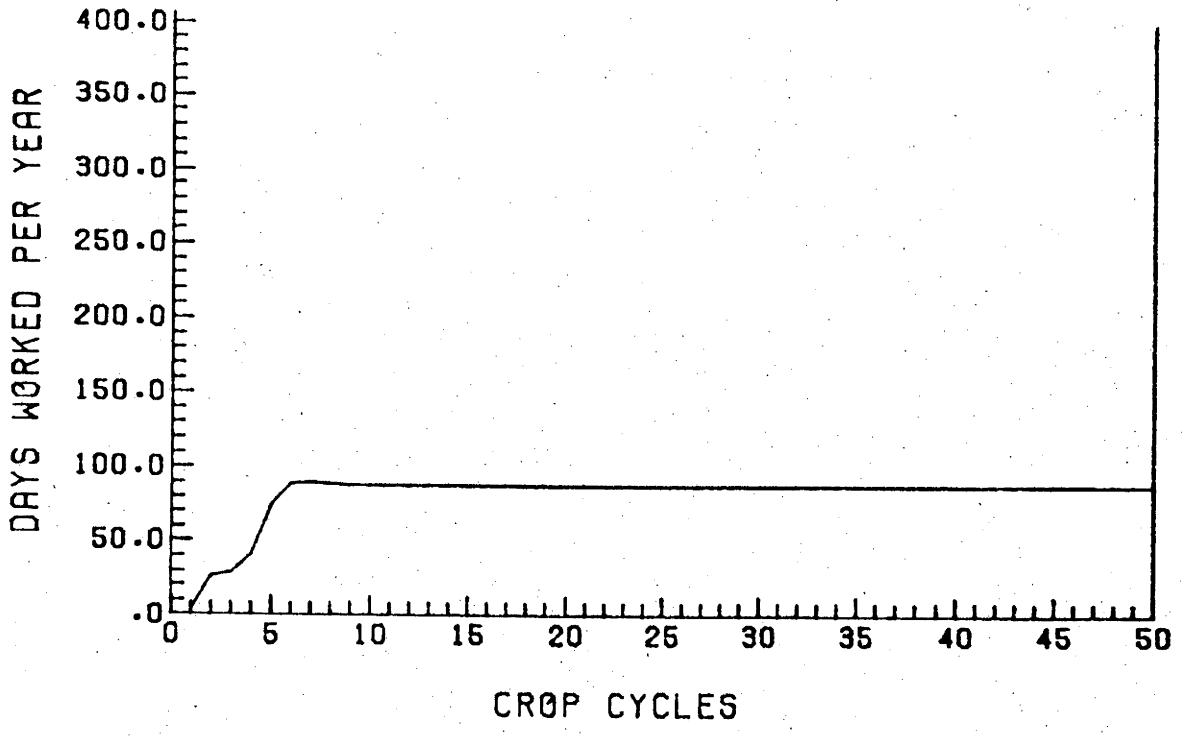

FIG. 6.5.29: SUBSISTENCE + GENERAL RESISTANCE: COST THROUGH TIME

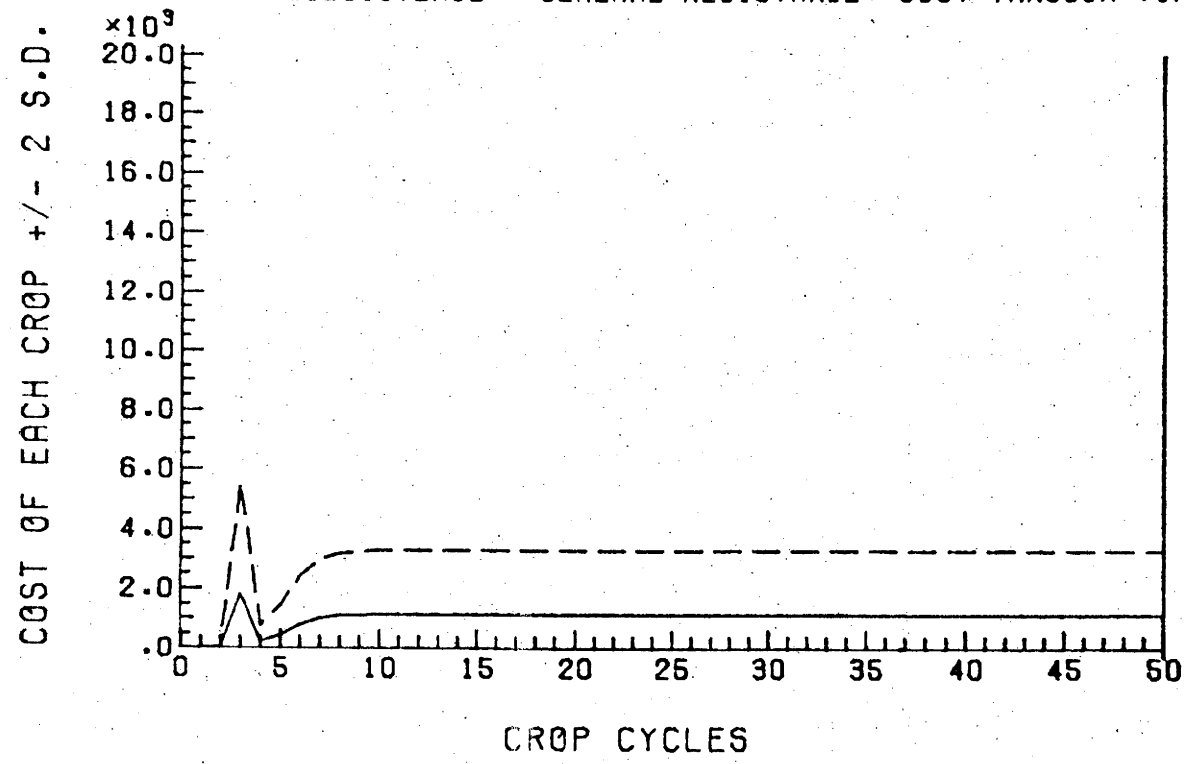

FIG. 6.5.30: SUBSISTENCE + GENERAL RESISTANCE: YIELD THROUGH TIME

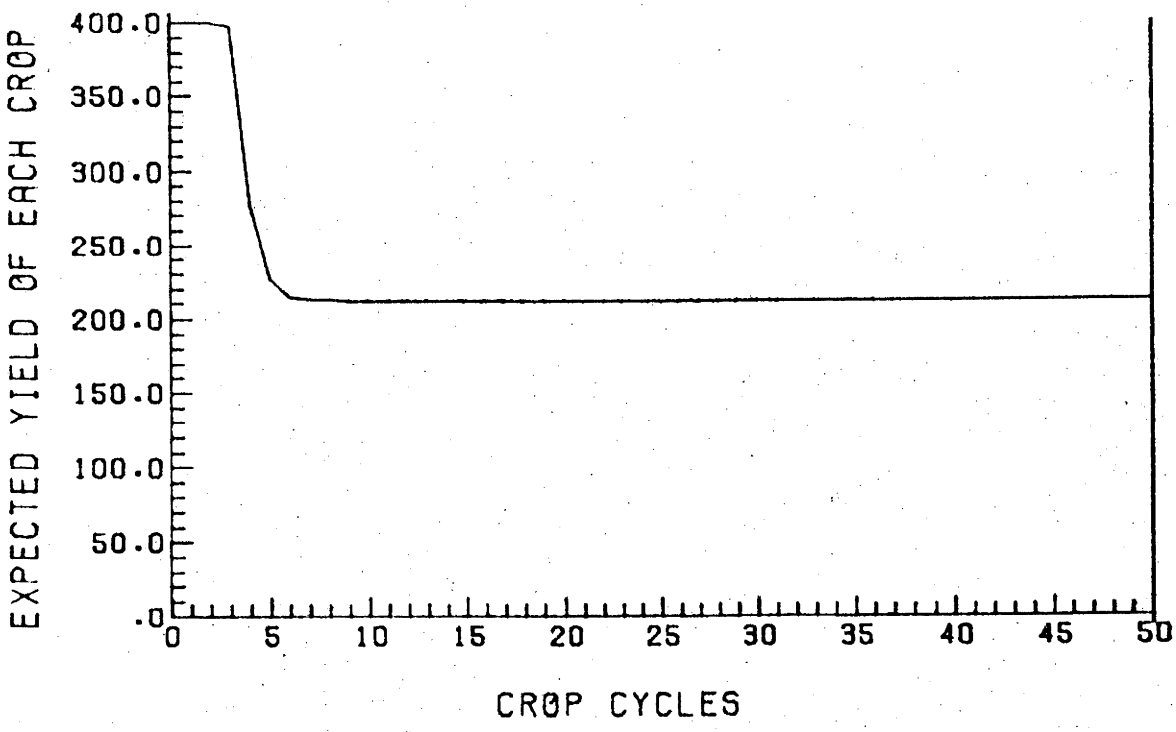

FIG. 6.4.31: SUBSISTENCE + GENERAL RESISTANCE: STORAGE THROUGH TIME

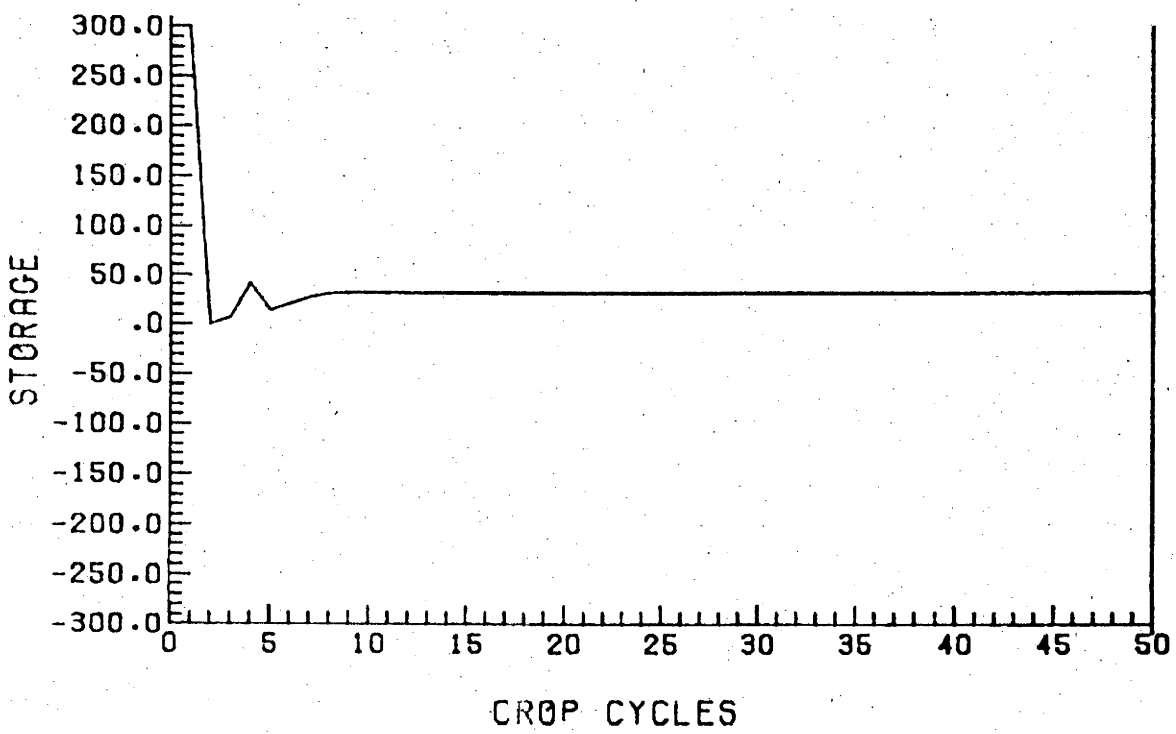

FIG. 6.5.32A: GENERAL RESISTANCE MULTILINE
F(A) USED VS STORAGE

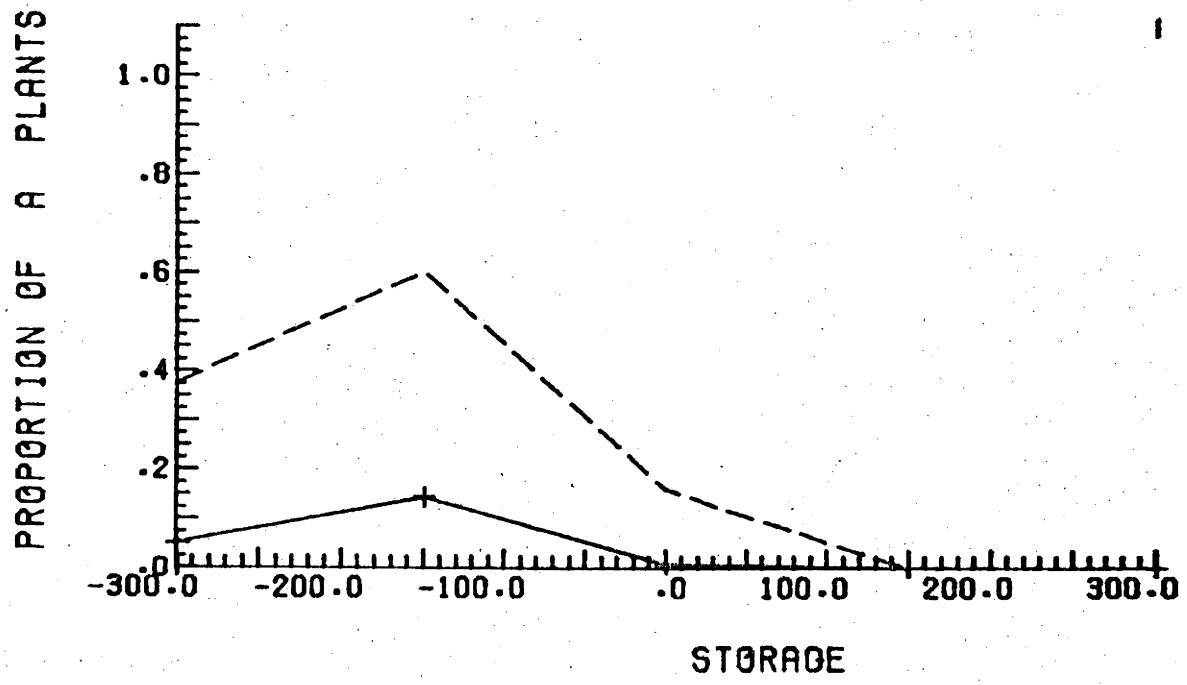

FIG. 6.5.32B: GENERAL RESISTANCE MULTILINE
F(A) USED VS AABB

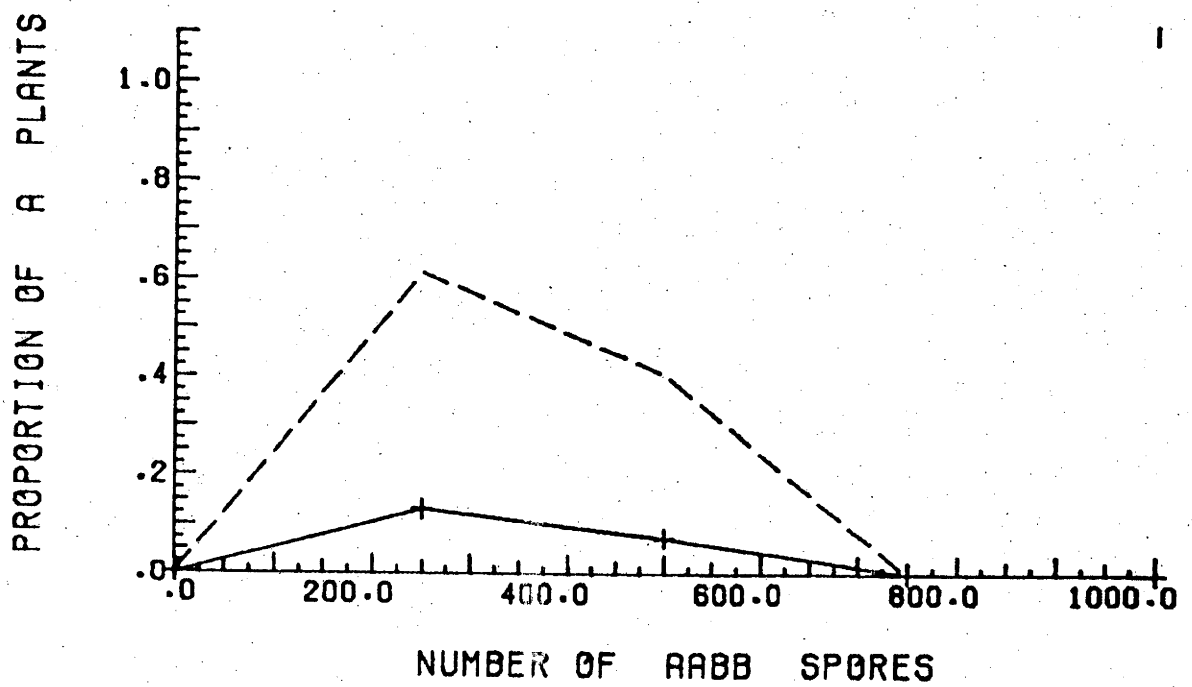

FIG. 6.5.33: GENERAL RESISTANCE MULTILINE STORAGE OUTPUT VS AABB

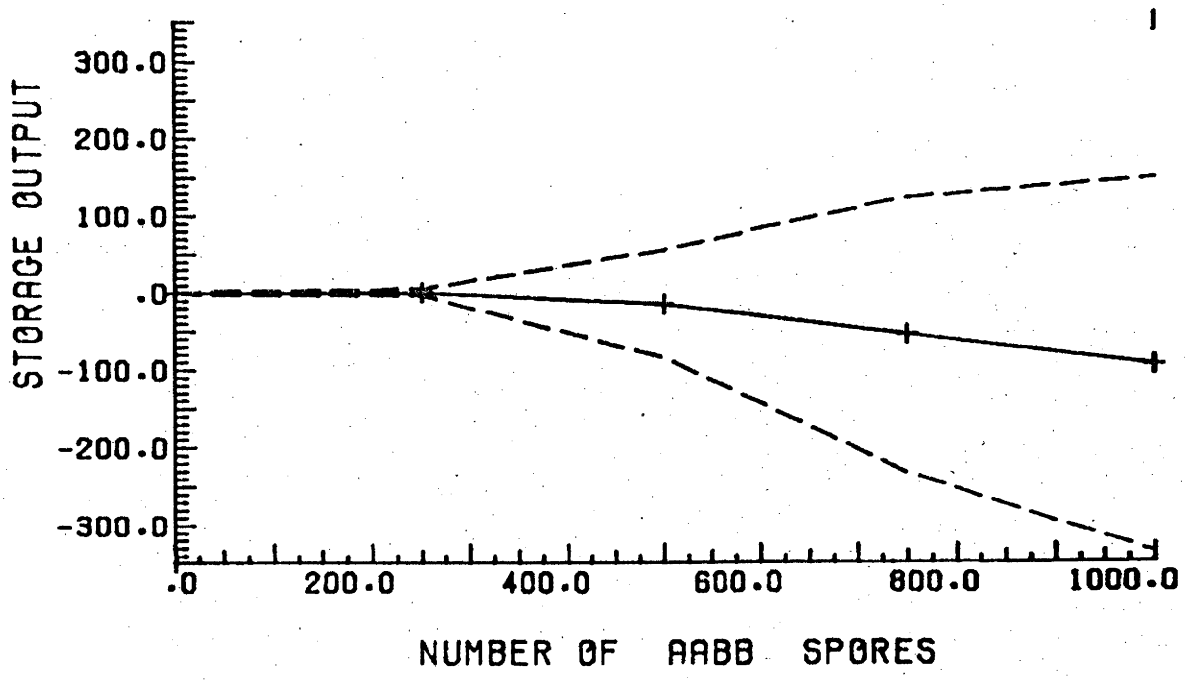

FIG. 6.5.34: GENERAL RESISTANCE MULTILINE AABB INPUT/OUTPUT

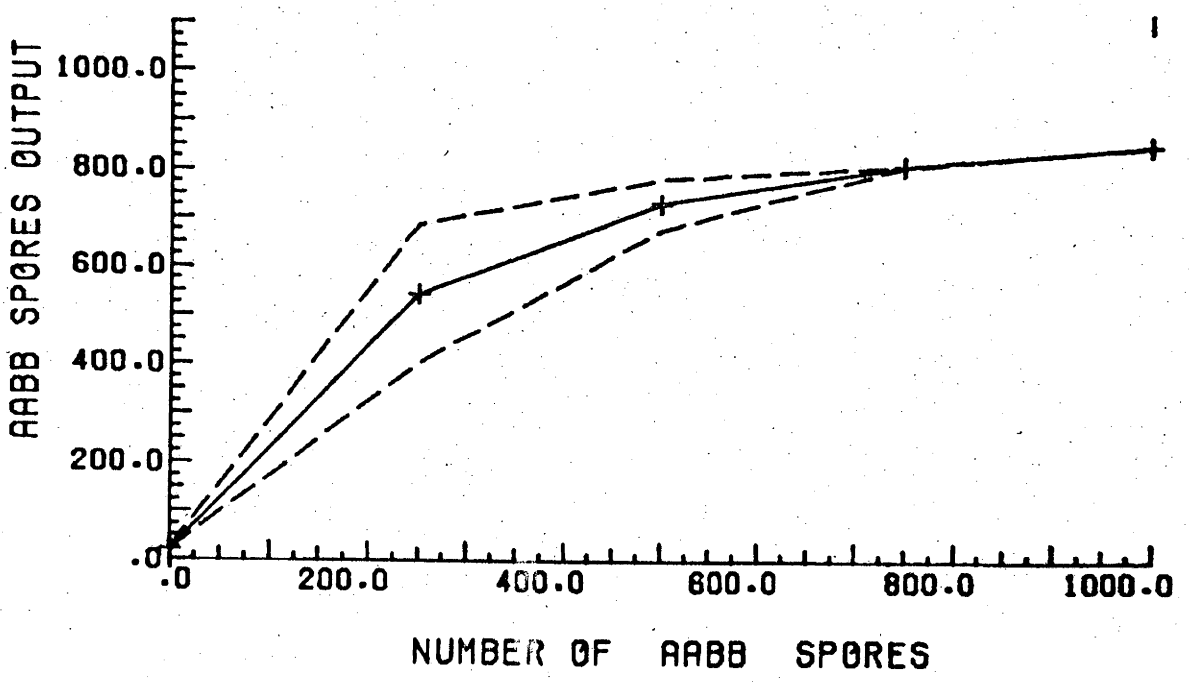

FIG. 6.5.35: GENERAL RESISTANCE MULTILINE
WORK DONE VS AABB

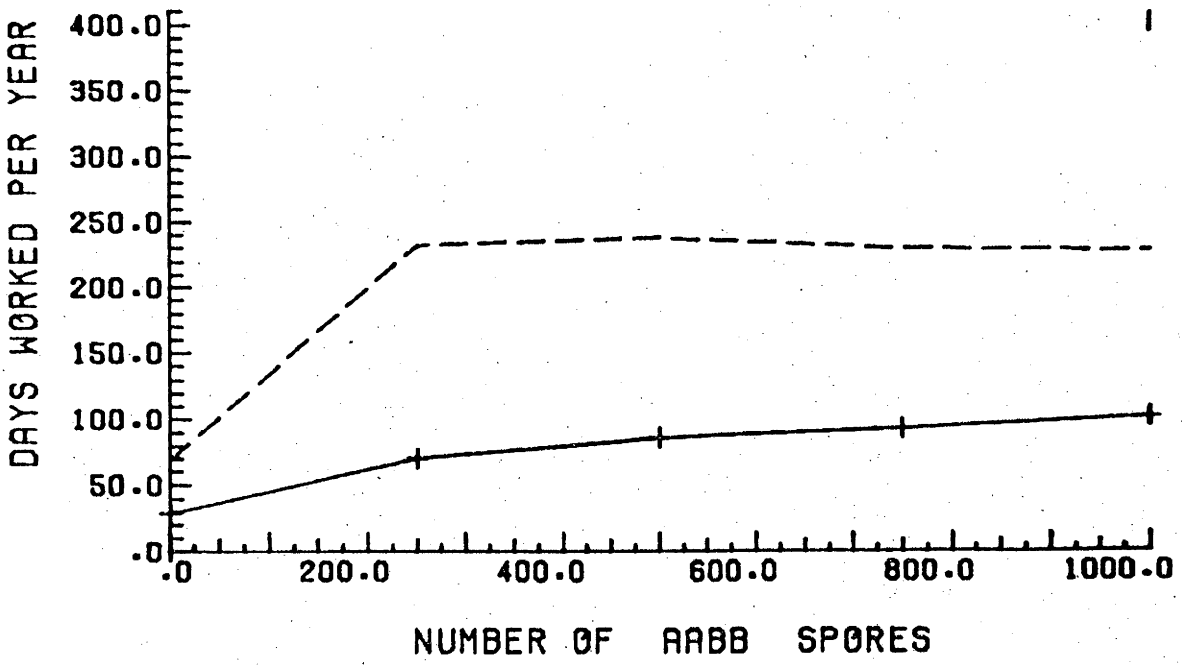

FIG. 6.5.36A: GENERAL RESISTANCE MULTILINE
AA SPORES USED VS STORAGE

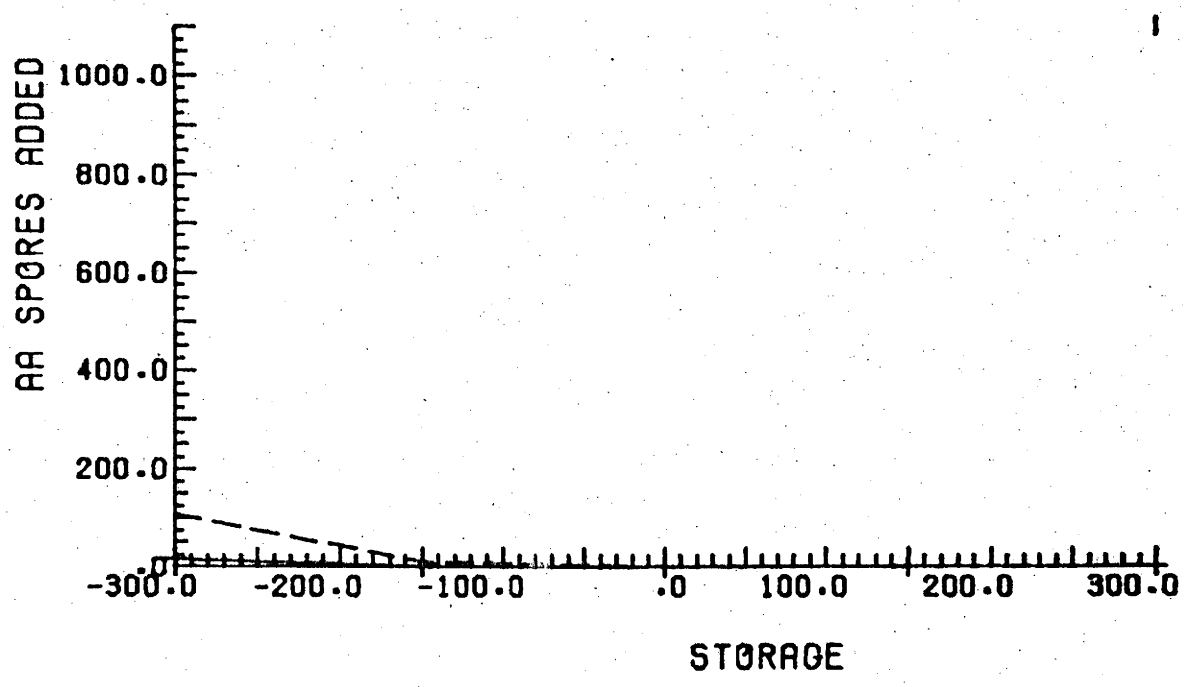

FIG. 6.5.36B: GENERAL RESISTANCE MULTILINE
AA ADDED VS AABB

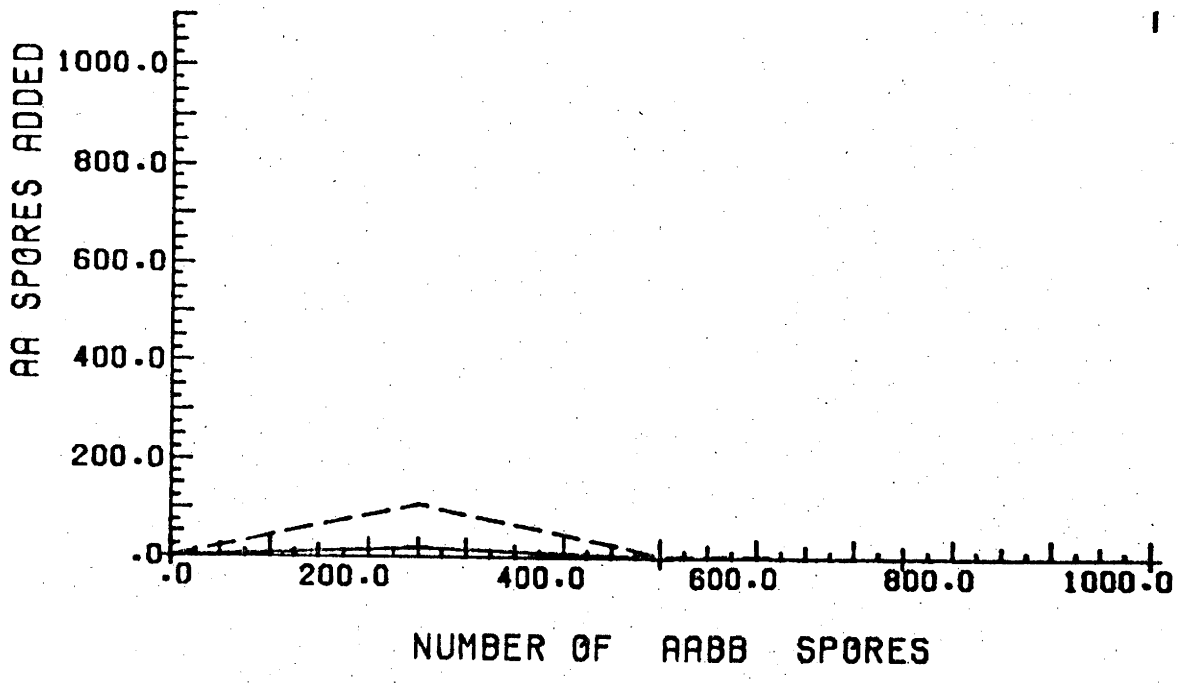

FIG. 6.5.37: REGRESSION OF SUBLINE PLANTED ON
SIMPLER RACES PRESENT

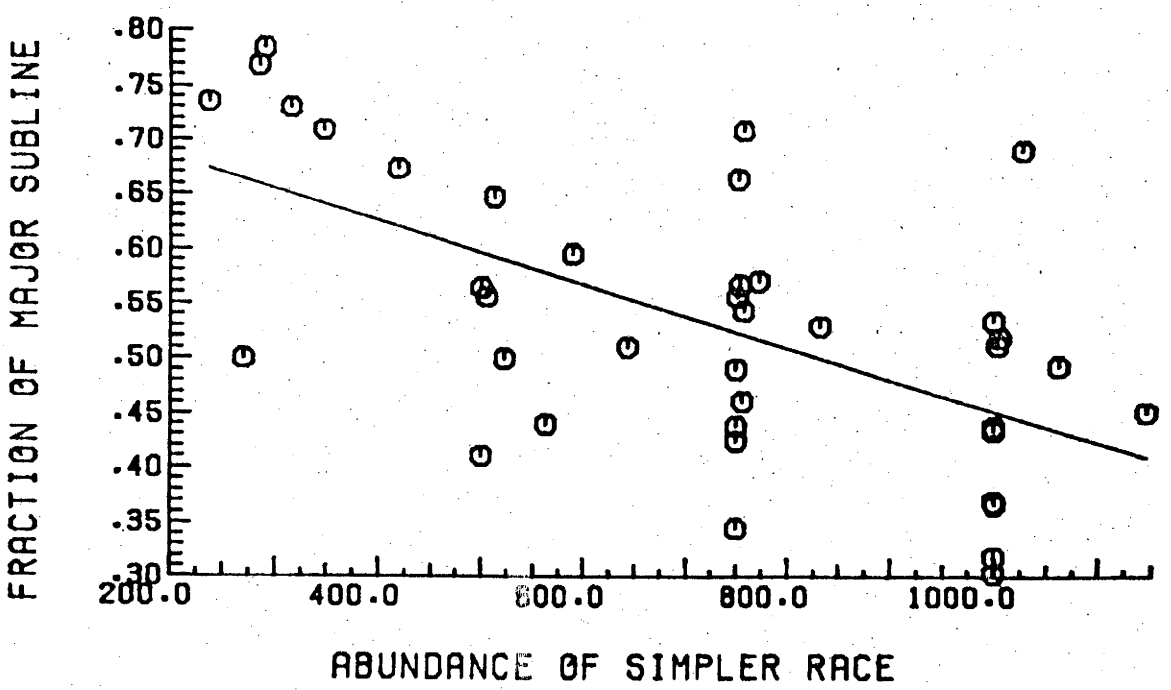

FIG. 6.5.38A: GENERAL RESISTANCE MULTILINE: SUPERLINE THROUGH TIME

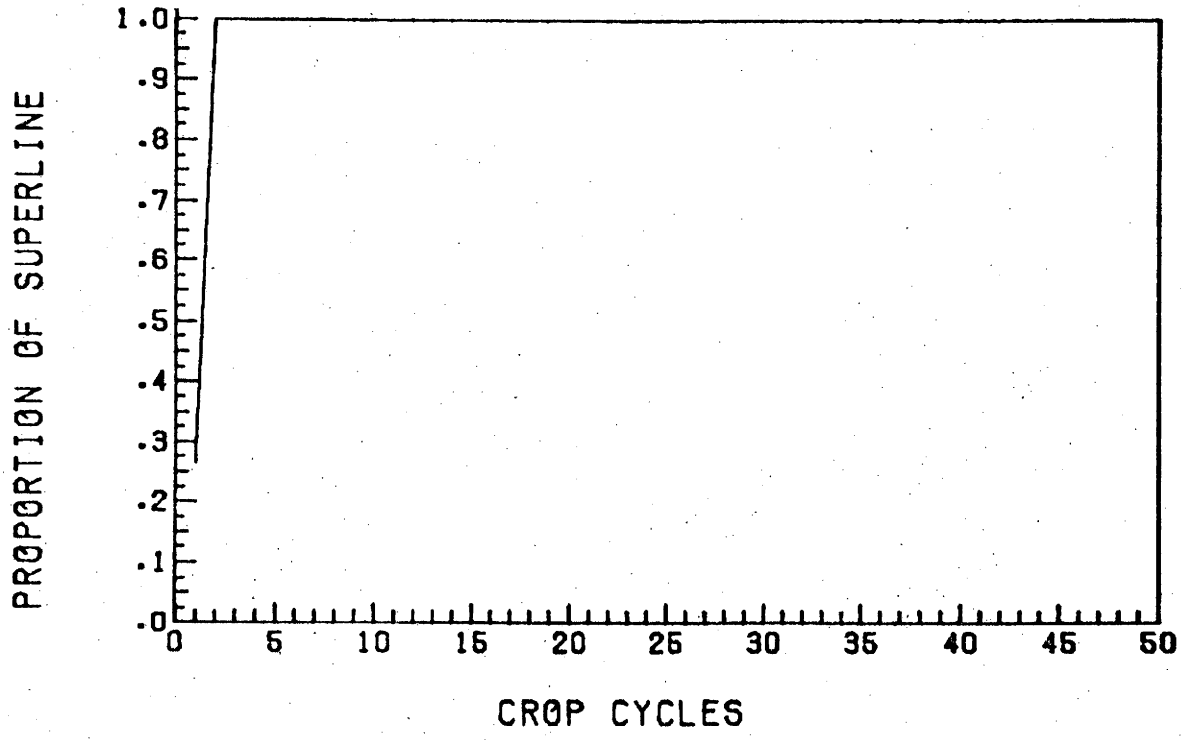

FIG. 6.5.38B: GENERAL RESISTANCE MULTILINE: SUPERRACE THROUGH TIME

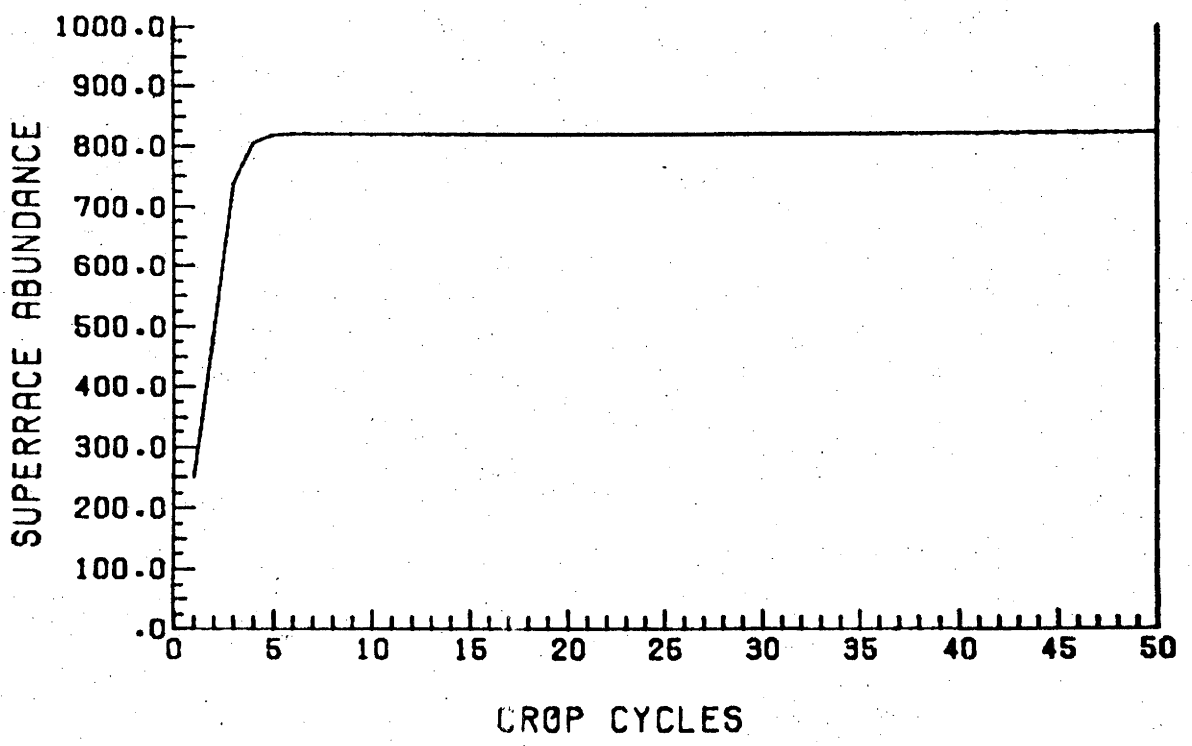

FIG. 6.5.40: CROSSPROTECTION MULTILINE STORAGE OUTPUT VS AABB

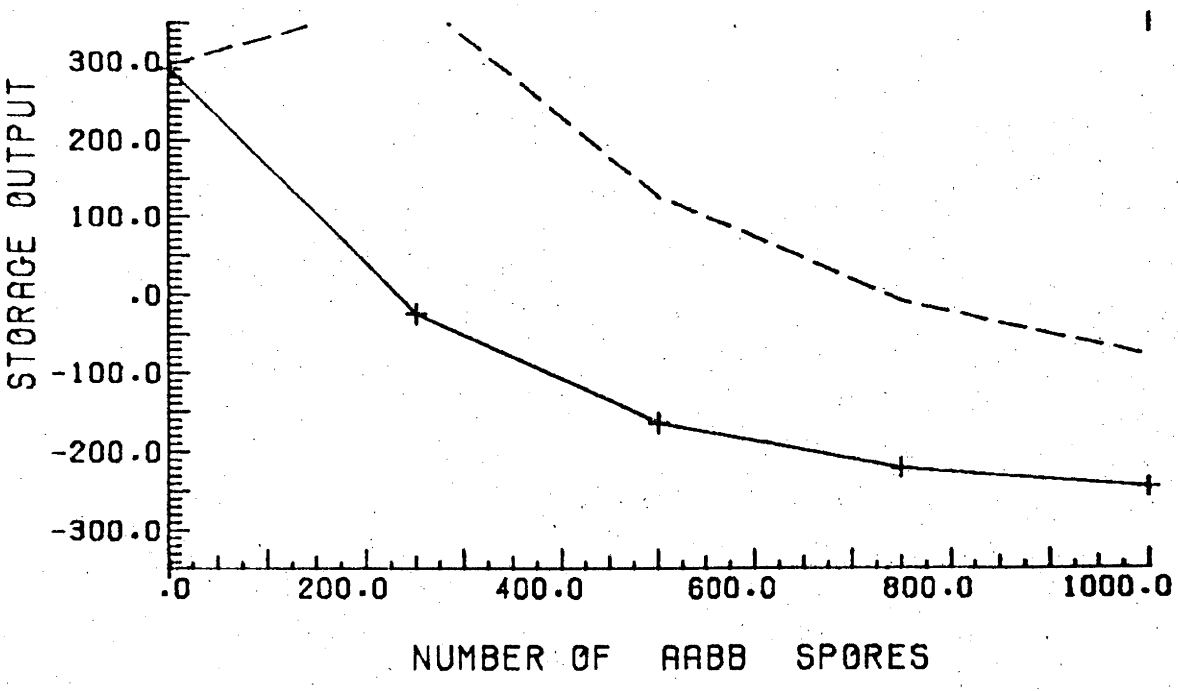

FIG. 6.5.41: CROSSPROTECTION MULTILINE AABB INPUT/OUTPUT

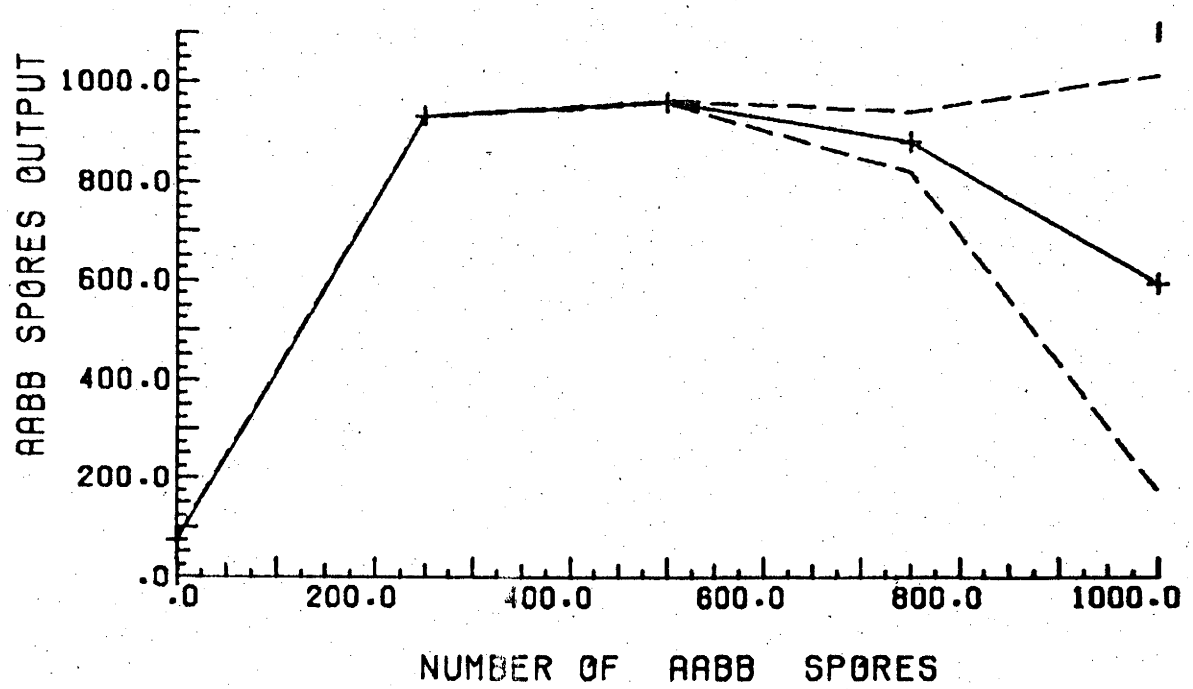

FIG. 6.5.42: CROSSPROTECTION MULTILINE
WORK DONE VS AABB

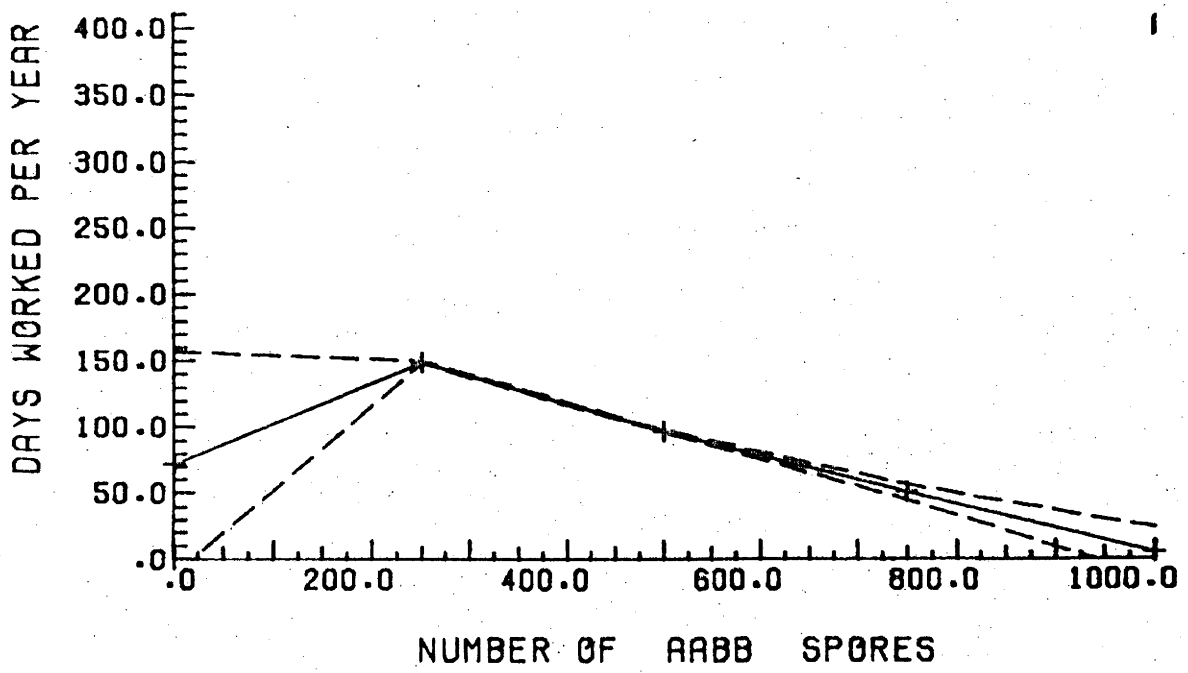

FIG. 6.5.43A: CROSSPROTECTION MULTILINE
AA SPORES USED VS STORAGE

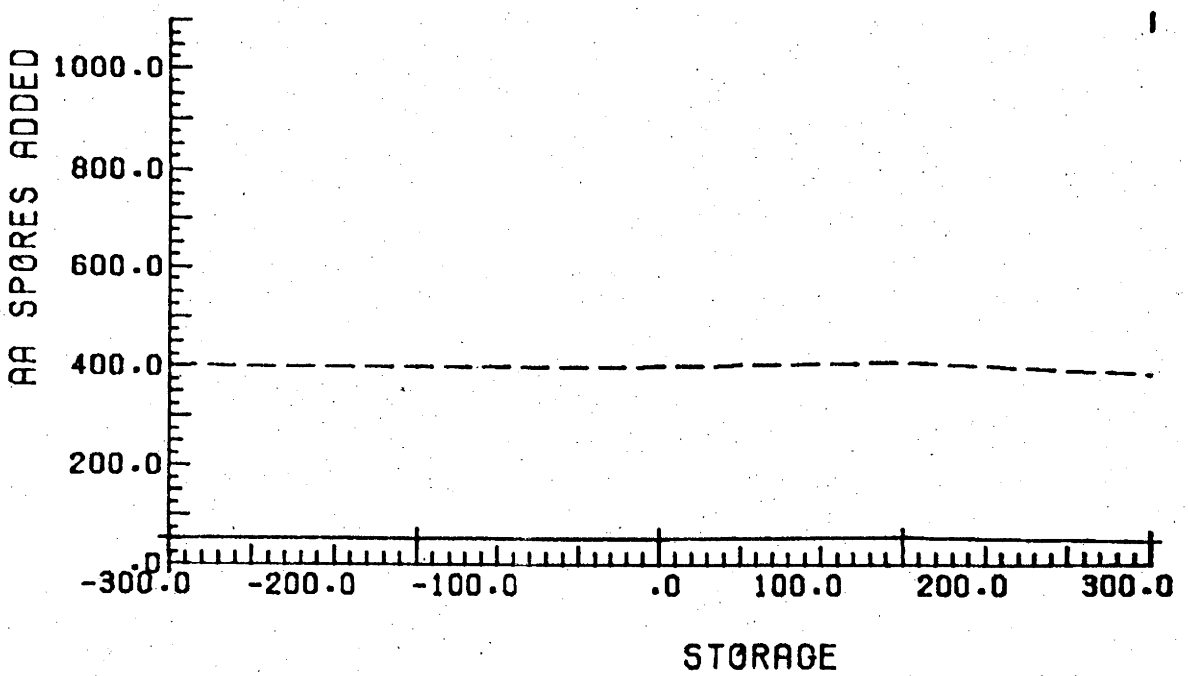

FIG. 6.5.43B: CROSSPROTECTION MULTILINE
AA ADDED VS AABB

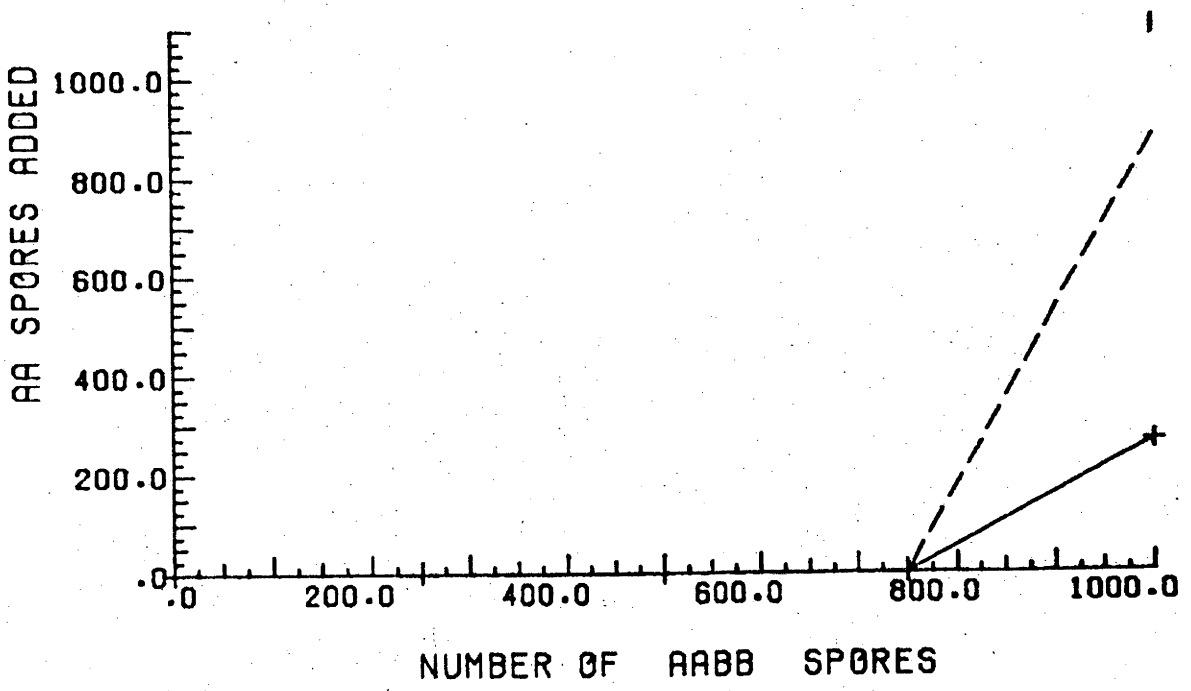

FIG. 6.5.44: REGRESSION OF CHANGE IN WORKRATE ON
SIMPLER LINES PLANTED

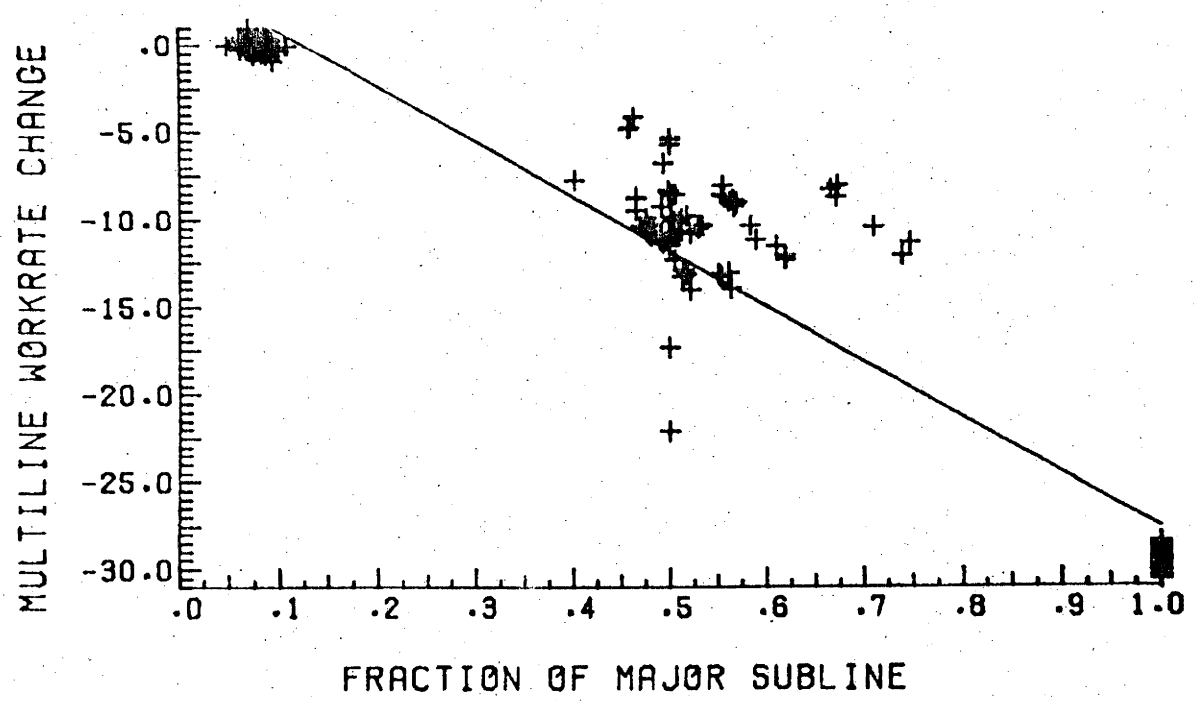

FIG. 6.5.45A: CROSSPROTECTION MULTILINE: SUPERLINE VS TIME

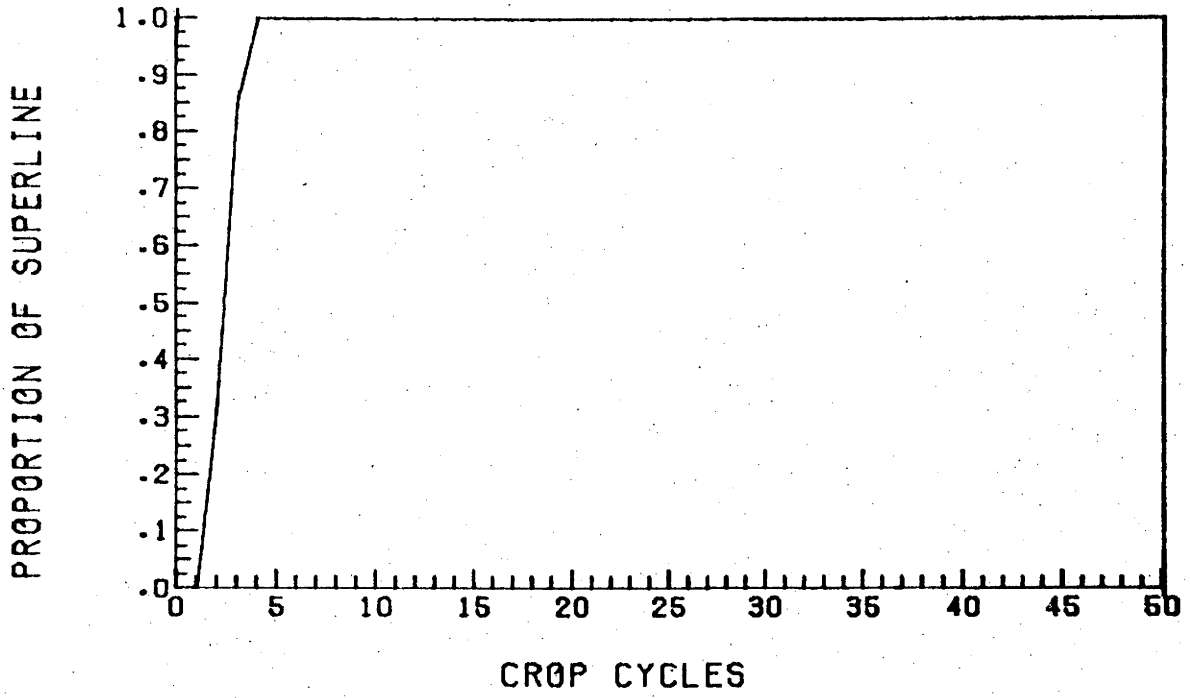

FIG. 6.5.45B: CROSSPROTECTION MULTILINE: COST THROUGH TIME

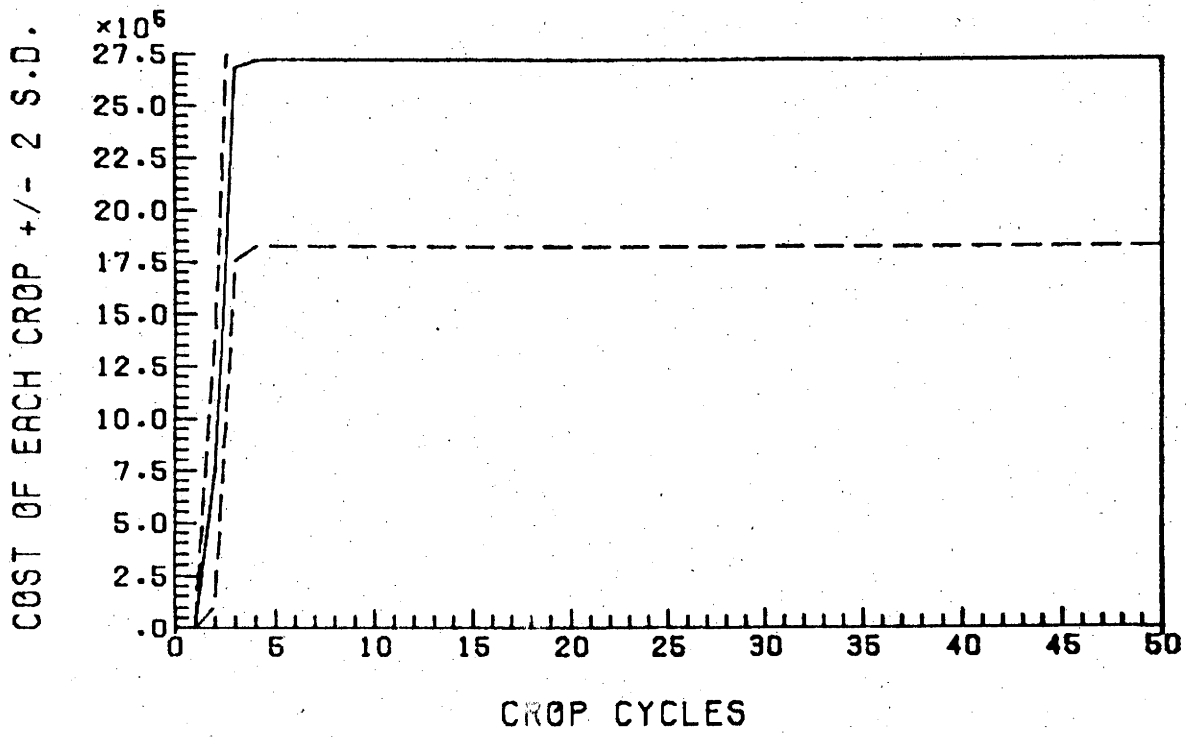

FIG. 6.5.45C: CROSSPROTECTION MULTILINE: YIELD THROUGH TIME

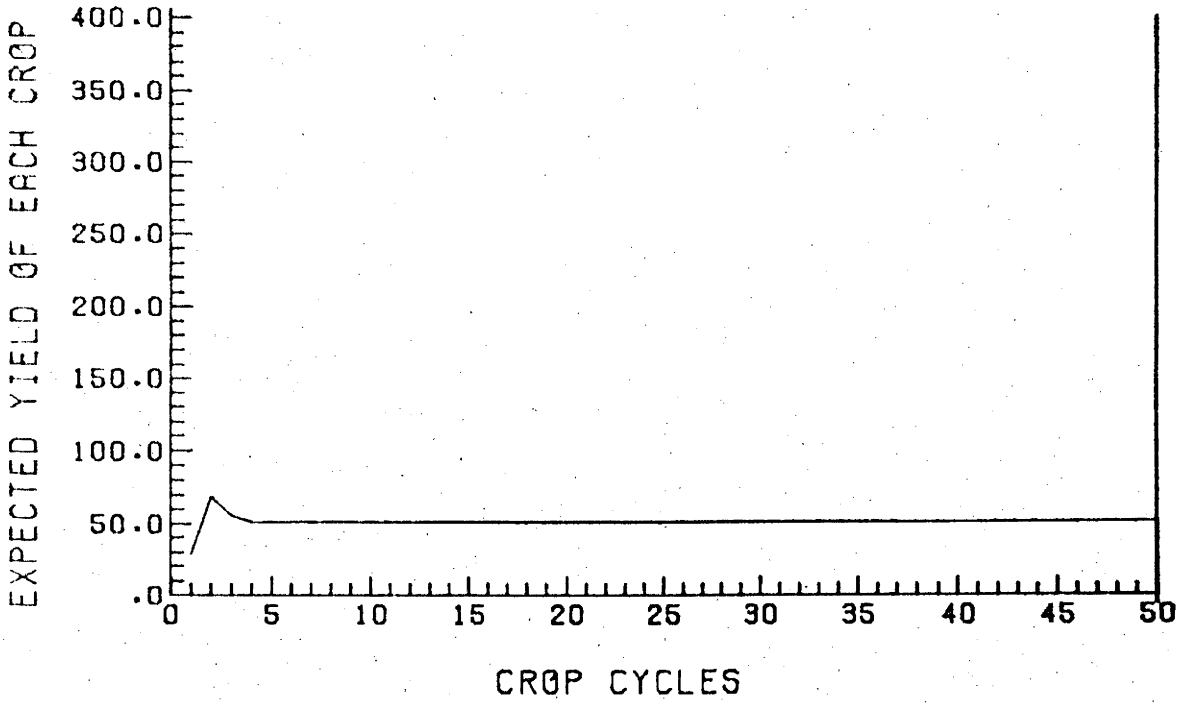

FIG. 6.5.45D: CROSSPROTECTION MULTILINE: STORAGE THROUGH TIME

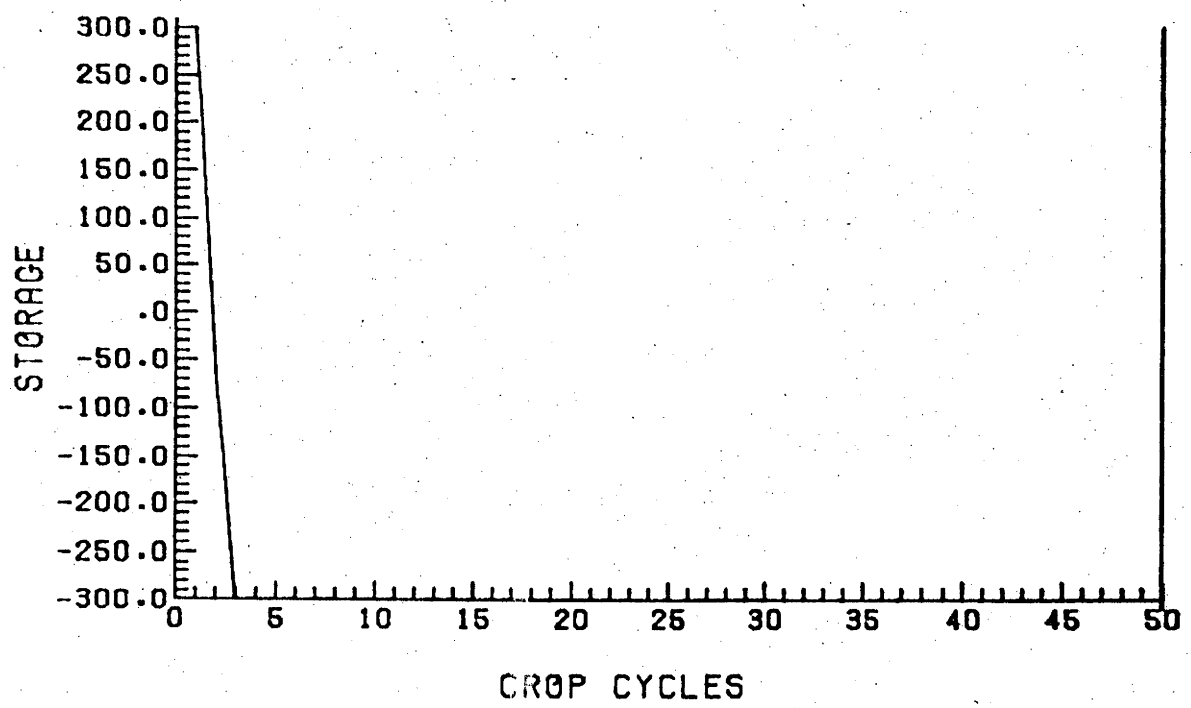

FIG. 6.5.45E: CROSSPROTECTION MULTILINE: SUPERRACE VS TIME

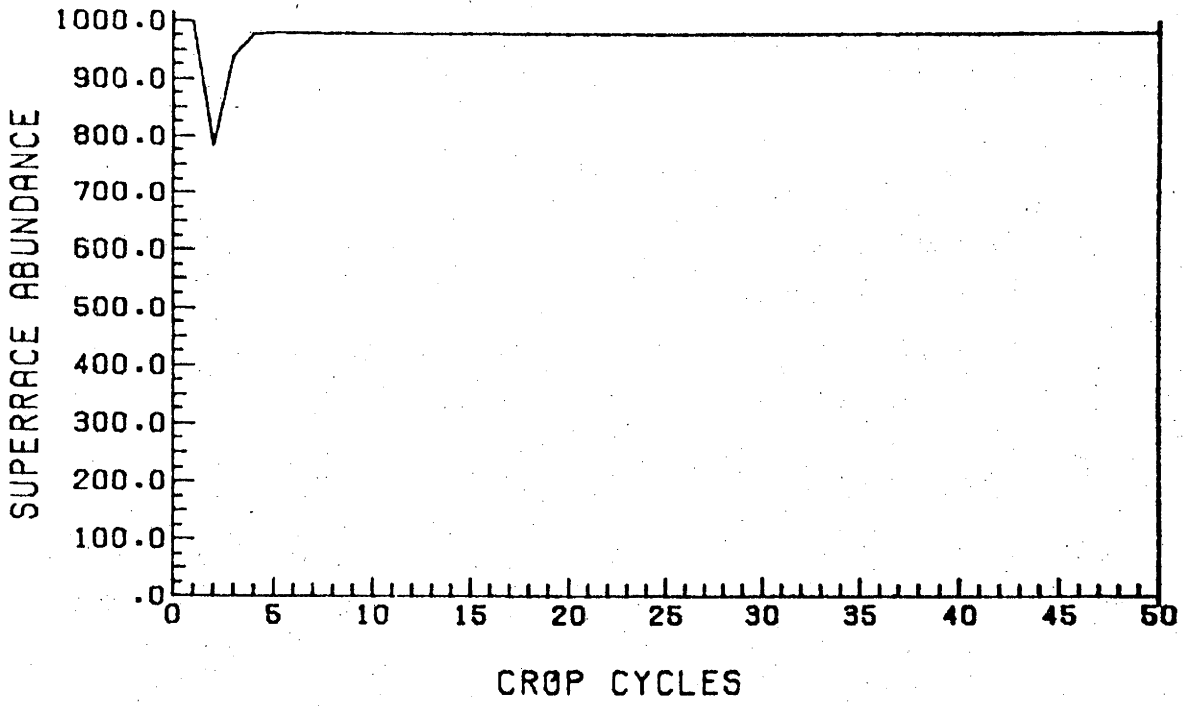

FIG. 6.5.45F: CROSSPROTECTION MULTILINE: WORK THROUGH TIME

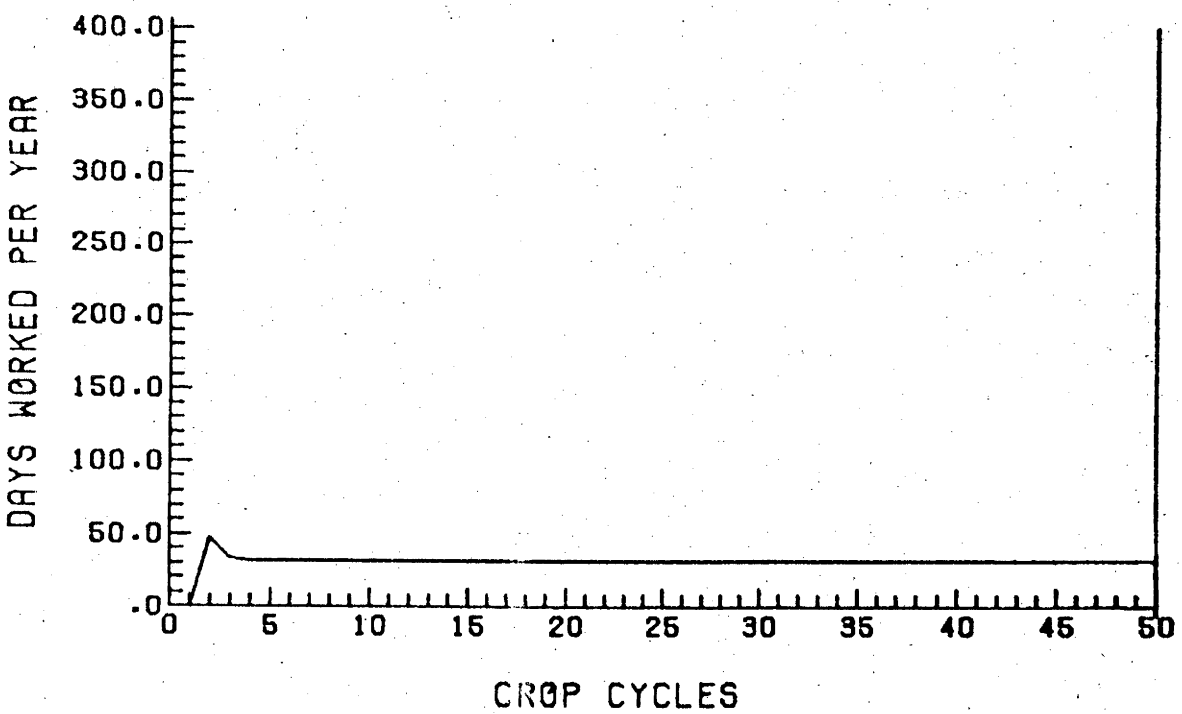

FIG. 6.5.46A: SUBSISTENCE + STABILISING SELECTION
DAYS WORKED VS STORAGE

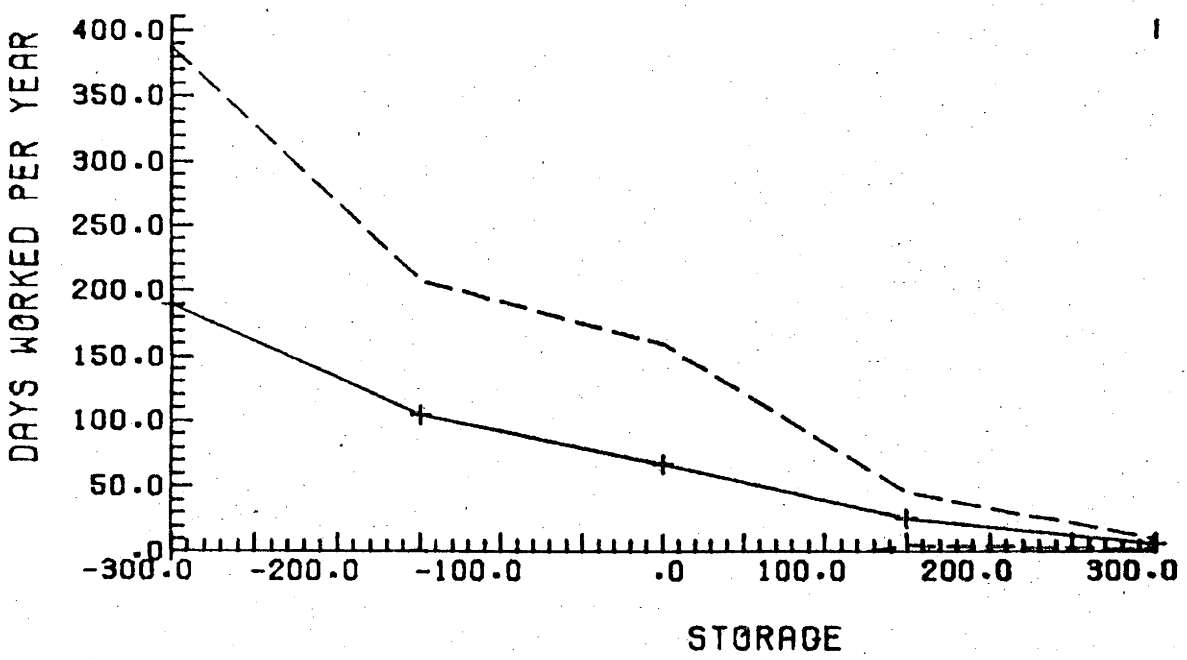

FIG. 6.5.46B: SUBSISTENCE + STABILISING SELECTION
DAYS WORKED VS AABB

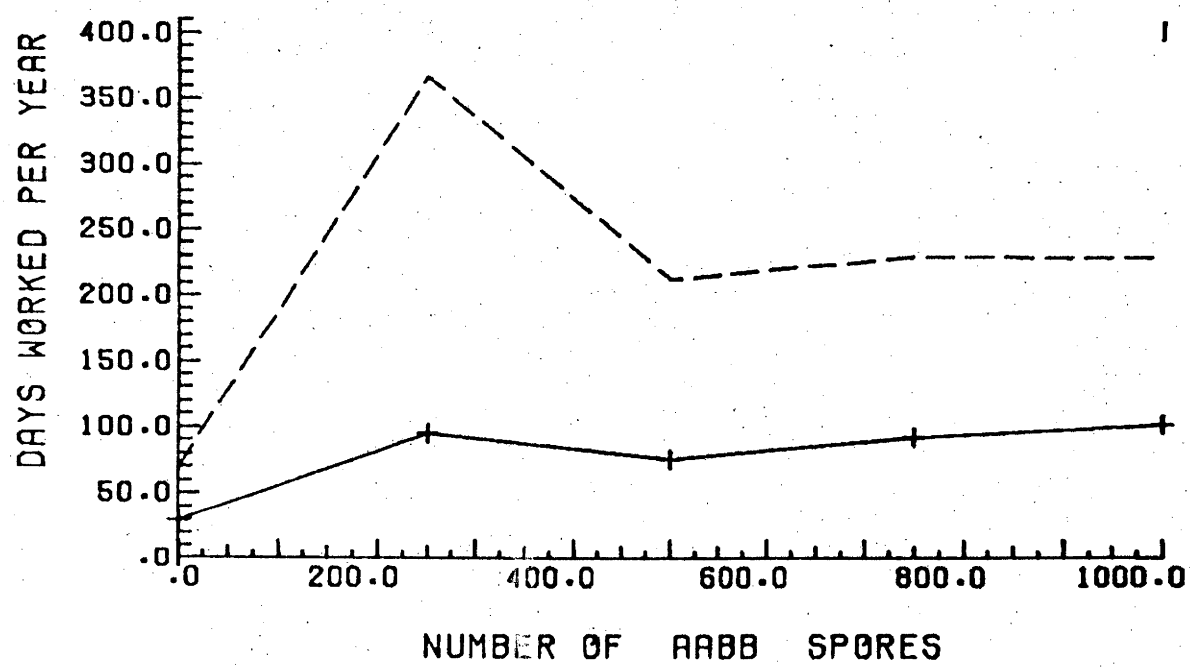

FIG. 6.5.47A: SUBSISTENCE + STABILISING SELECTION
E(COST) VS STORAGE

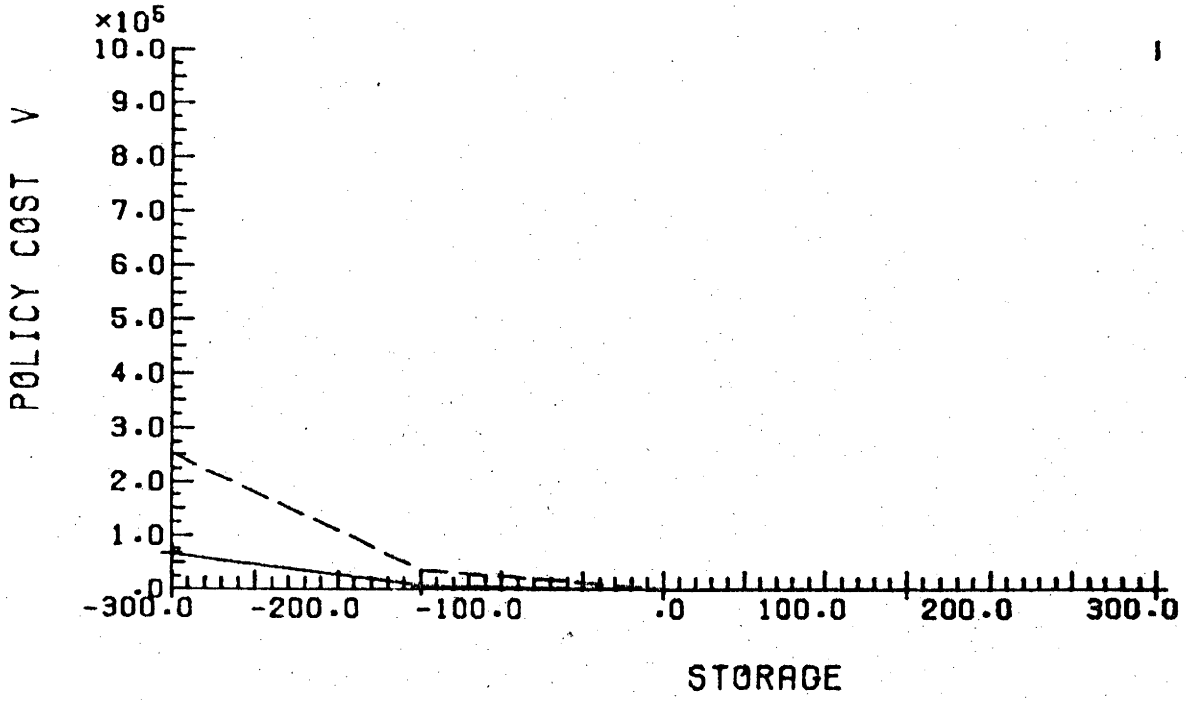

FIG. 6.5.47B: SUBSISTENCE + STABILISING SELECTION
E(COST) VS AABB

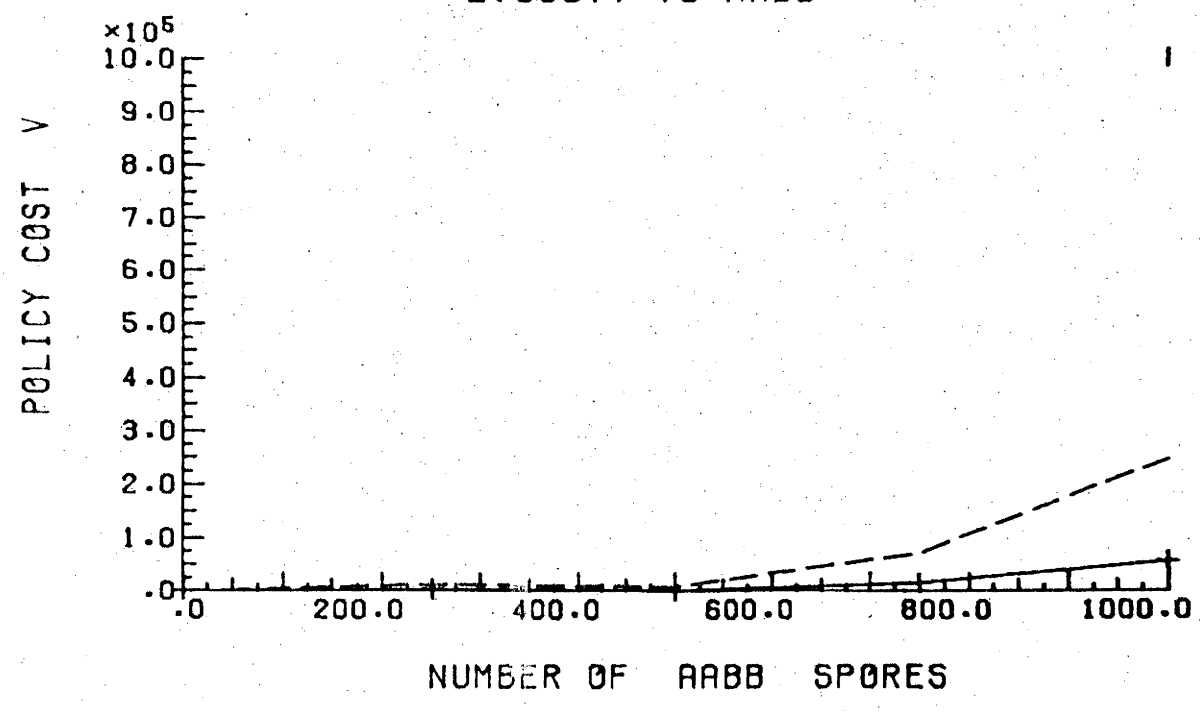

FIG. 6.5.48A: SUBSISTENCE + STABILISING SELECTION
STORAGE INPUT/OUTPUT

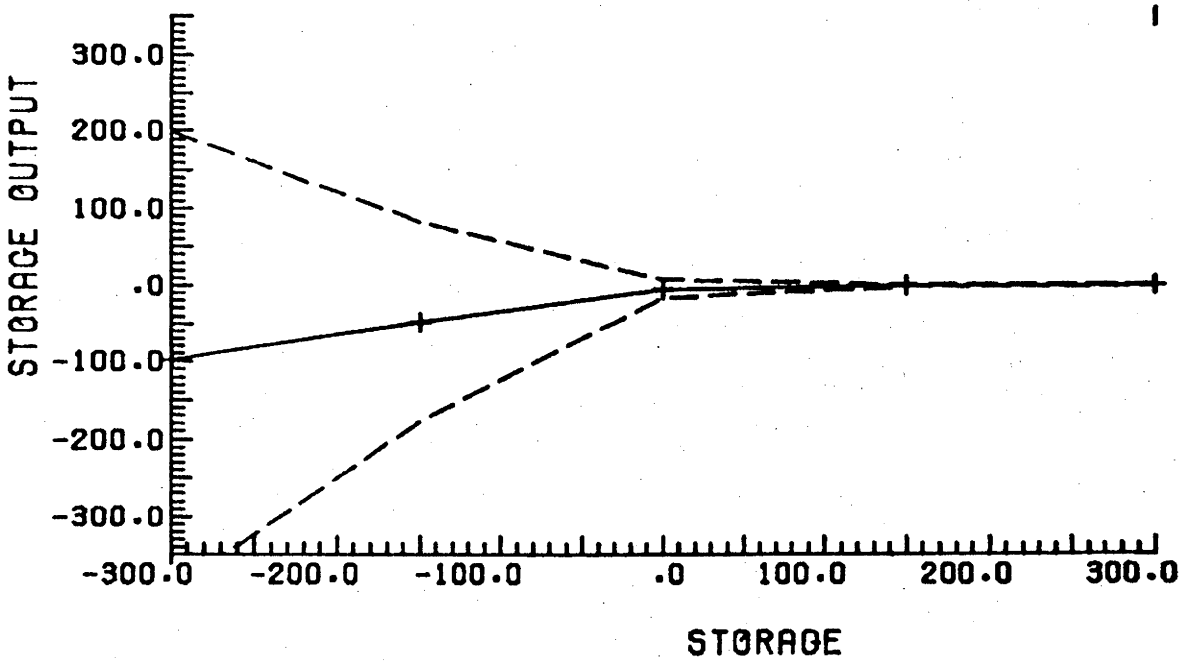

FIG. 6.5.48B: SUBSISTENCE + STABILISING SELECTION
STORAGE LEFT VS AABB

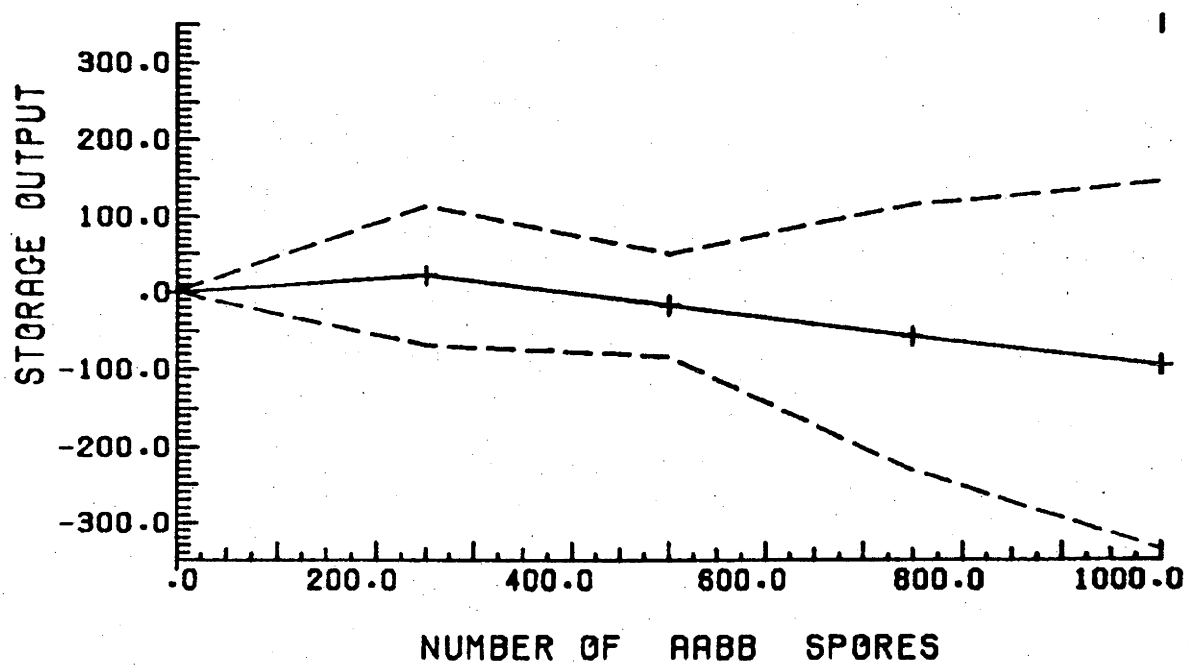

FIG. 6.5.49: SUBSISTENCE + STABILISING SELECTION
SUPERRACE INPUT/OUTPUT

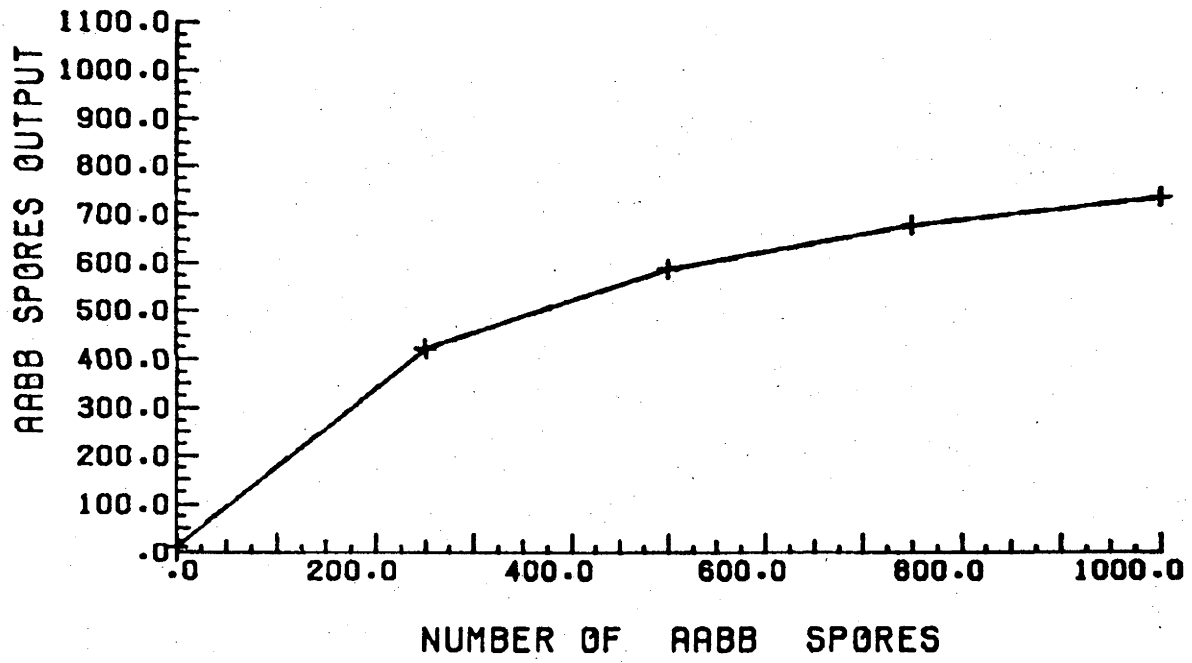

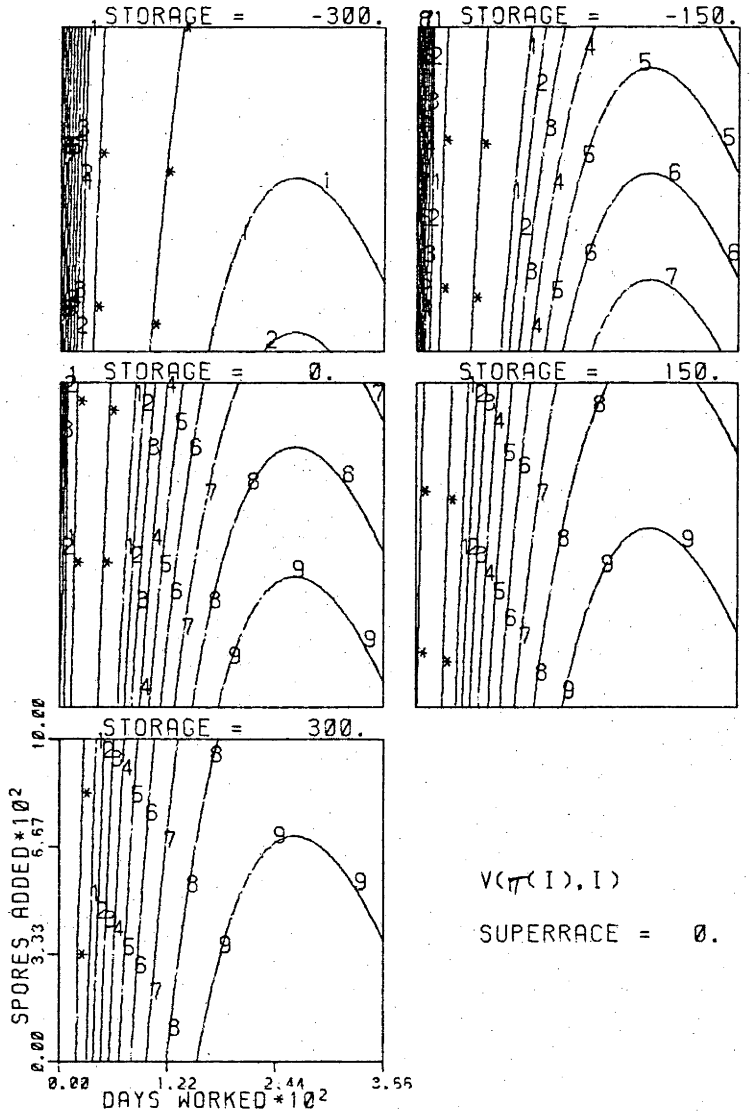

FIG. 6.5.50· SUBSISTENCE + STABILISING SELECTION

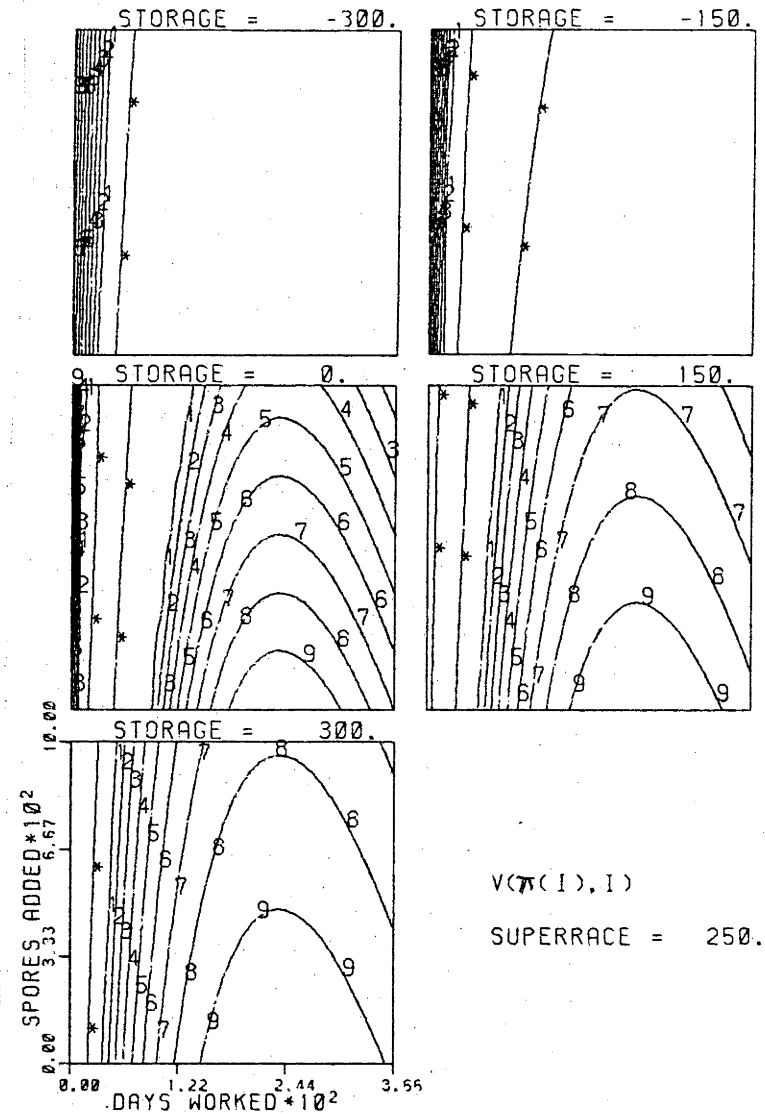

FIG. 6.5.51· SUBSISTENCE + STABILISING SELECTION

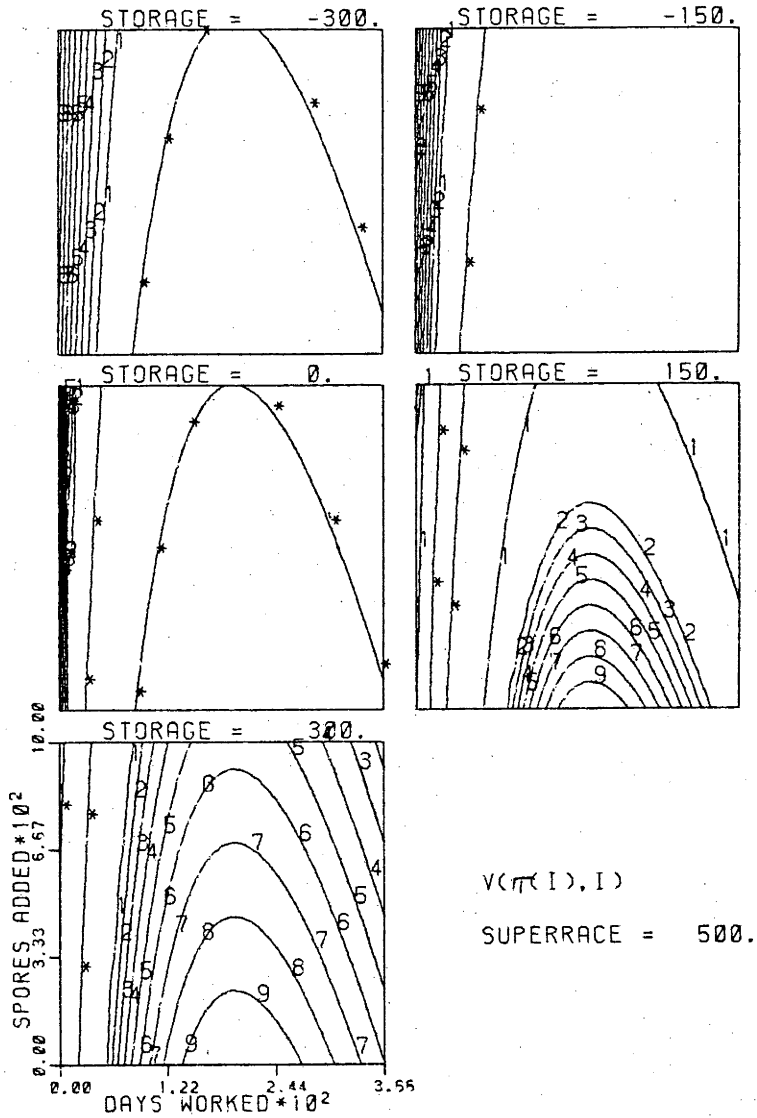

FIG. 6.5.52: SUBSISTENCE + STABILISING SELECTION

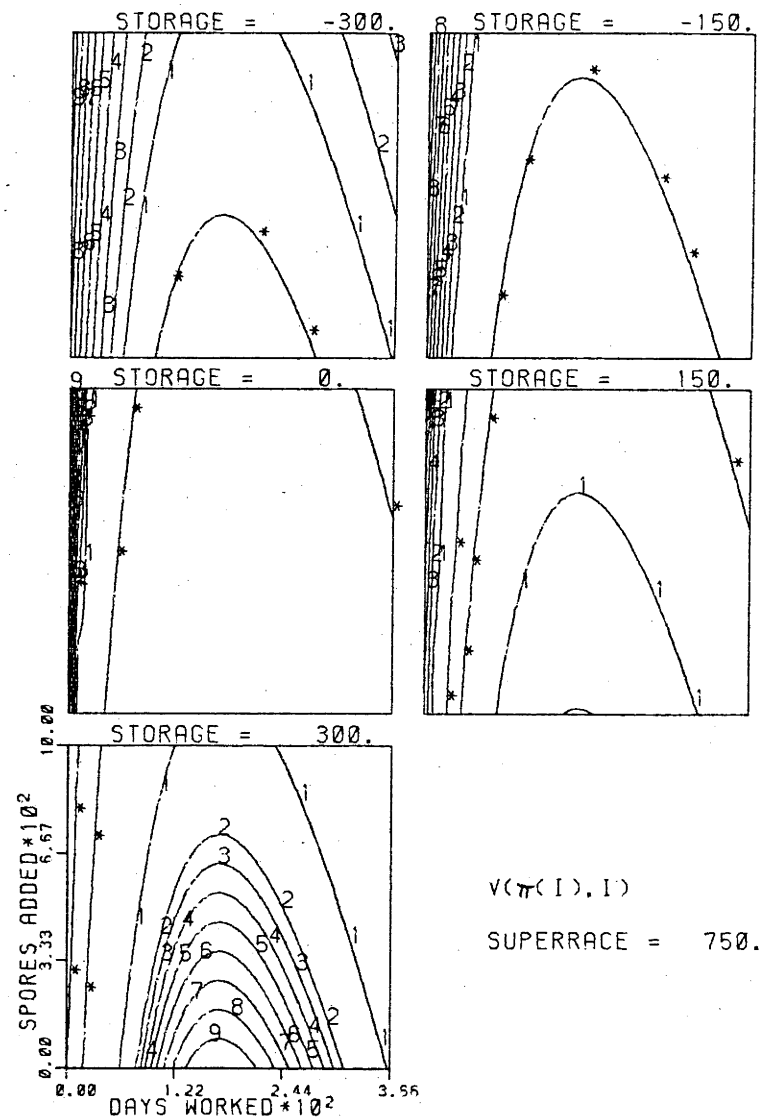

FIG. 6.5.53: SUBSISTENCE + STABILISING SELECTION

FIG. 6.5.55: SUBSISTENCE + STABILISING SELECTION: WORK THROUGH TIME

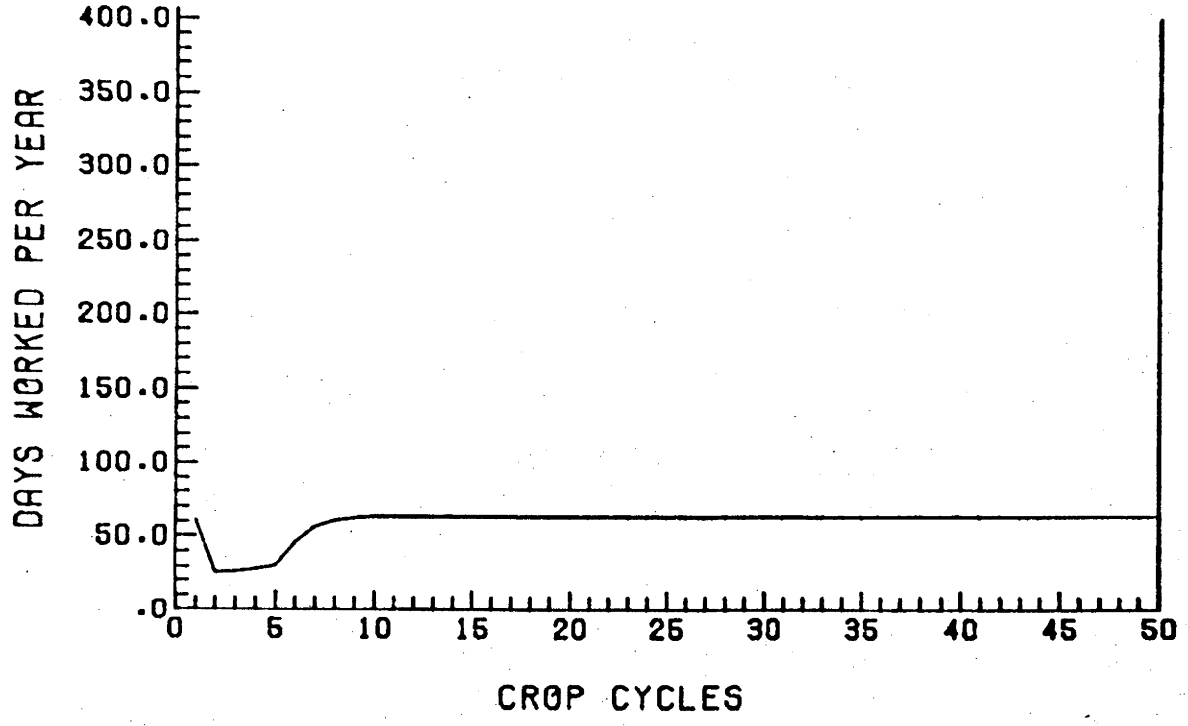

FIG. 6.5.56: SUBSISTENCE + STABILISING SELECTION: COST THROUGH TIME

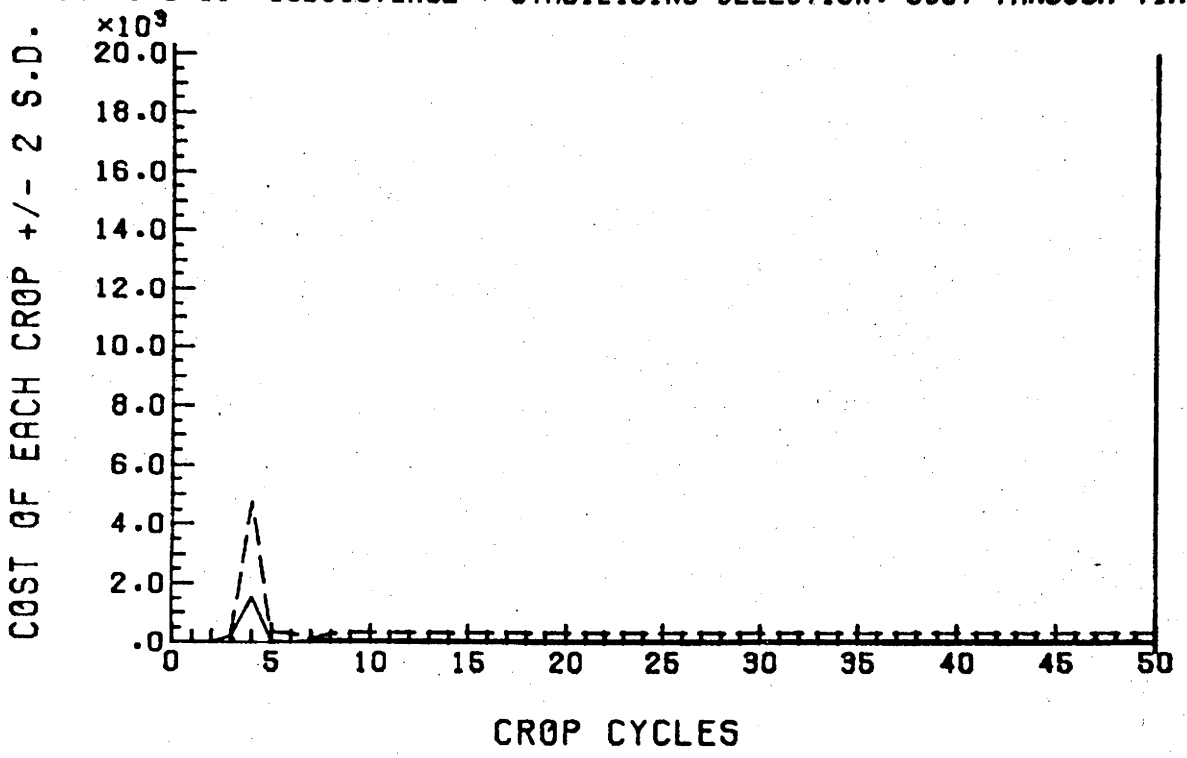

FIG. 6.5.57: SUBSISTENCE + STABILISING SELECTION: YIELD THROUGH TIME

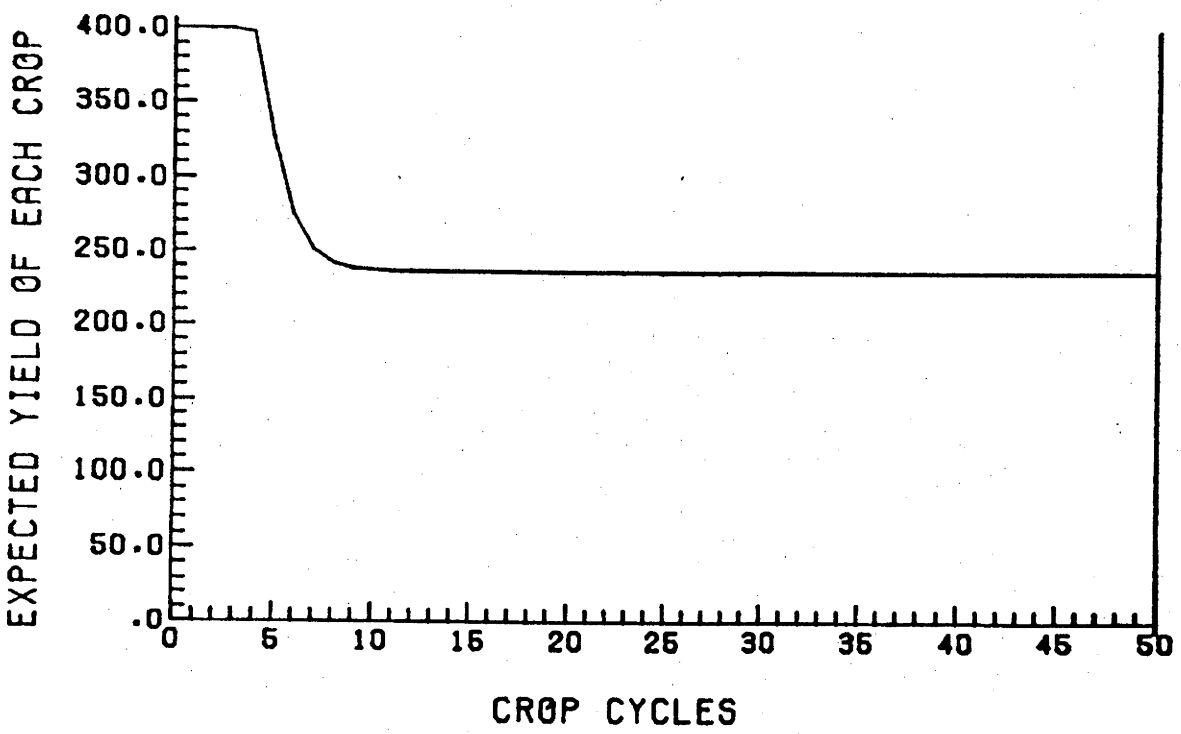

FIG. 6.4.58: SUBSISTENCE + STABILISING SELECTION: STORAGE THROUGH TIME

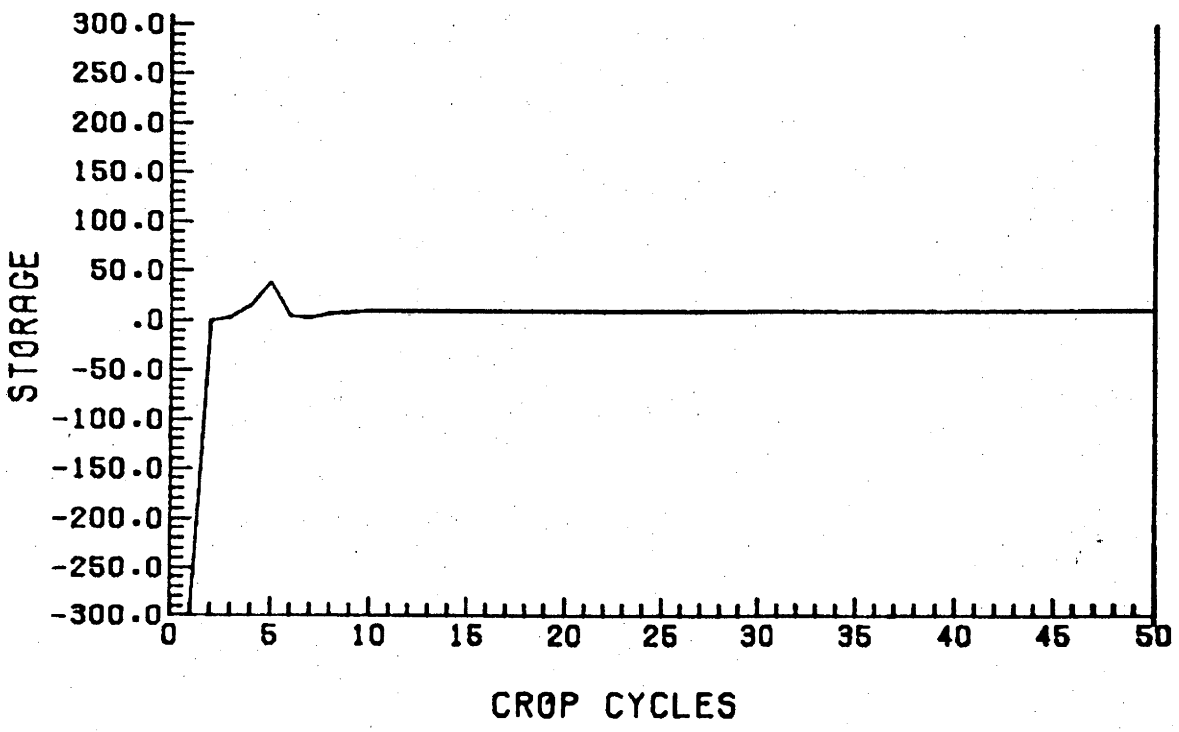

FIG. 6.4.59: SUBSISTENCE + STABILISING SELECTION: SUPERRACE THROUGH TIME

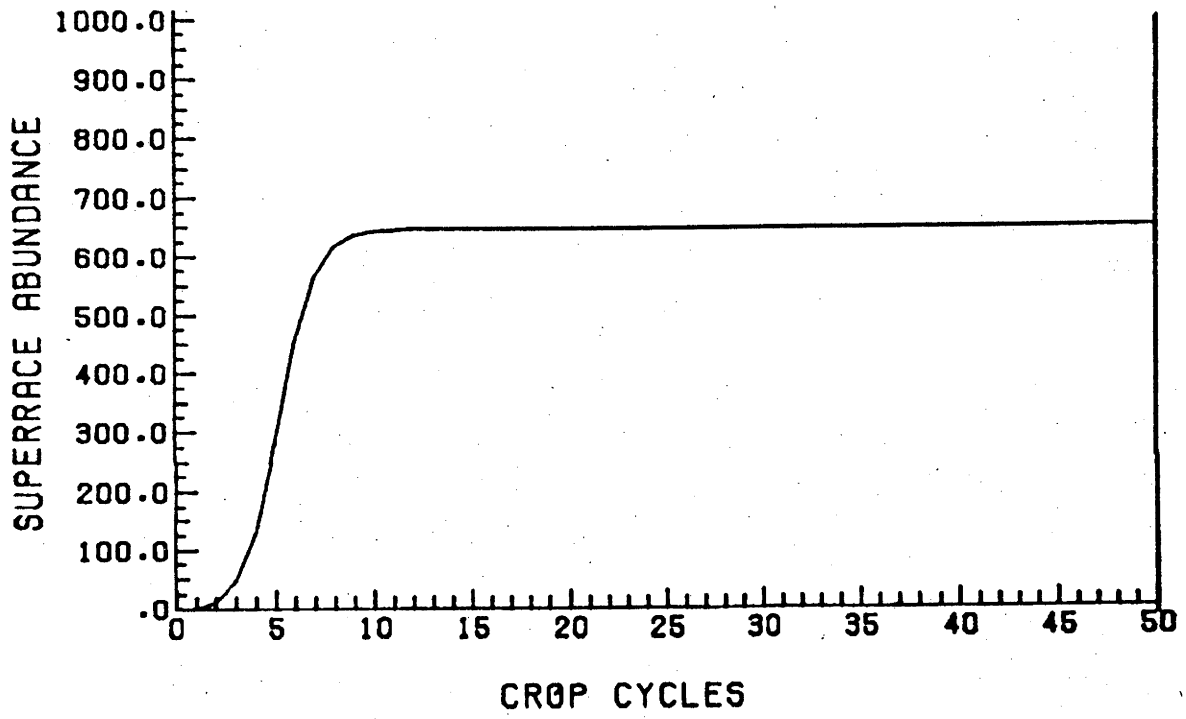

FIG. 6.5.60A: STABILISING SELECTION MULTILINE F(A) USED VS STORAGE

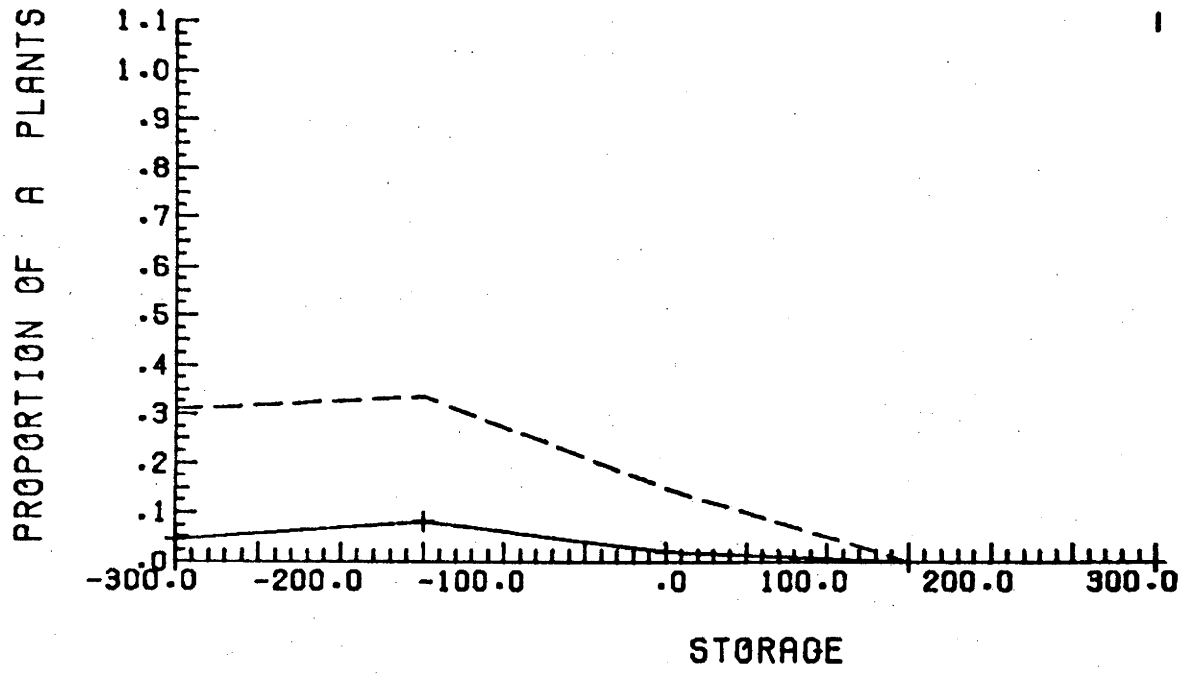

FIG. 6.5.60B: STABILISING SELECTION MULTILINE F(A) USED VS AABB

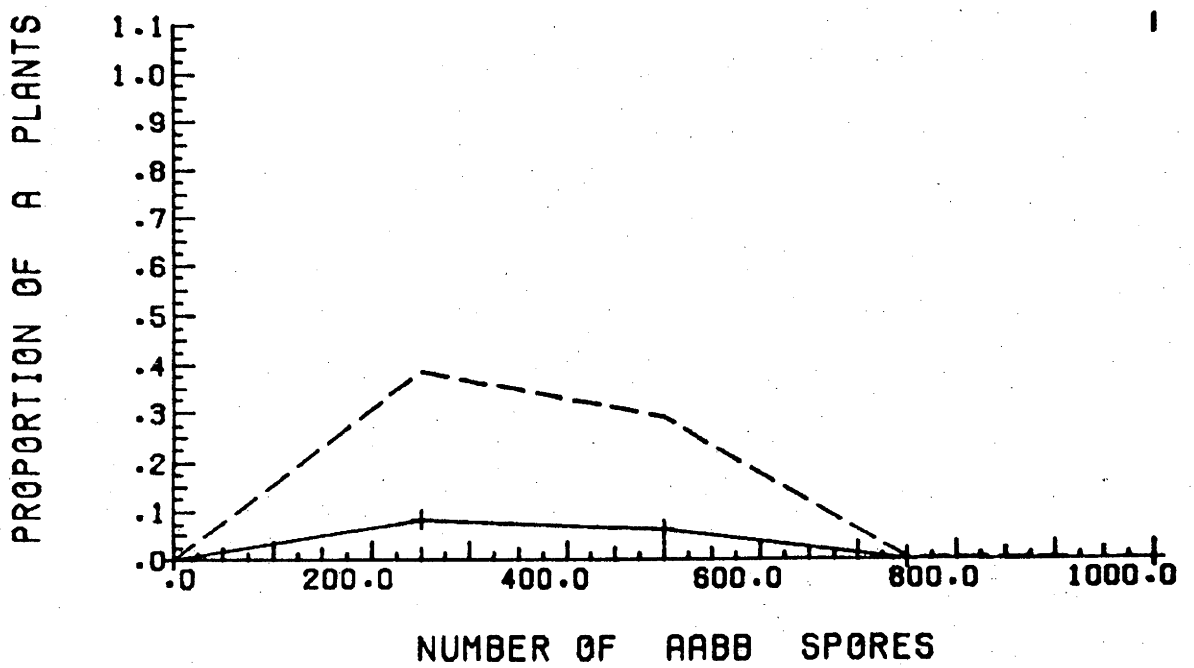

FIG. 6.5.63: STABILISING SELECTION MULTILINE
WORK DONE VS AABB

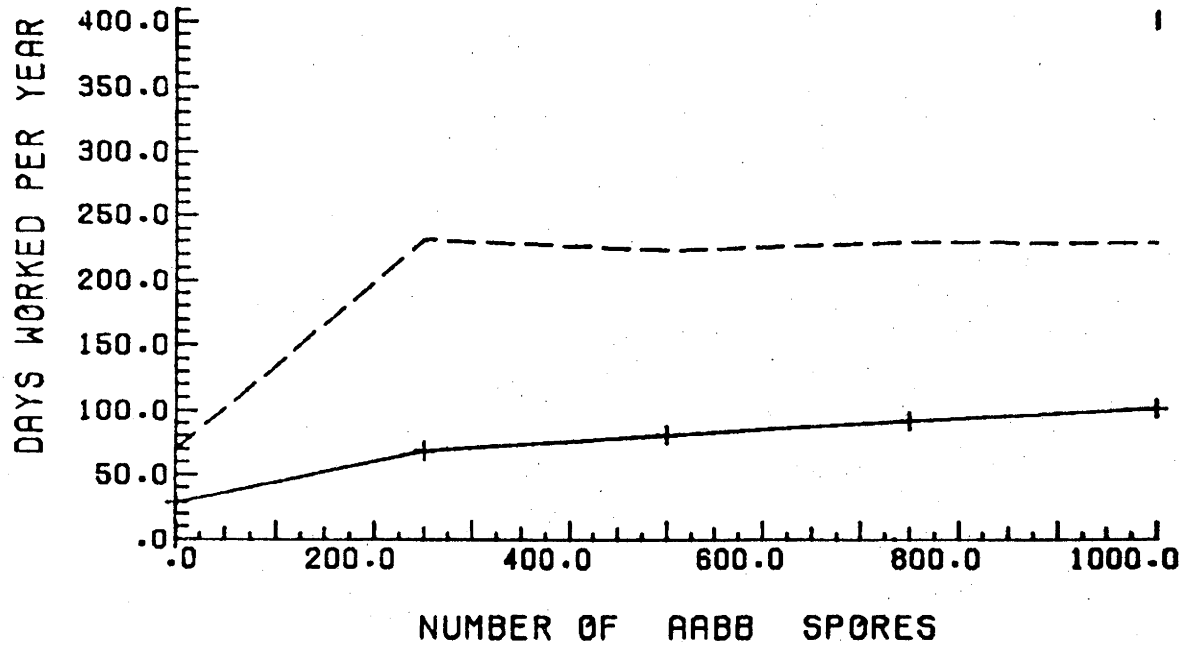

FIG. 6.5.64A: STABILISING SELECTION MULTILINE
AA SPORES USED VS STORAGE

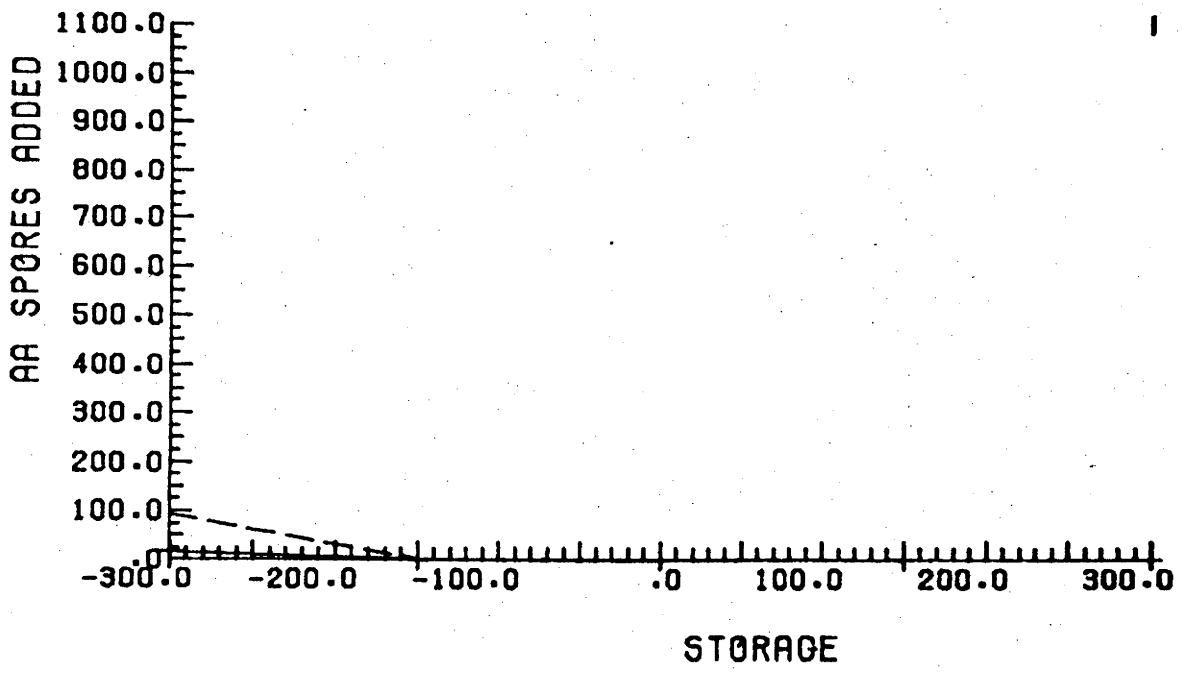

FIG. 6.5.64B: STABILISING SELECTION MULTILINE
AA ADDED VS AABB

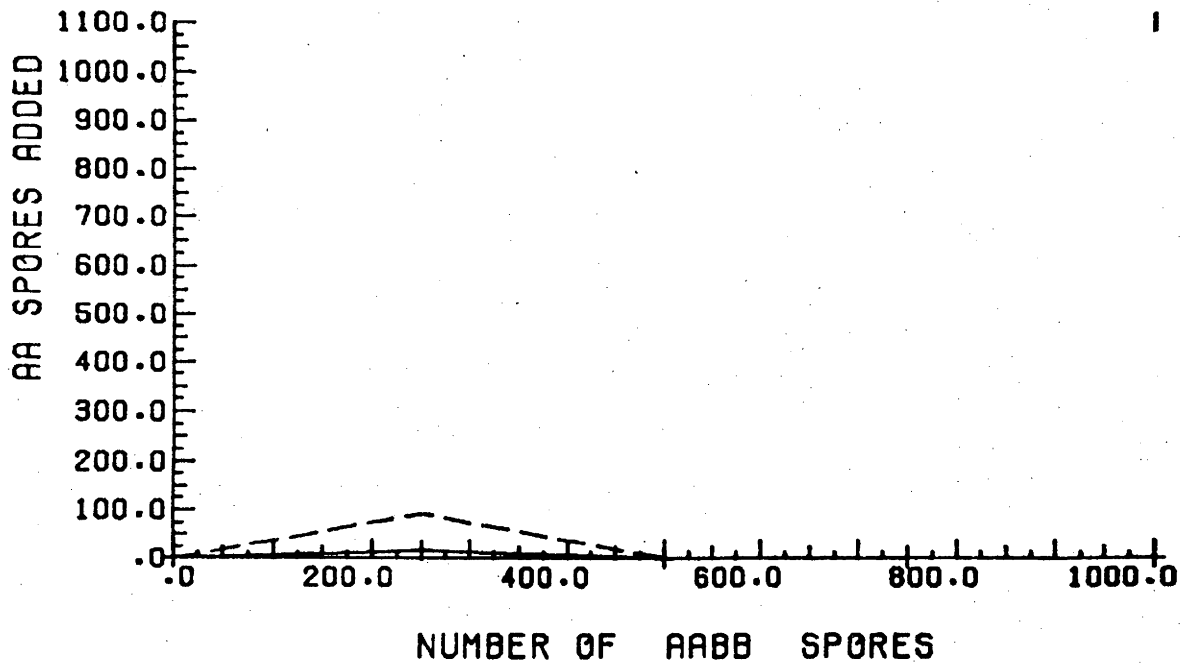

FIG. 6.5.65: REGRESSION OF SUBLINE USED ON
SIMPLER RACE PRESENT

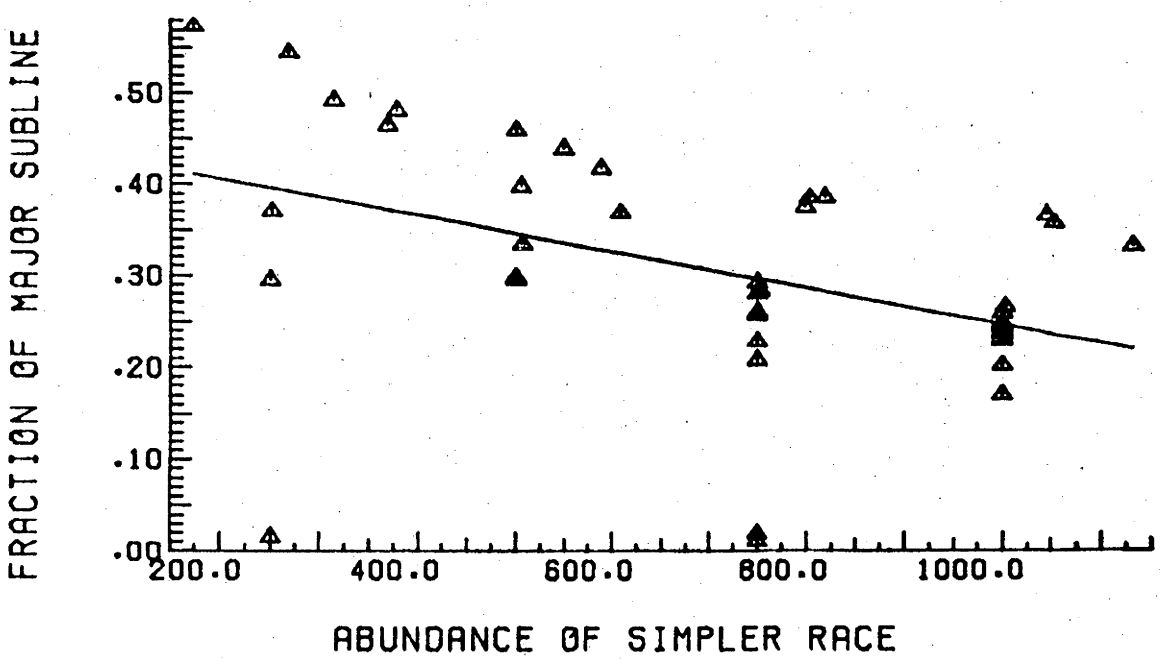

FIG. 6.5.66A: STABILISING SELECTION MULTILINE SUPERLINE VS TIME

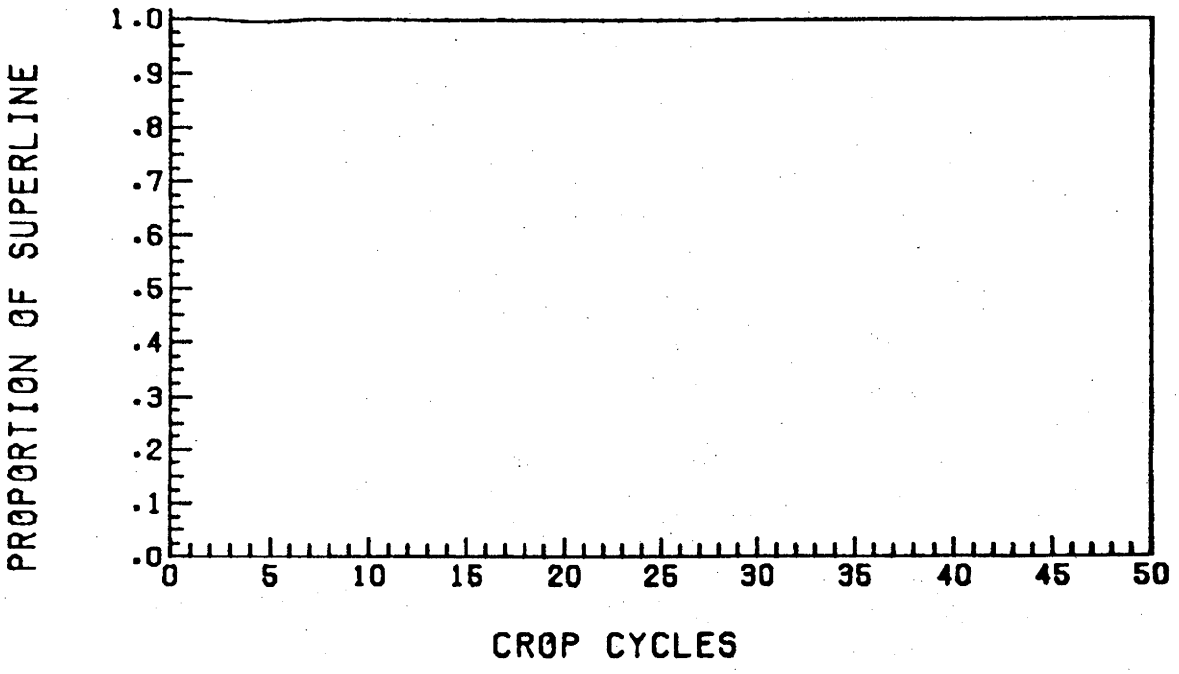

FIG. 6.5.66B: STABILISING SELECTION MULTILINE SUPERRACE VS TIME

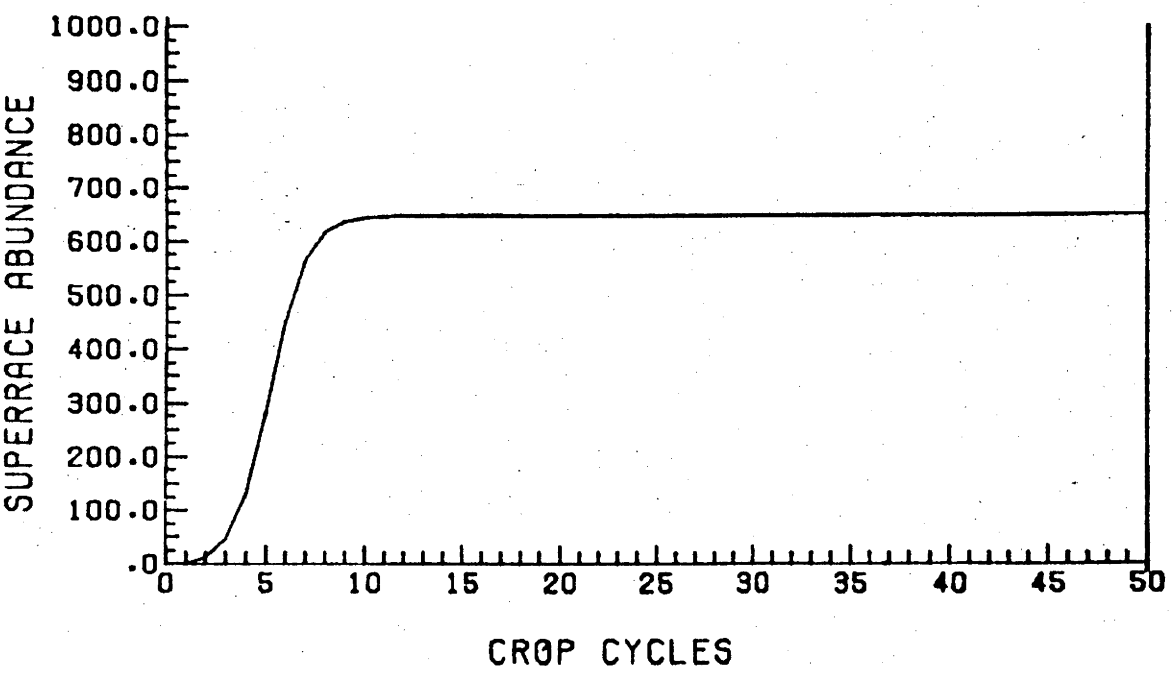

FIG. 6.5.67A: SUBSISTENCE + GROWTH
DAYS WORKED VS STORAGE

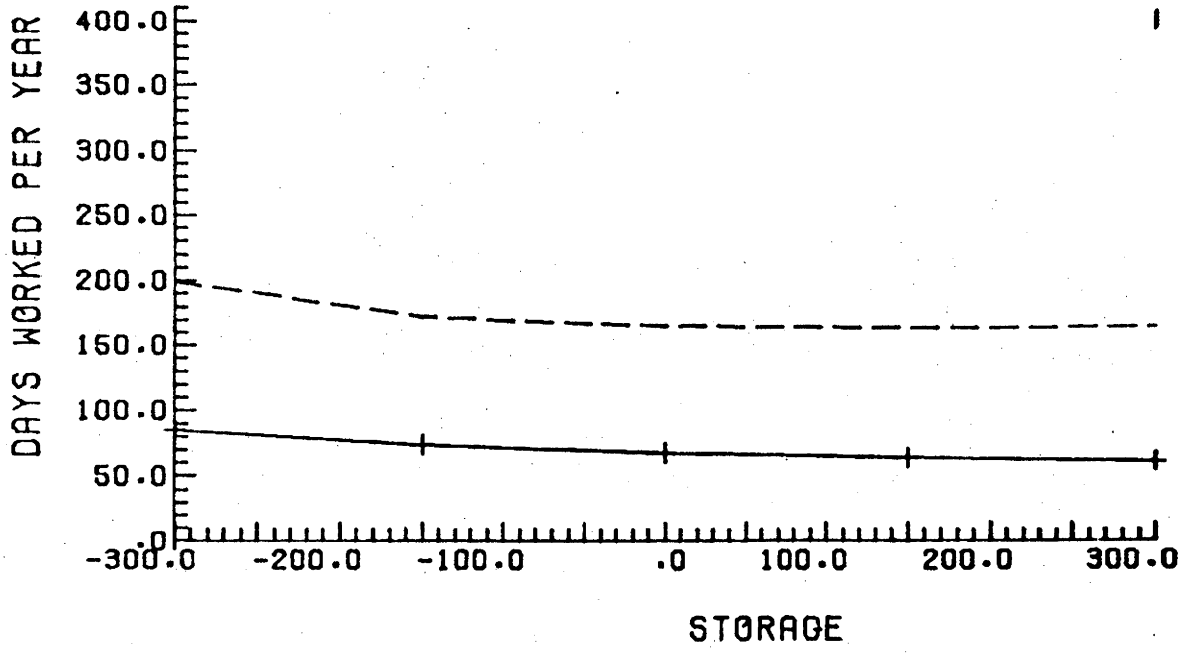

FIG. 6.5.67B: SUBSISTENCE + GROWTH
DAYS WORKED VS AABB

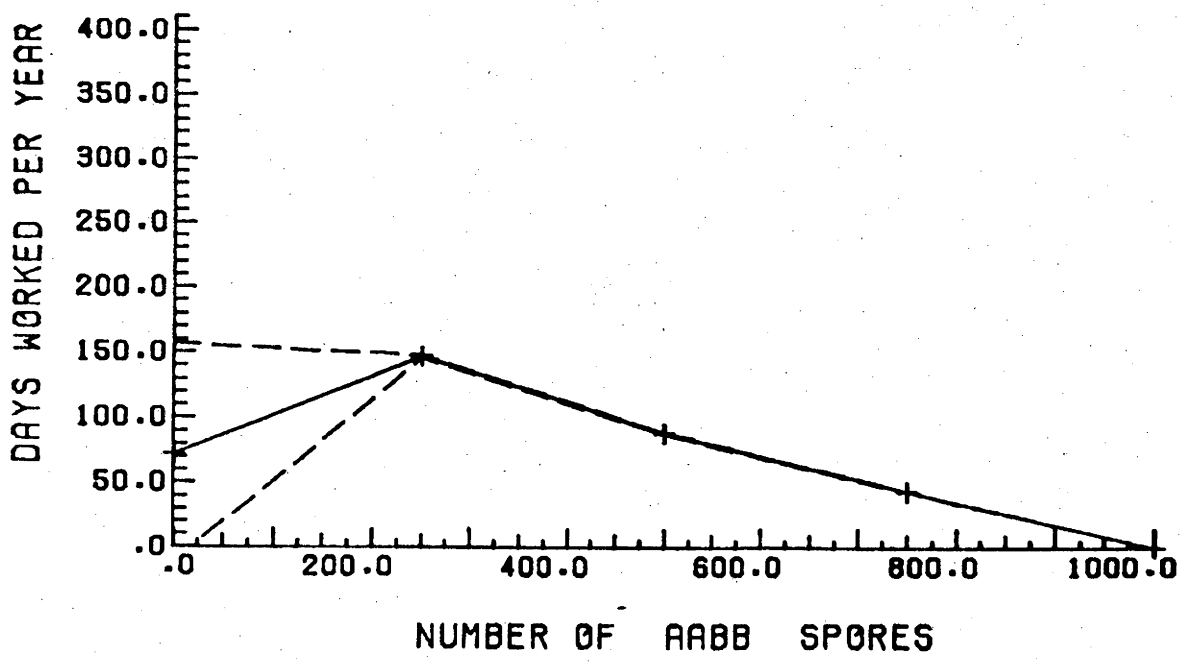

FIG. 6.5.68A: SUBSISTENCE + GROWTH
E(COST) VS STORAGE

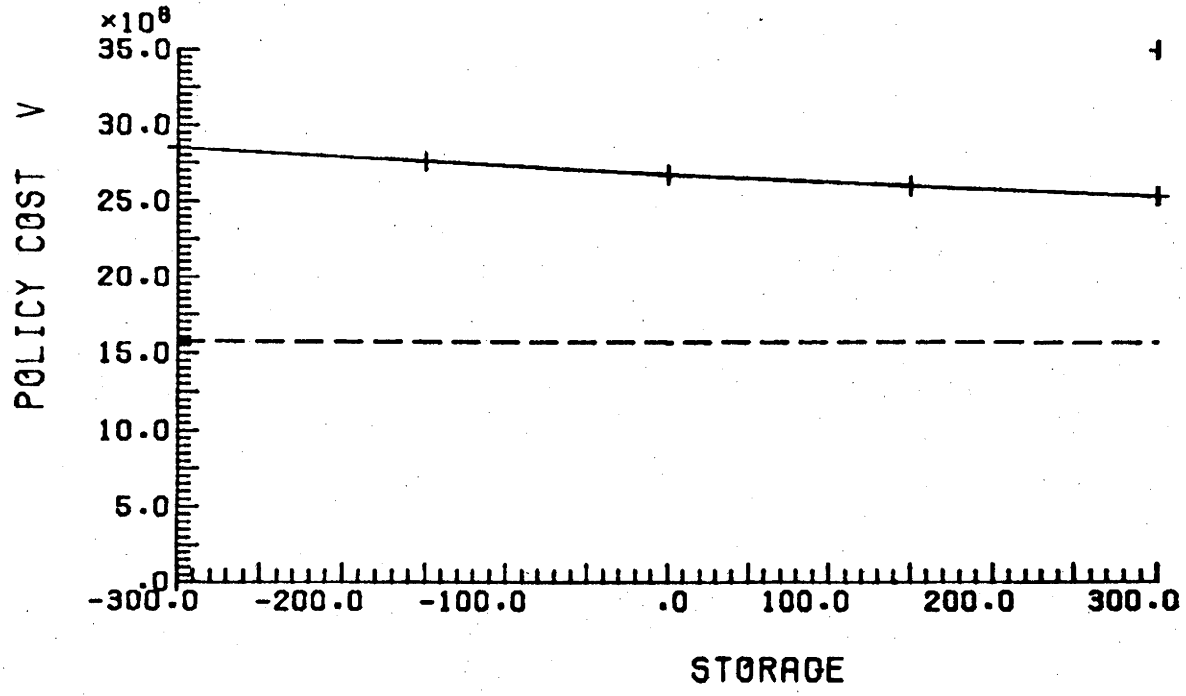

FIG. 6.5.68B: SUBSISTENCE + GROWTH
E(COST) VS AABB

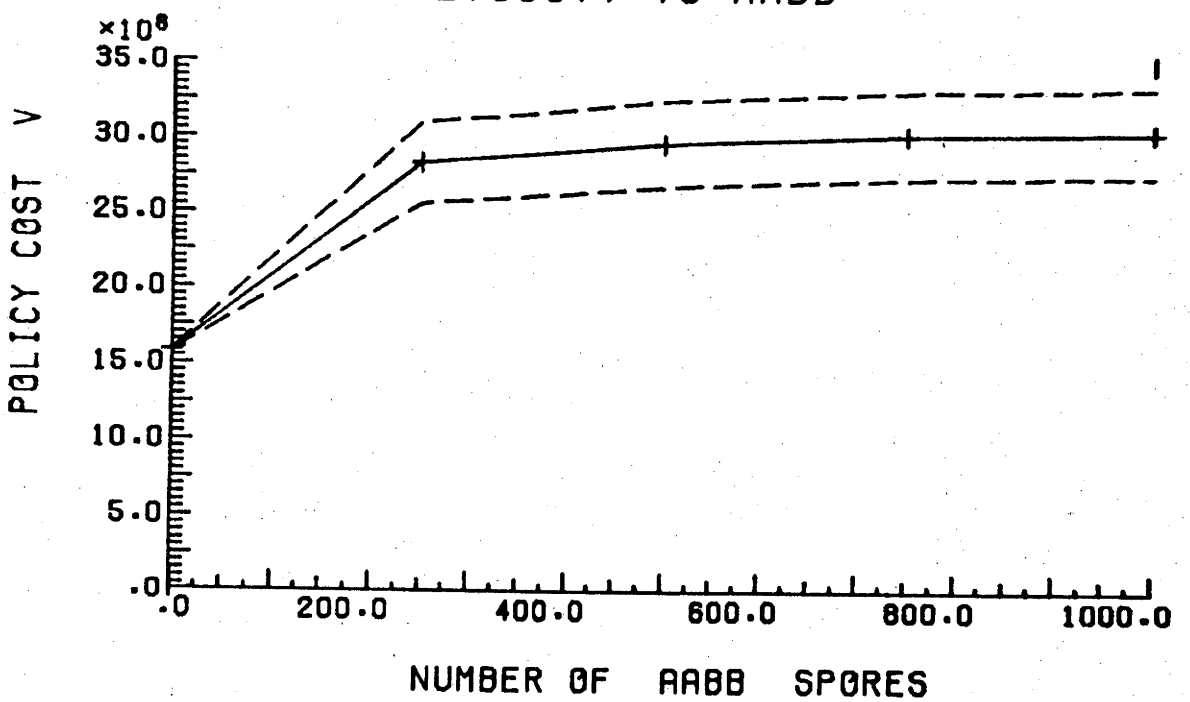

FIG. 6.5.69A: SUBSISTENCE + GROWTH
STORAGE INPUT/OUTPUT

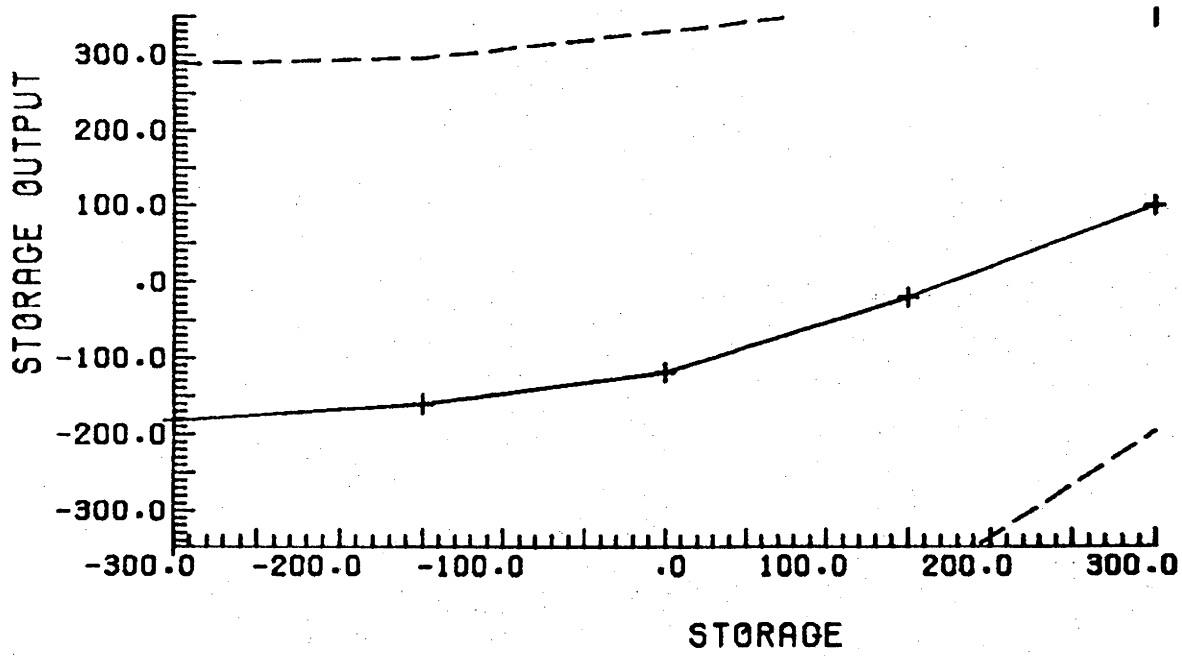

FIG. 6.5.69B: SUBSISTENCE + GROWTH
STORAGE OUTPUT VS AABB

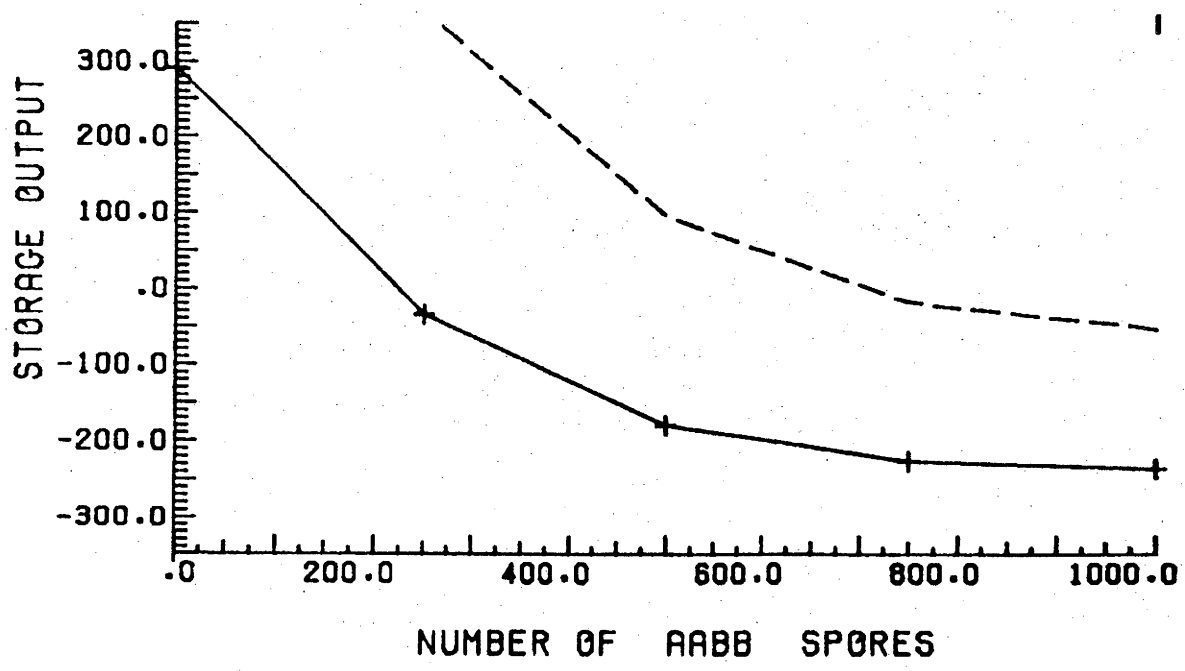

FIG. 6.5.70: SUBSISTENCE + GROWTH
SUPERRACE INPUT/OUTPUT

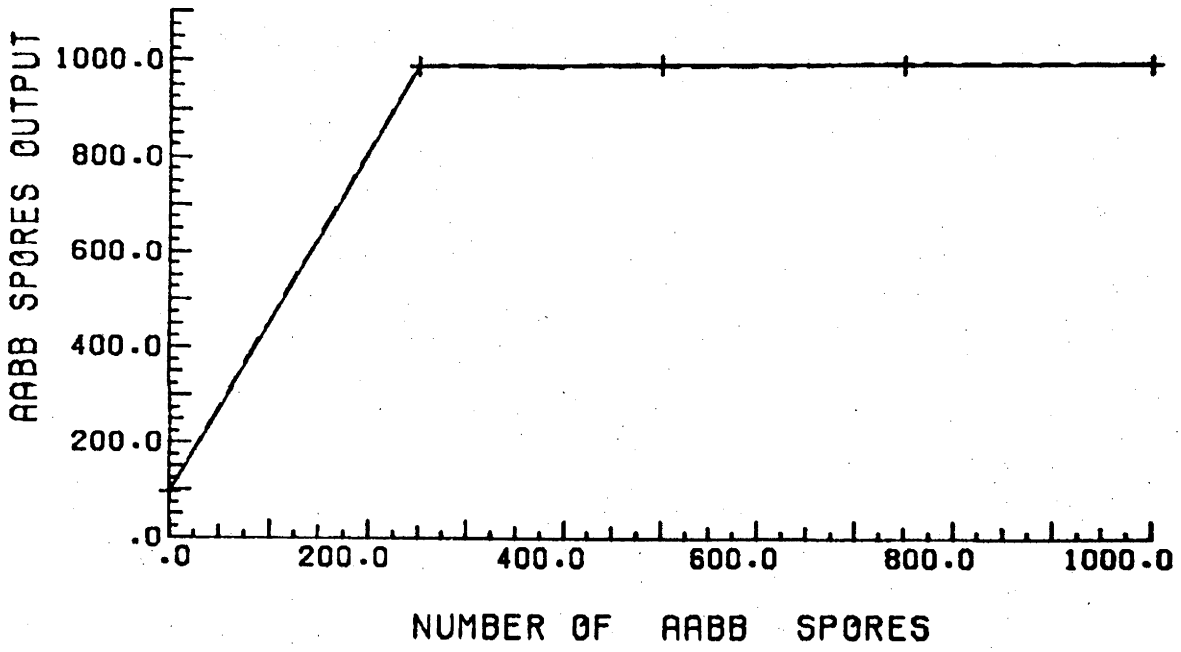

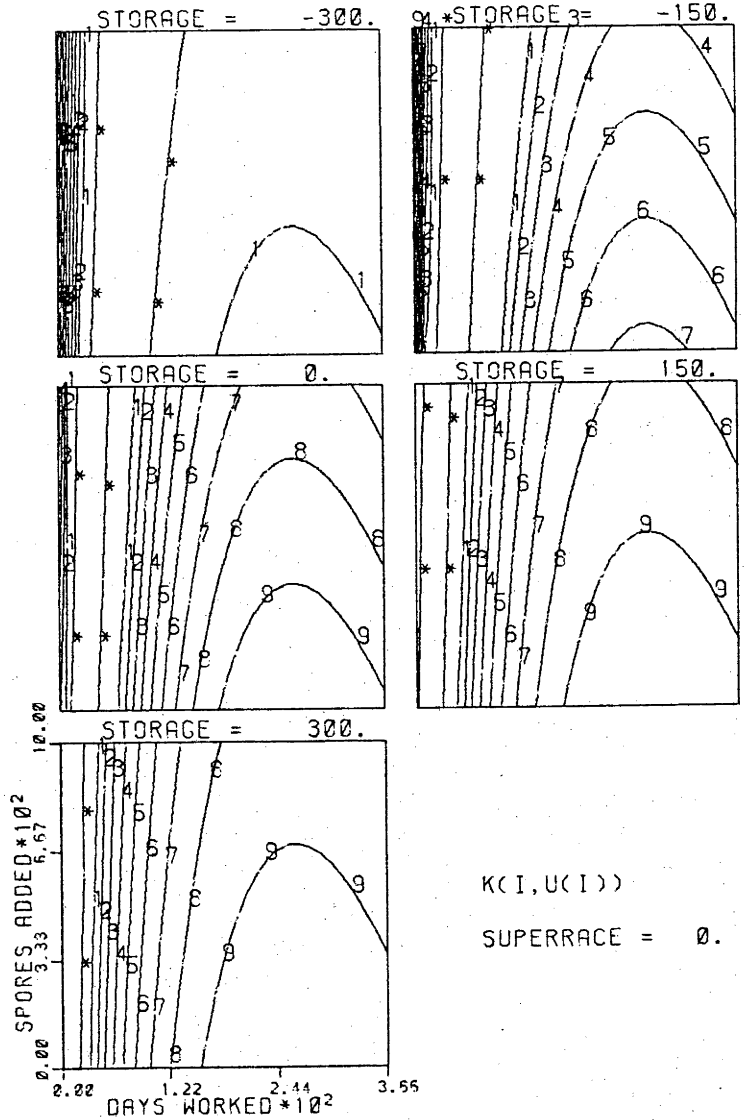

FIG. 6.5.71A: SUBSISTENCE WITH GROWING PLANT

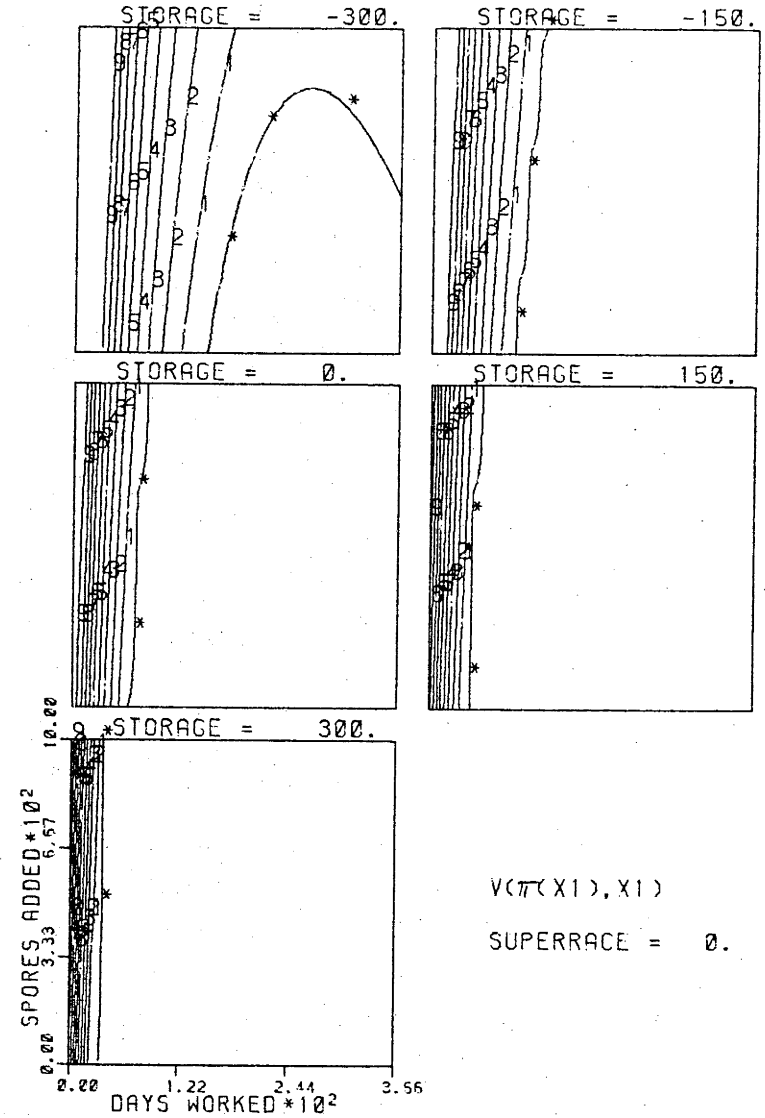

FIG. 6.5.71B: SUBSISTENCE WITH GROWING PLANT

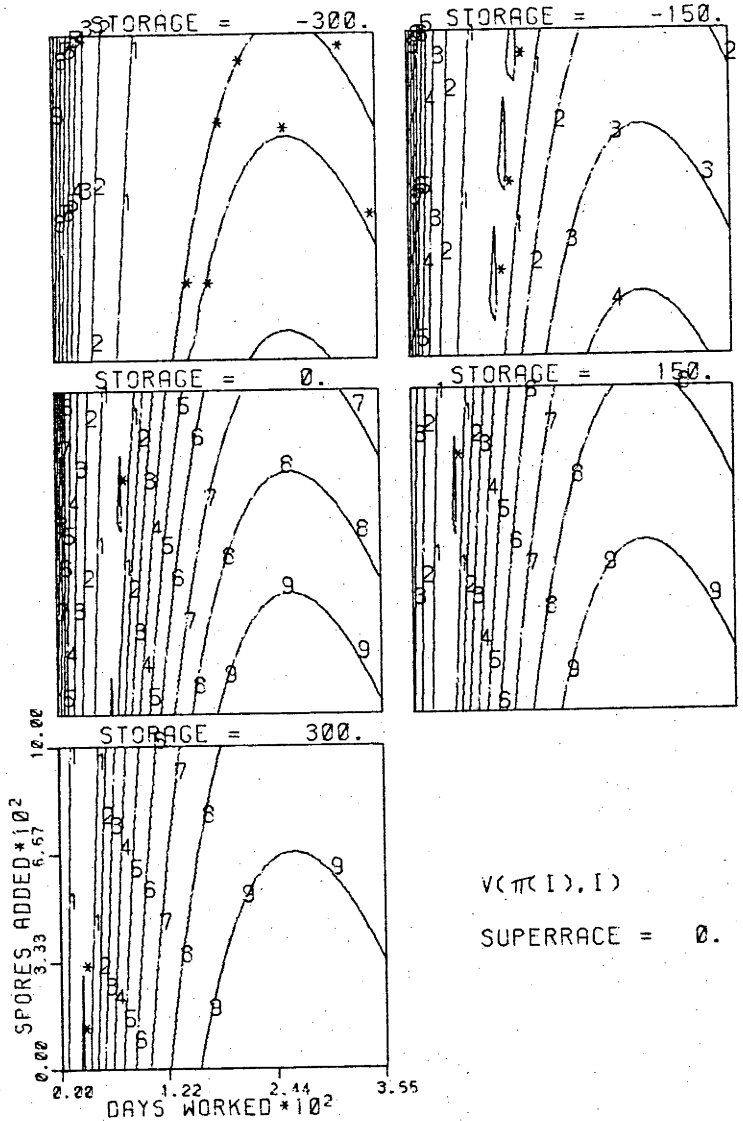

FIG. 6.5.71C: SUBSISTENCE WITH GROWING PLANT

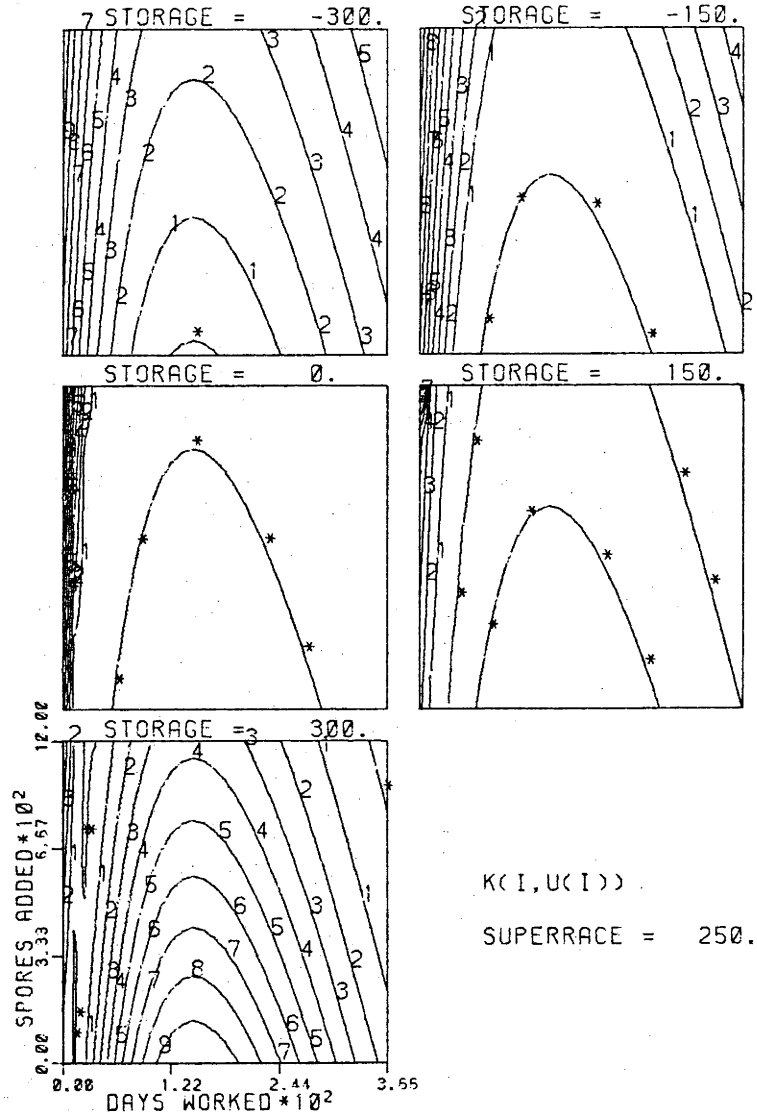

FIG. 6.5.72A: SUBSISTENCE WITH GROWING PLANT

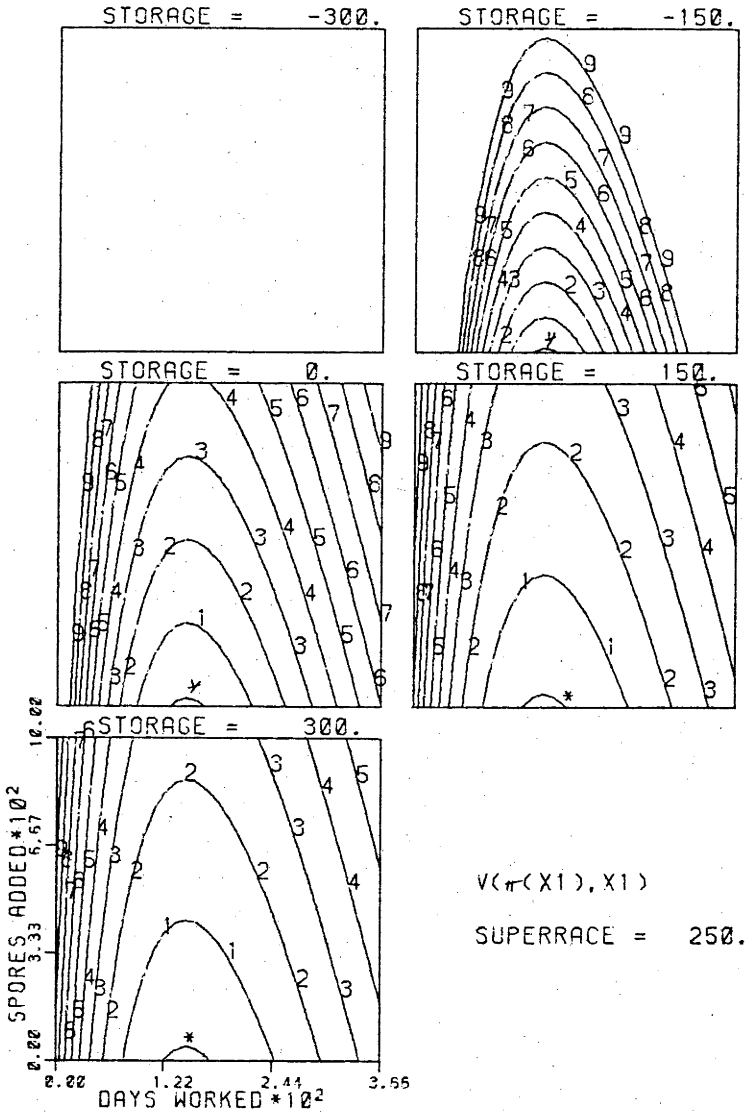

FIG. 6.5.72B: SUBSISTENCE WITH GROWING PLANT

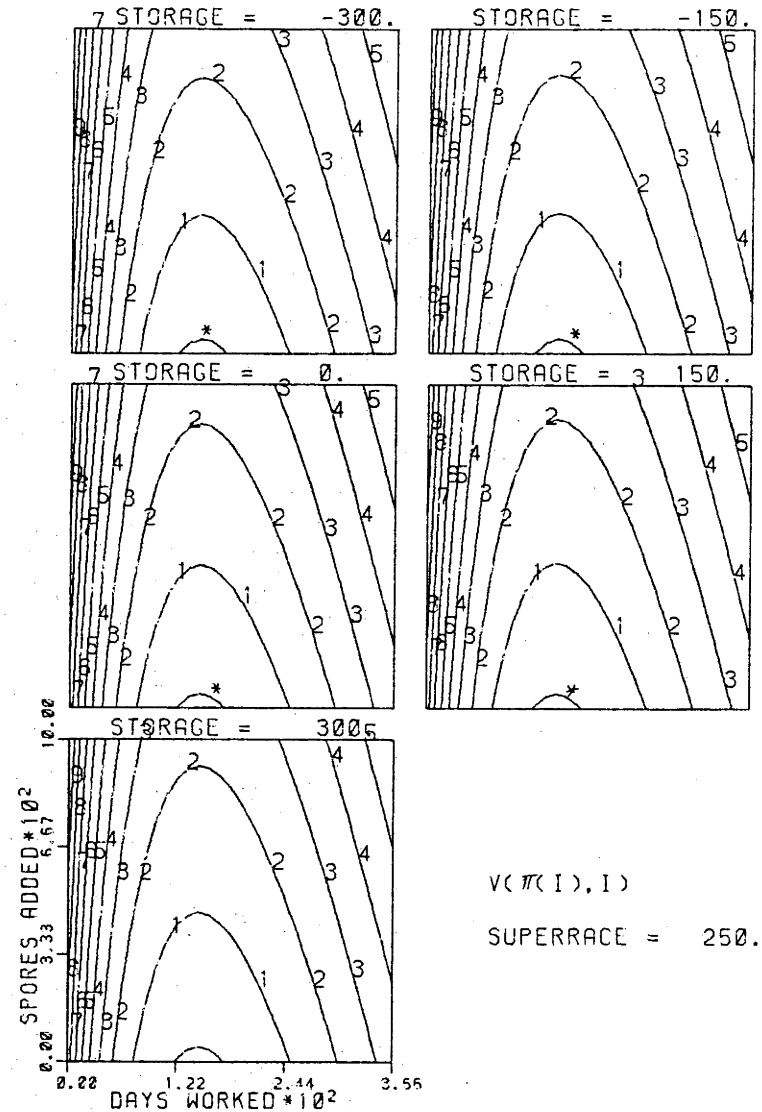

FIG. 6.5.72C: SUBSISTENCE WITH GROWING PLANT

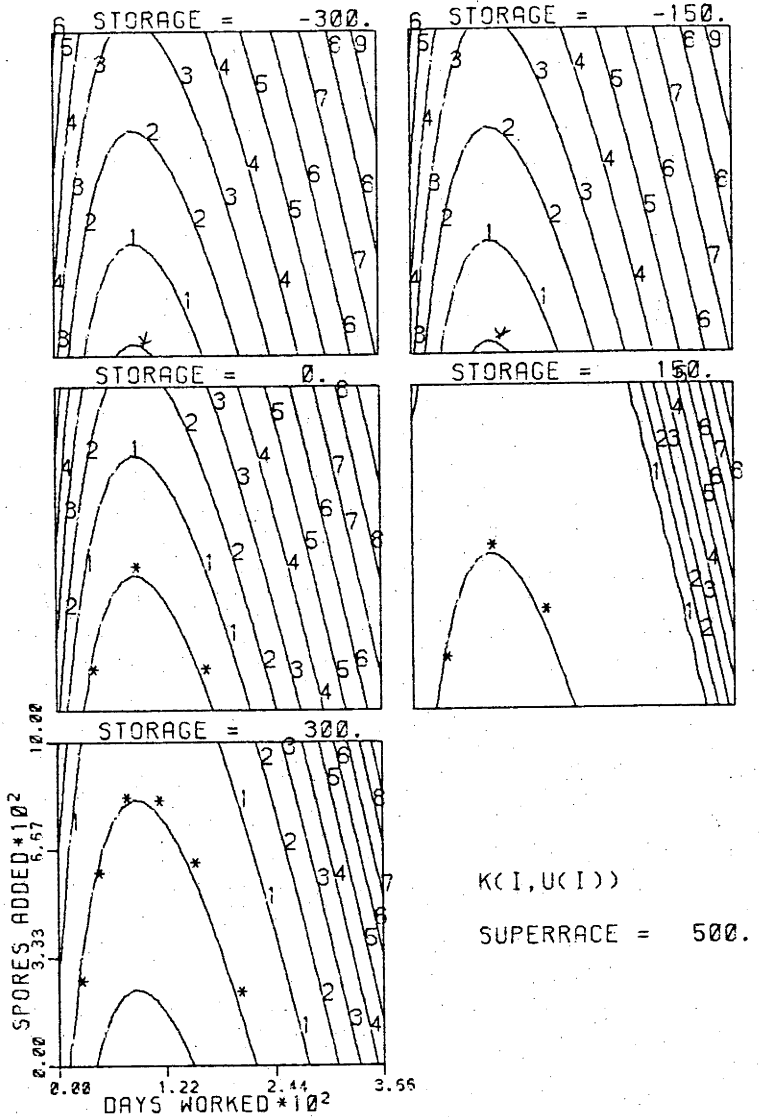

FIG. 6.5.73A: SUBSISTENCE WITH GROWING PLANT

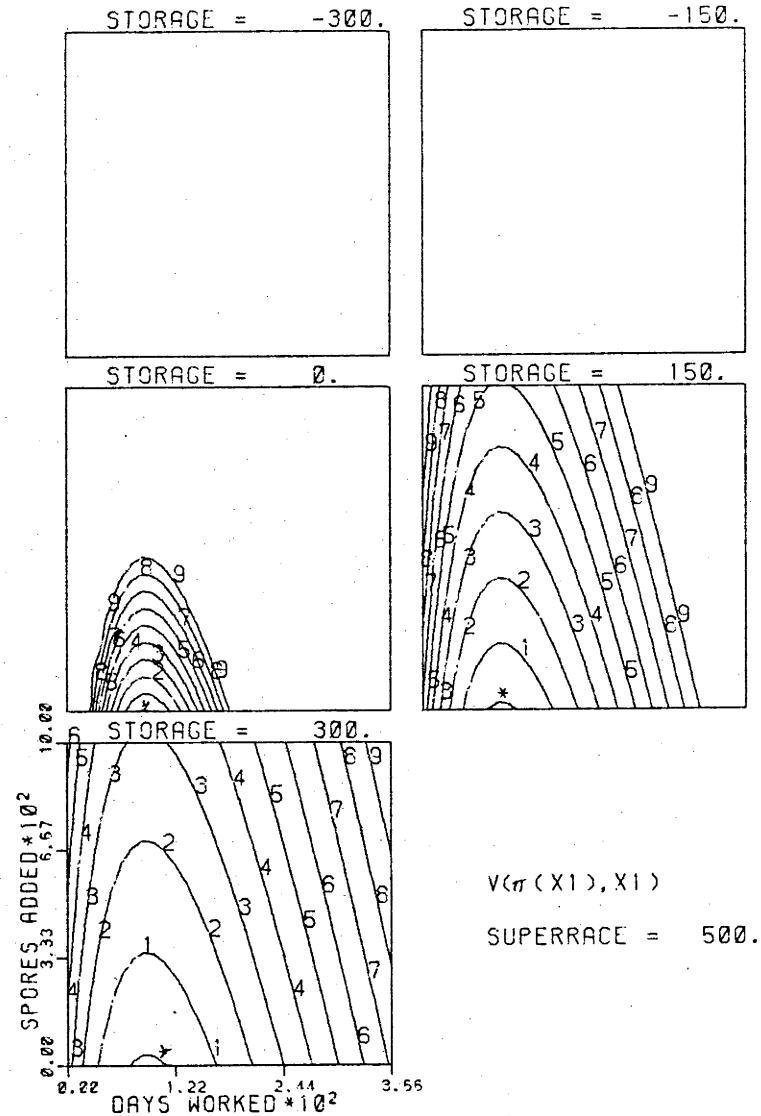

FIG. 6.5.73B: SUBSISTENCE WITH GROWING PLANT

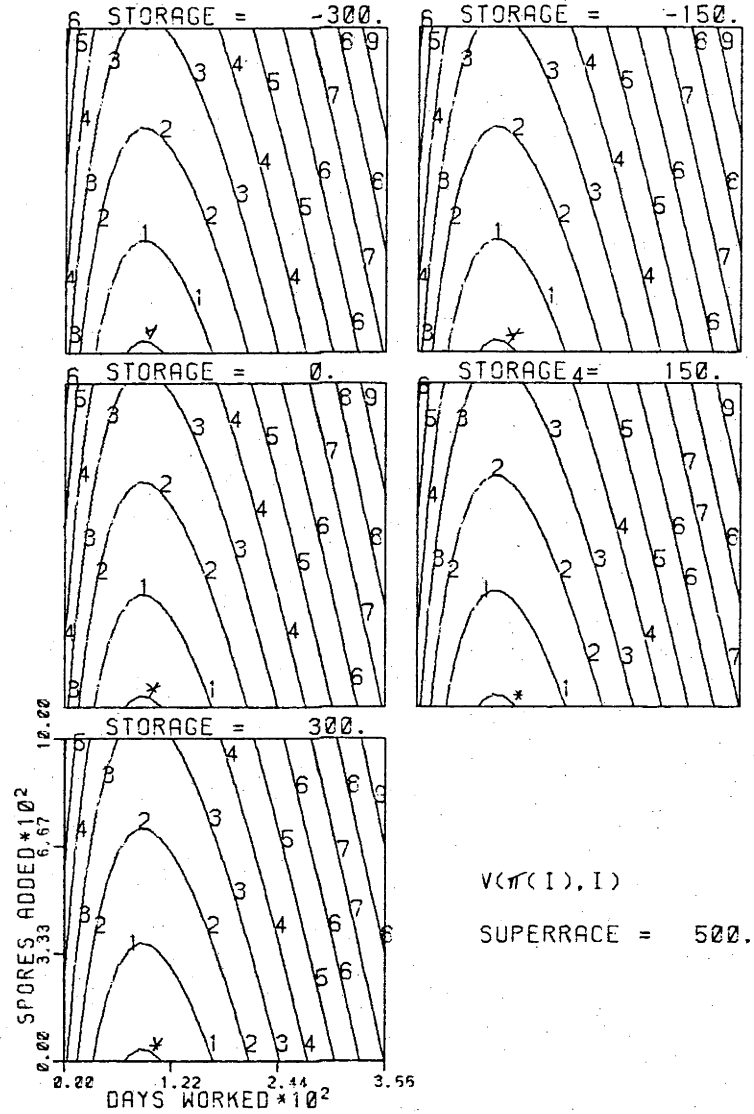

FIG. 6.5.73C: SUBSISTENCE WITH GROWING PLANT

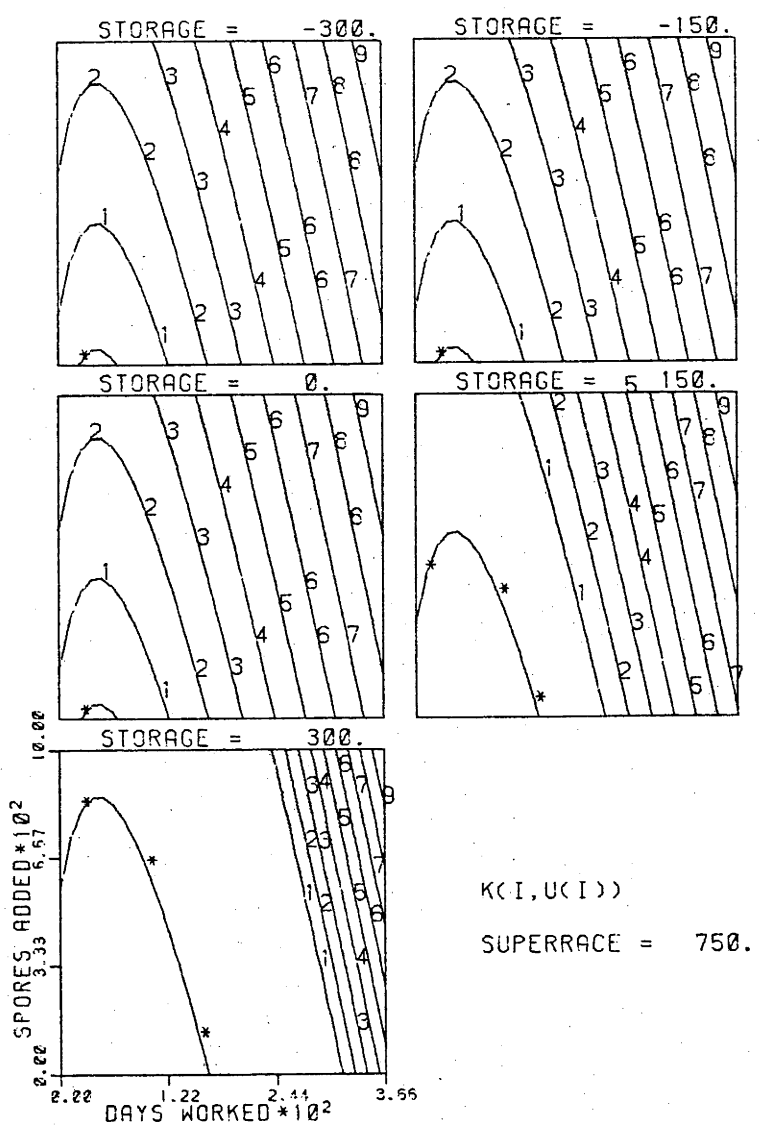

FIG. 6.5.74A: SUBSISTENCE WITH GROWING PLANT

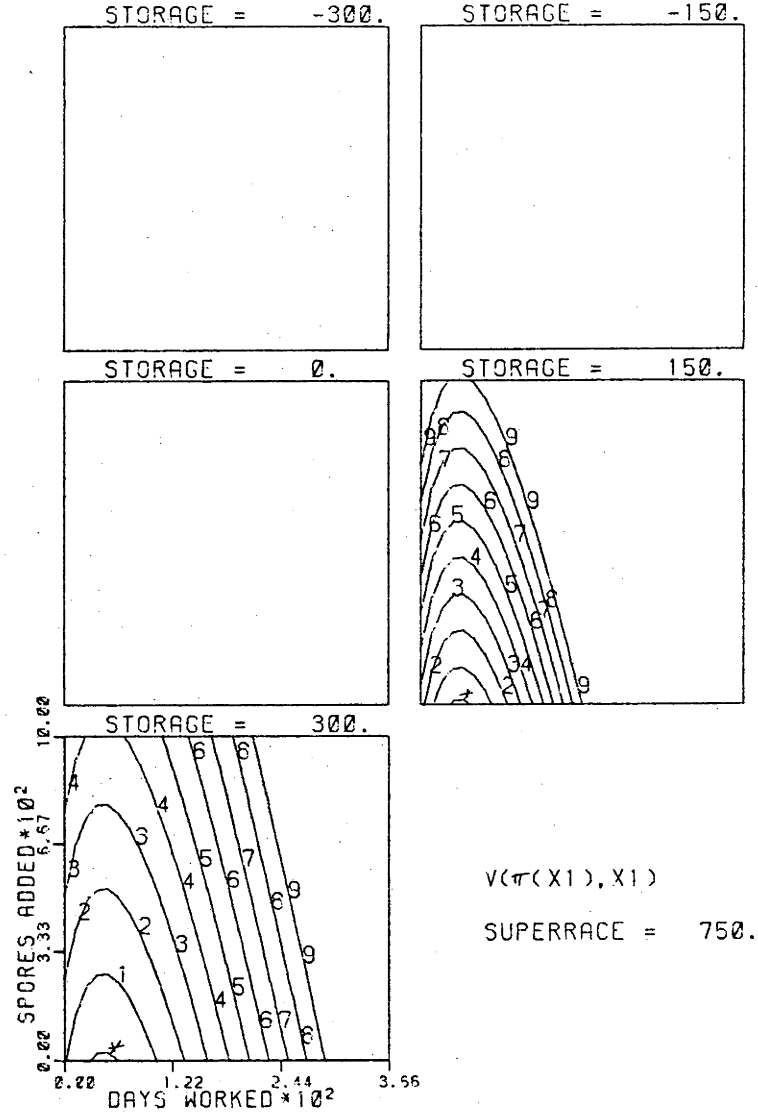

FIG. 6.5.74B: SUBSISTENCE WITH GROWING PLANT

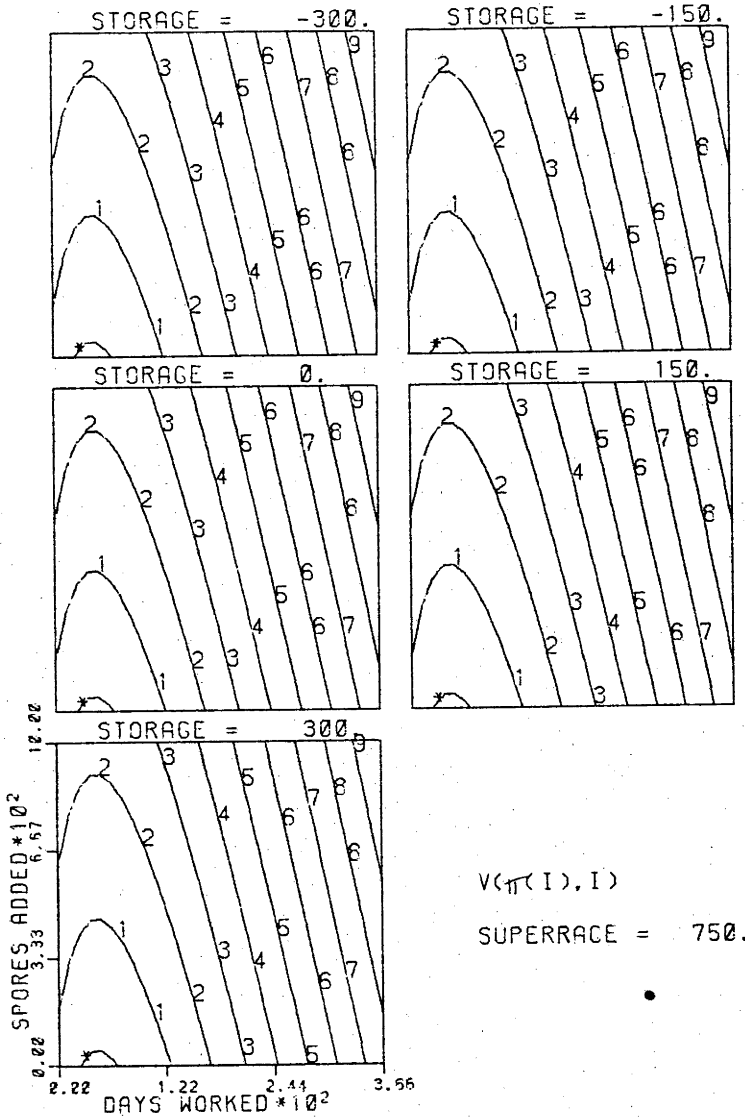

FIG. 6.5.74C: SUBSISTENCE WITH GROWING PLANT

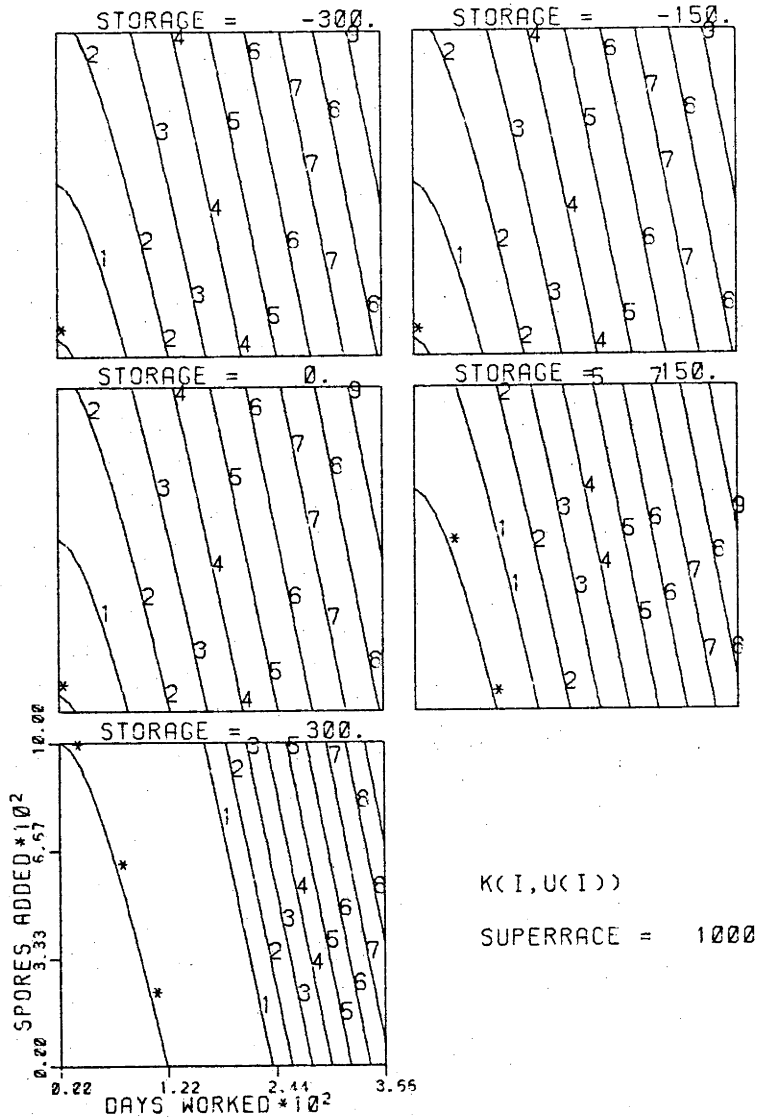

FIG. 6.5.75A: SUBSISTENCE WITH GROWING PLANT

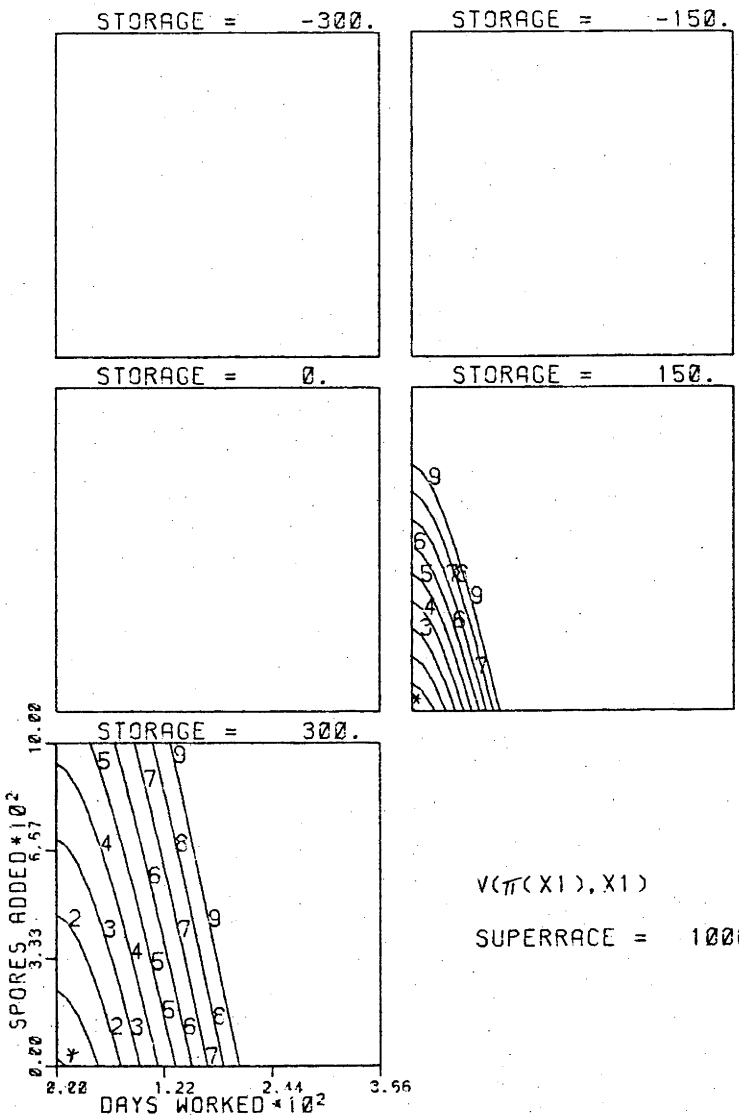

FIG. 6.5.75B: SUBSISTENCE WITH GROWING PLANT

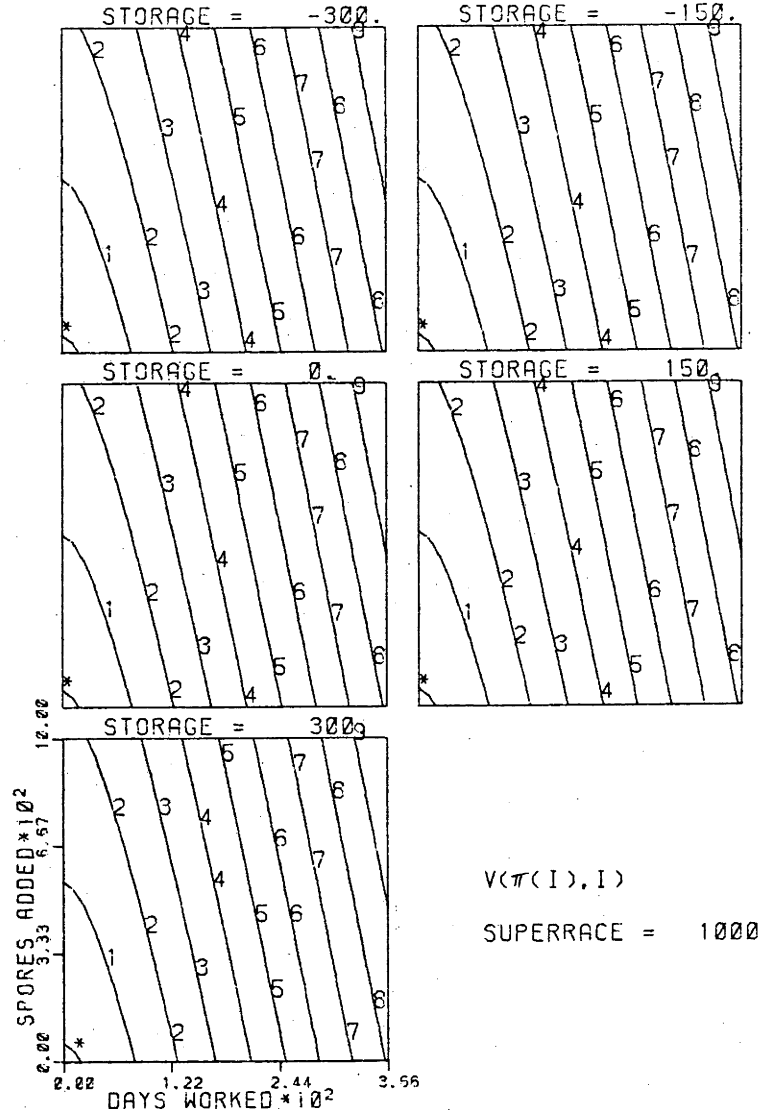

FIG. 6.5.75C: SUBSISTENCE WITH GROWING PLANT

FIG. 6.5.76: SUBSISTENCE + GROWTH: WORK THROUGH TIME

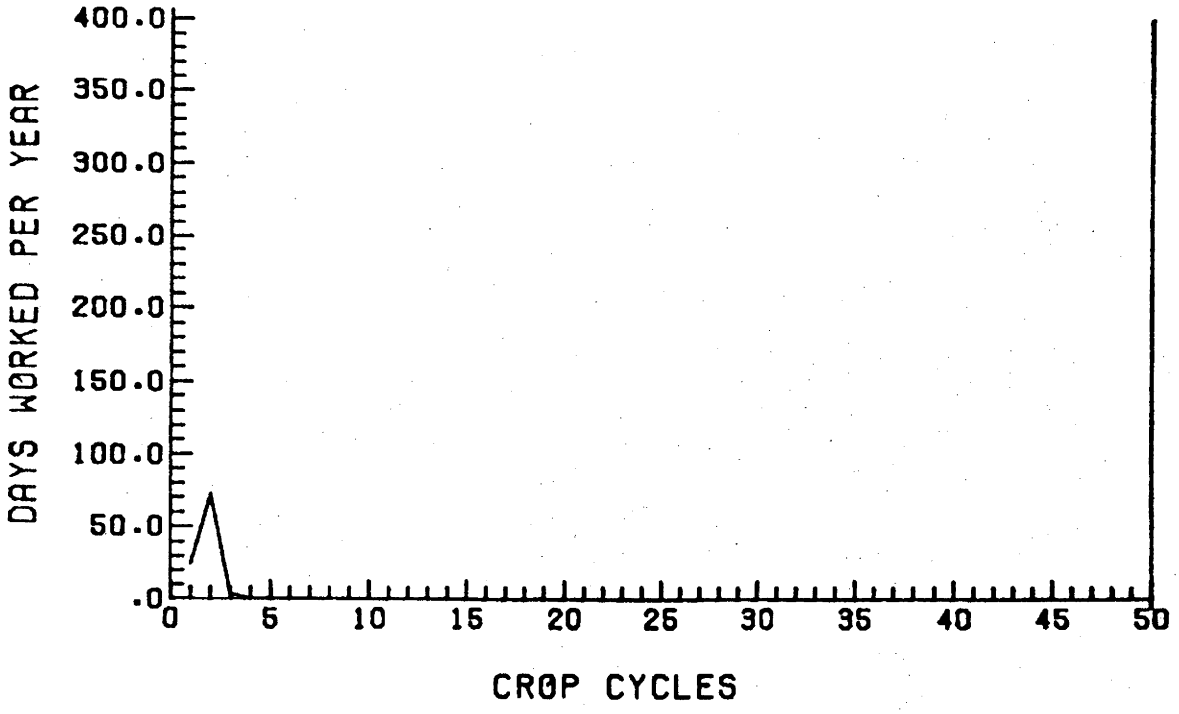

FIG. 6.5.77: SUBSISTENCE + GROWTH: COST THROUGH TIME

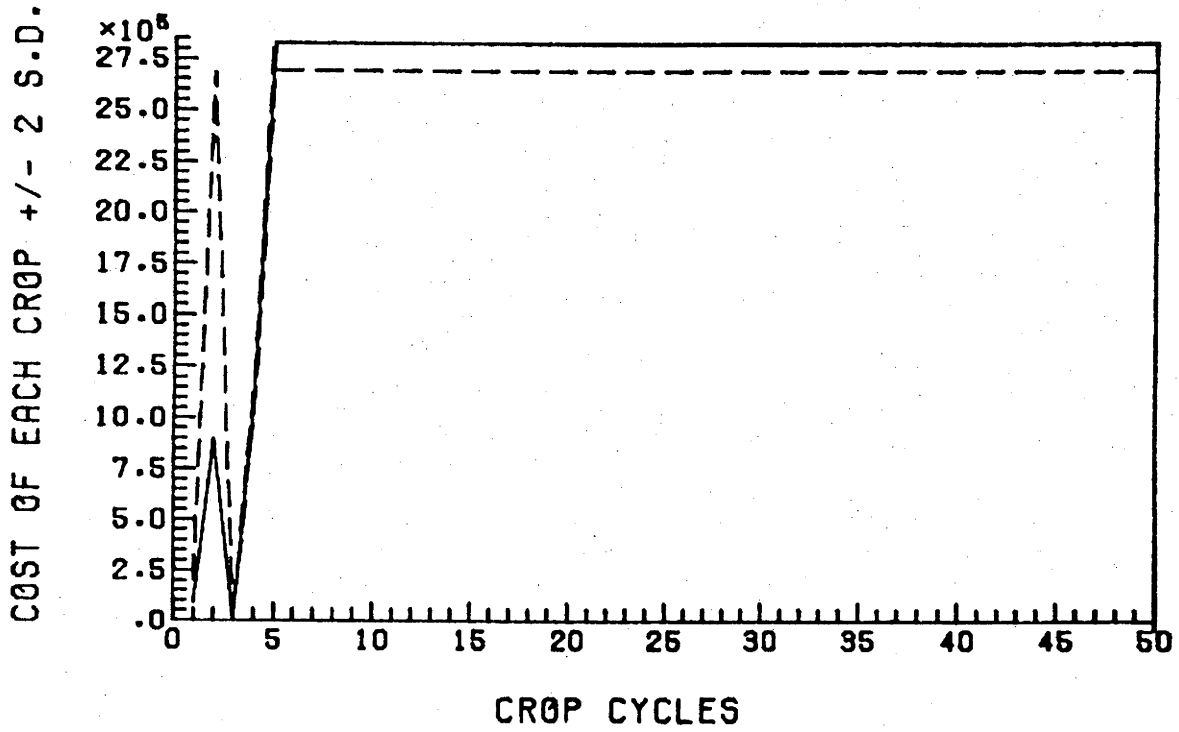

FIG. 6.5.78: SUBSISTENCE + GROWTH: YIELD THROUGH TIME

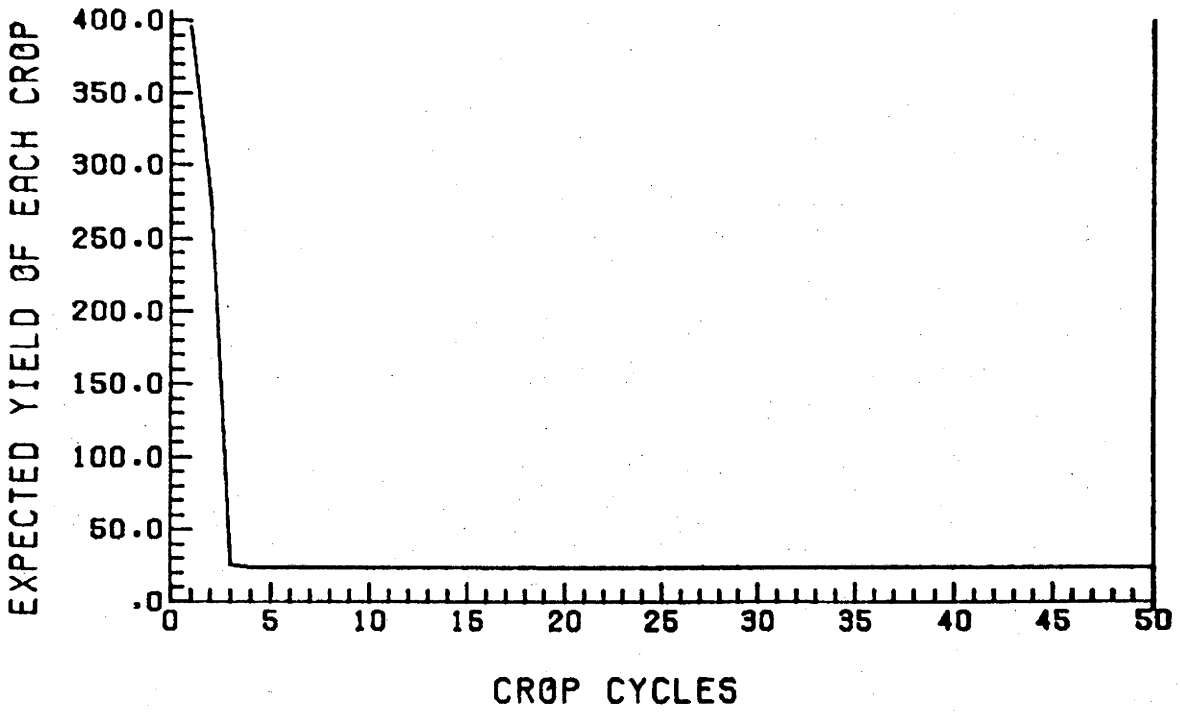

FIG. 6.5.79: SUBSISTENCE + GROWTH: STORAGE THROUGH TIME

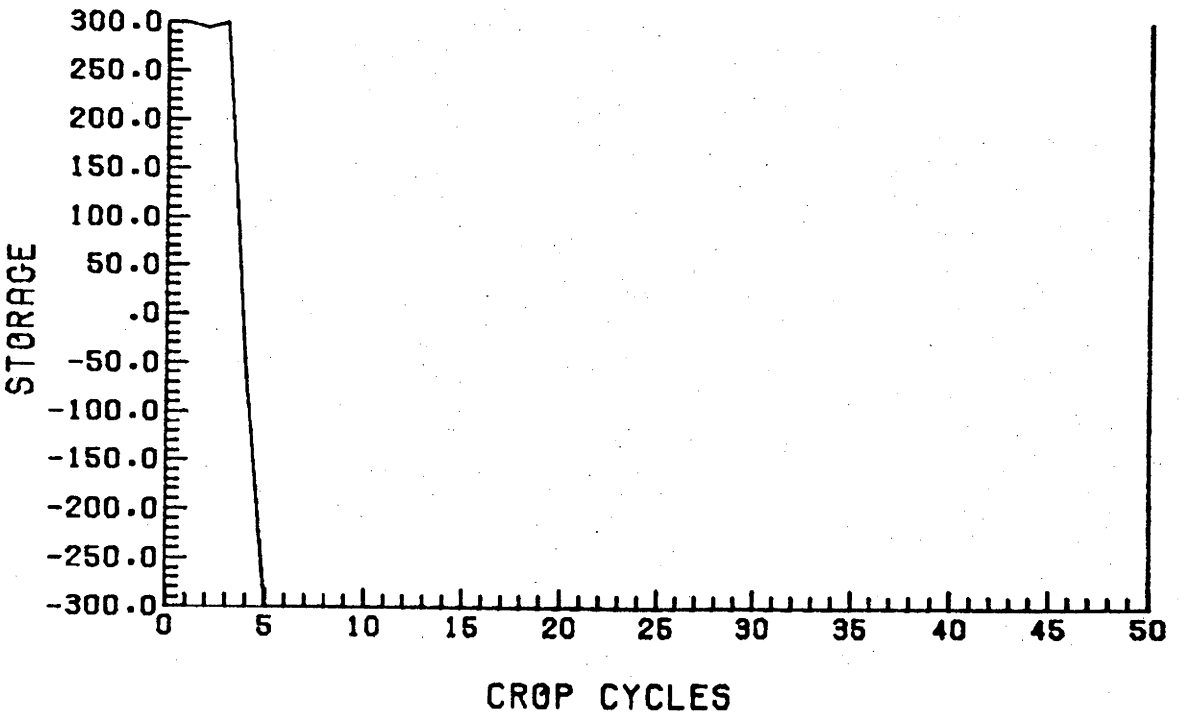

FIG. 6.5.80: SUBSISTENCE + GROWTH: SUPERRACE THROUGH TIME

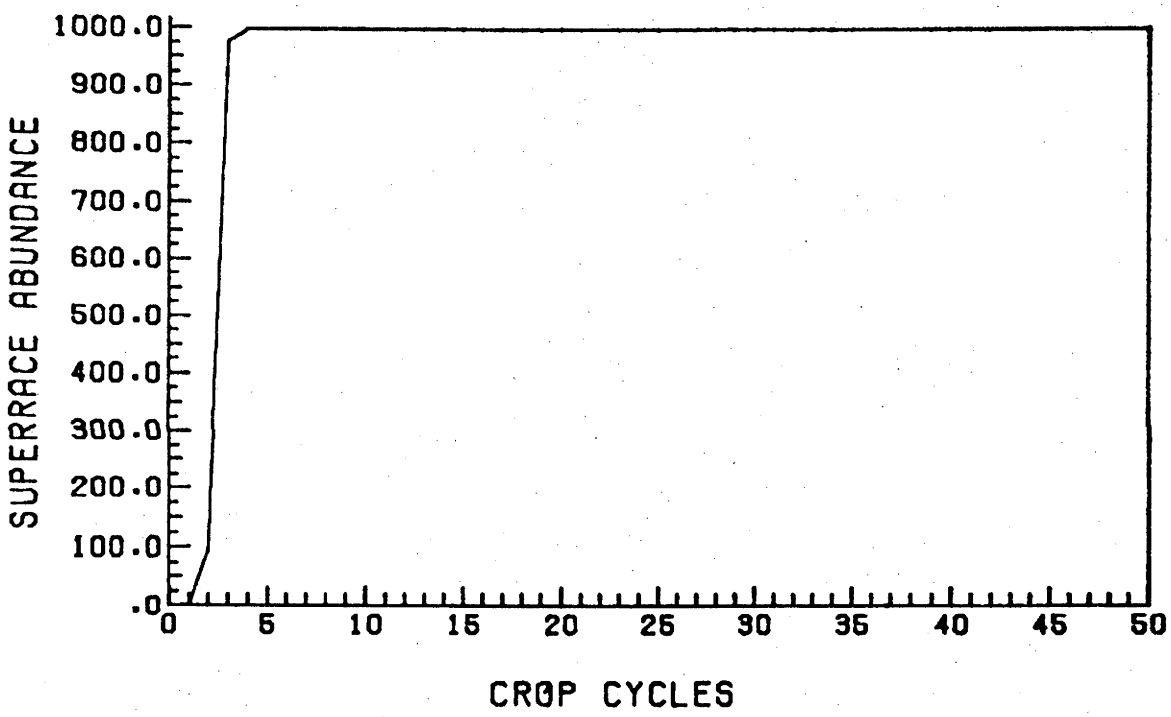

FIG. 6.5.81A: REGRESSION OF SUBLINE USED ON SUPERRACE ABUNDANCE AT MULTILINE POINTS

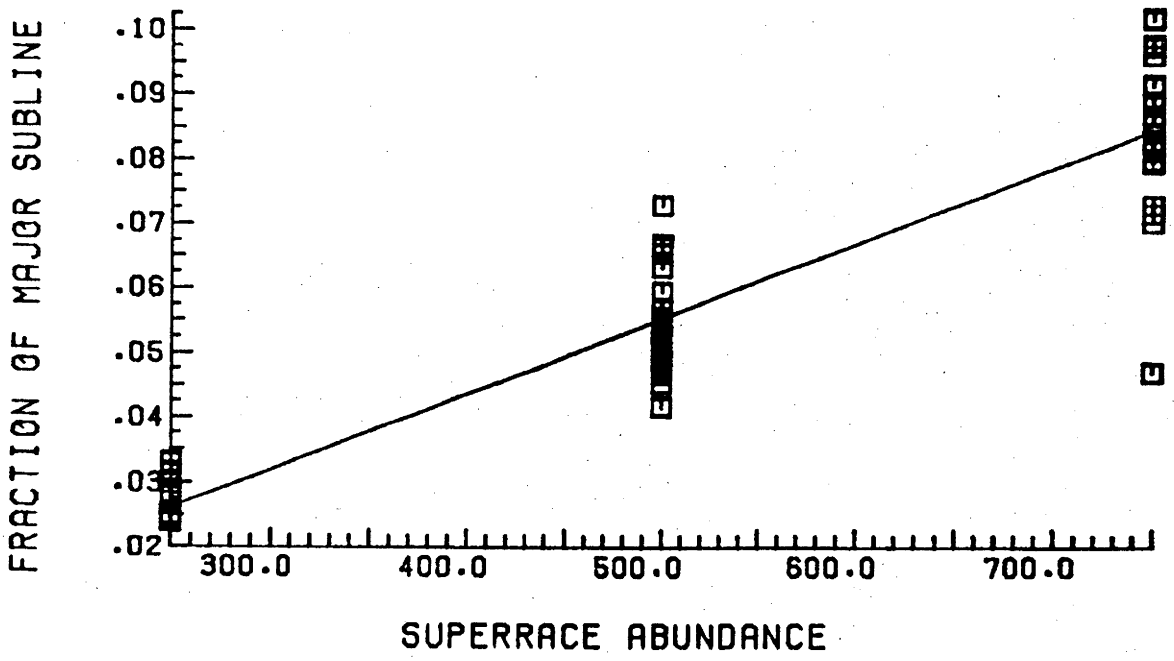

FIG. 6.5.81B: REGRESSION OF CHANGE IN WORK RATE ON CHANGE IN YIELD AT MULTILINE POINTS

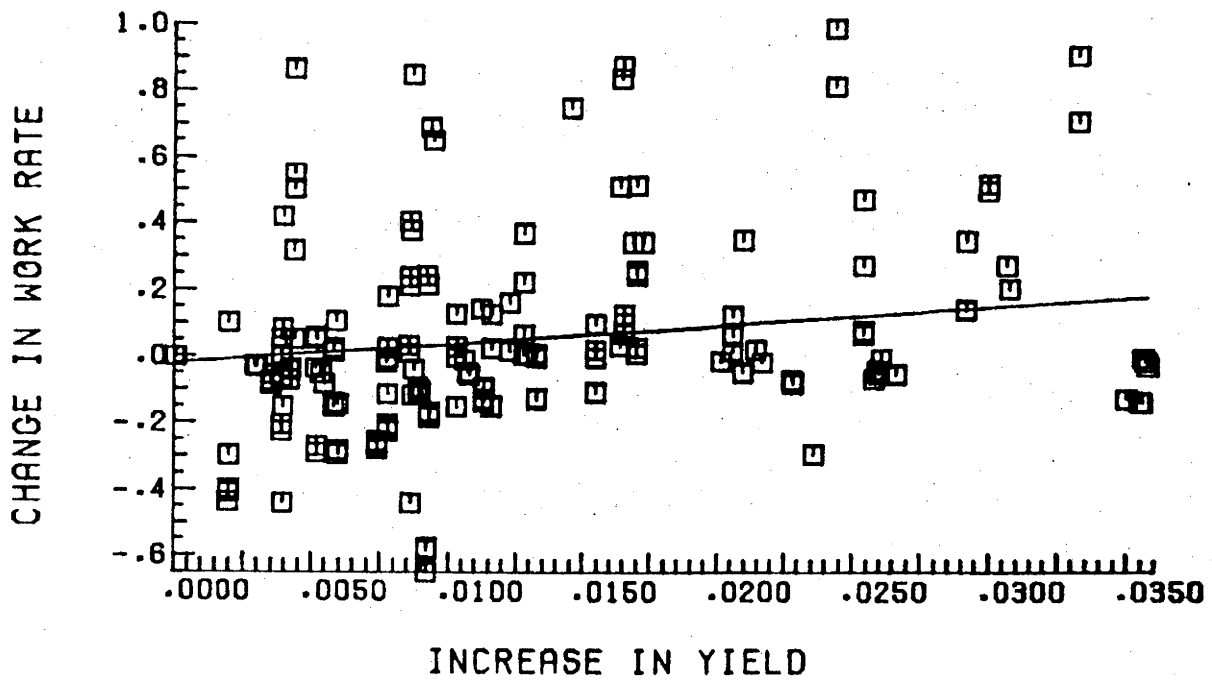

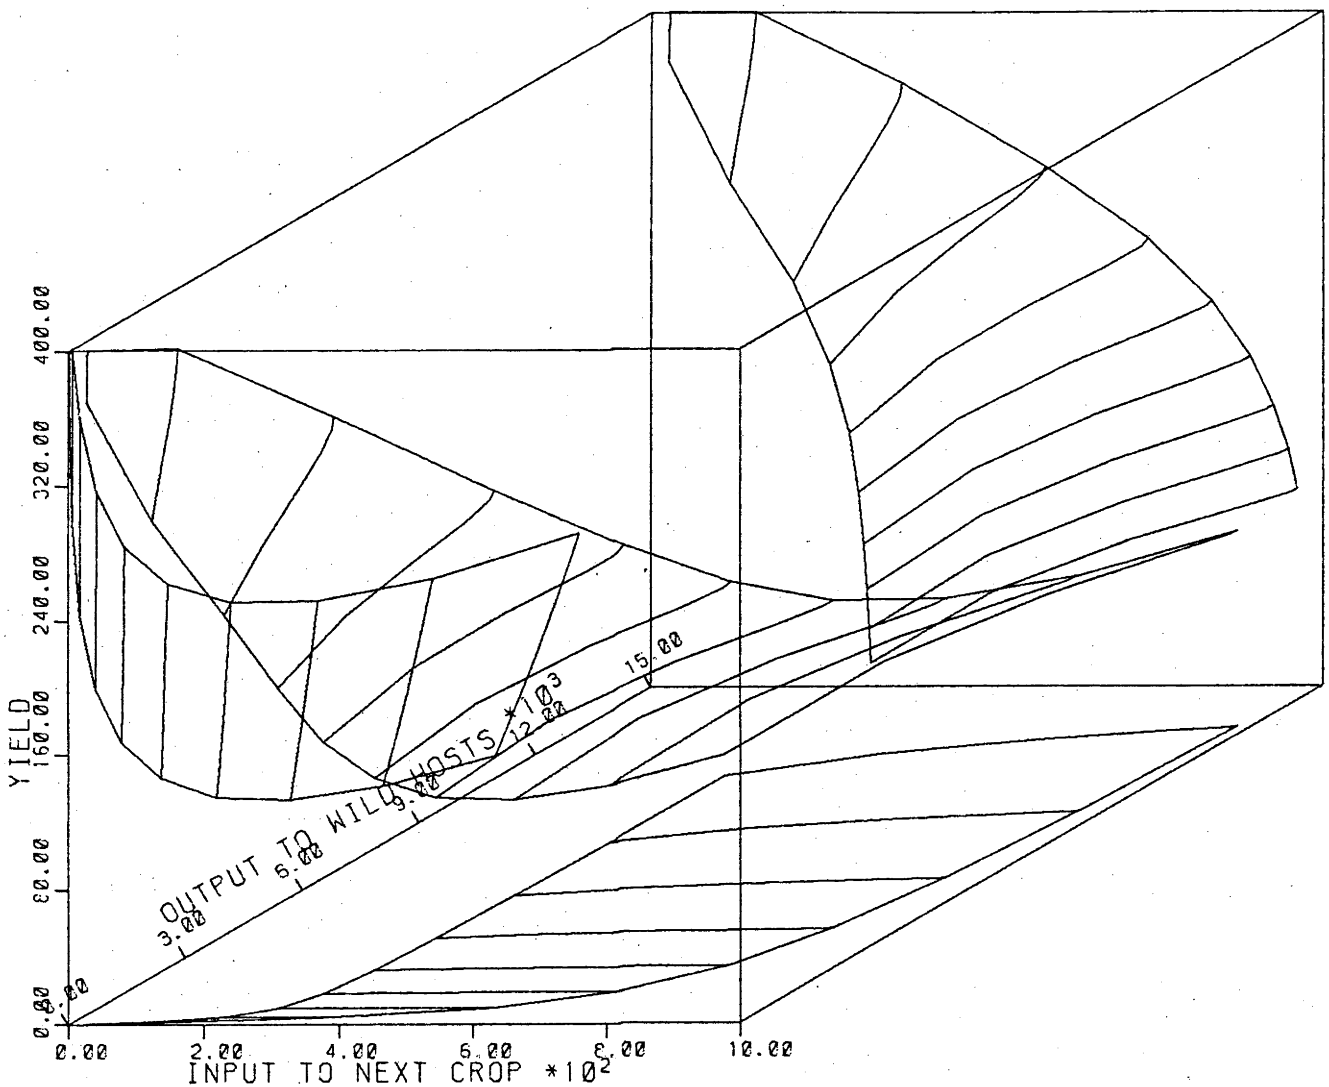

FIG. 6.6.1: TRADEOFFS WITH GENERAL RESISTANCE

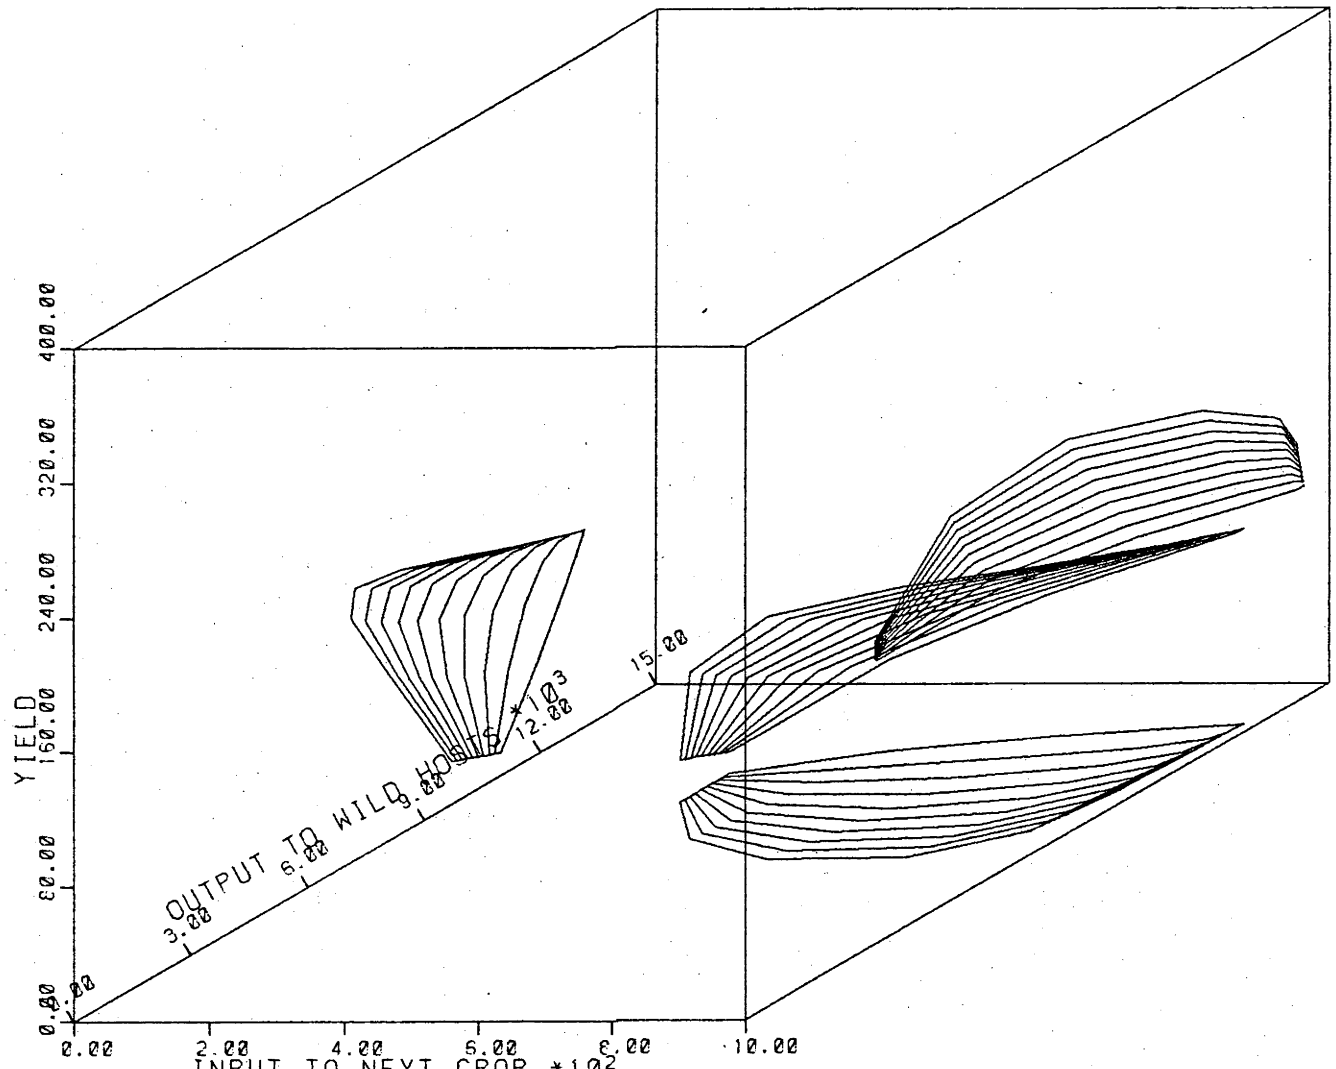

FIG. 6.6.2: TRADEOFFS WITH CROSSPROTECTION

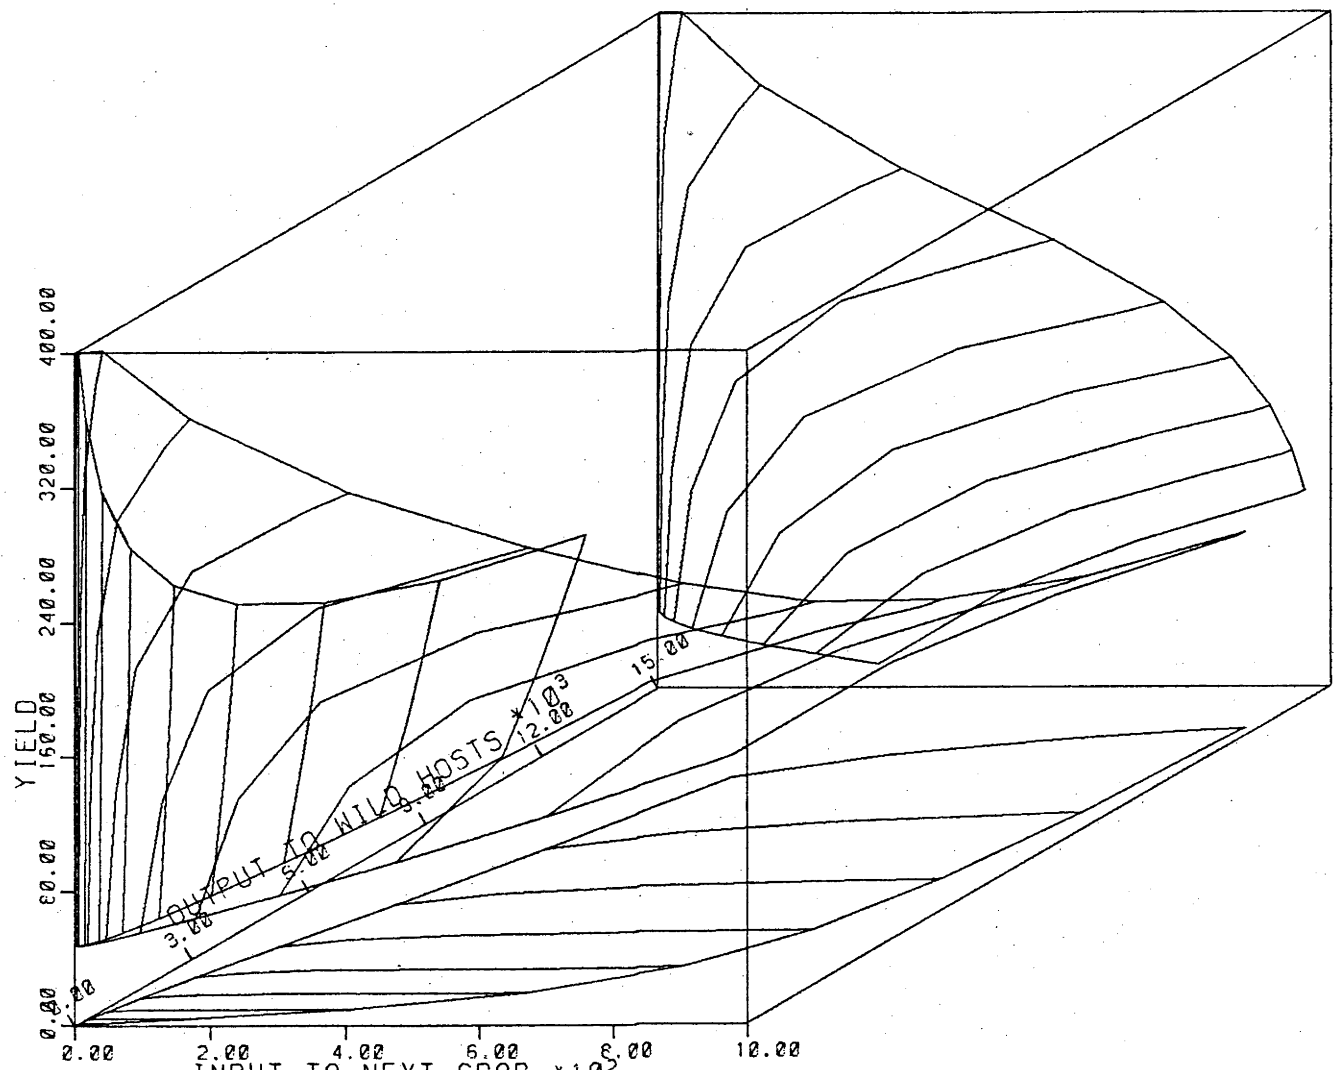

FIG 6.6.3: TRADEOFFS WITH STABILISING SELECTION

7. Conclusions

7.1. Introduction

In this Chapter four things will be done. The conclusions of the study will be summarised. The limitations of the study and the ways in which they might be reduced in further work will be discussed. The prospects for applying the lessons of the study in practice will be considered. Finally, a perspective on the different approaches to breeding crops for disease resistance will be given.

7.2. Conclusions of the Study

The conclusions of the study emerge directly from the experimental results and may be summarised as follows. An examination of the results of the use of disease resistance genes using simulation modelling techniques, has shown that, from the point of view of optimal control theory, there is nothing inherently self-contradictory in the idea of using multiline crops despite the fact that they hasten the appearance of the superrace and probably increase the disease cover on the crop. The results however also suggest that the optimal use of multilines is transitional, unless certain special conditions obtain, in the sense that their use is to smooth the passage from a state in which the simpler races are abundant to one in which the superrace dominates. It appears that multilines would only be optimal in the long term if some host-pathogen interaction such as an extremely strong cross-protection effect increases the yield of the multiline over the yield of the pure superline under the same superrace input, or if stabilising selection acting against the superrace outside the crop increases the equilibrium yield of the multiline over that of the superline. As suggested in section (3.4) the artificial use of spores of

the simpler races was an important feature of the use of multilines in the optimal policies. However the imposition of a cost on the use of spores prevented a direct test of Theorem 2 of Chapter 3.

The results have further suggested that the existence of stabilising selection is not, as had been supposed (Browning and Frey, 1969), necessary for there to be a use for multilines. The precise significance of stabilising selection for optimal gene use remains unclear. On the one hand there is a suggestion that its presence may increase the time period over which multilines are optimal, and even introduce the possibility of continuous use. On the other hand, in a multiline with stabilising selection the avoidable yield losses produced by using a multiline are likely to be increased. Moreover, a multiline which relies on stabilising selection for its advantage is likely to be more dependent on the behaviour of the pathogen population outside the crop than one which does not, and this external behaviour of the pathogen is the most difficult part of the crop/pathogen system to investigate quantitatively.

It has also been found that a cross-protection effect, in the form modelled here, is conducive to multiline optimality. The evidence is conflicting here too. On the one hand an extremely strong cross-protection effect is necessary for any obvious and easily-measured change in epidemic dynamics to take place. On the other hand, it appears that an effect little stronger than some that have been reported can be sufficient to make a quite radical change in the optimal policy. This latter result may mean that the baseline epidemic of the study was fortuitously close to being an epidemic for which multilines were optimal. This explanation is given some support by the fact that multilines were used in the (unconverged) subsistence baseline game.

The last result of this study which bears directly on the use of multilines is that the most important biological factor in determining whether multilines are ever optimal is related to the natural rate of increase from year to year of the superrace of the pathogen and the relation between disease cover and yield loss in the crop. This factor does not correspond precisely to any parameter of the model used, but it appears that if the natural saturation level of the superrace corresponds to an intermediate level of yield depression rather than to total loss of the crop, then multilines are more likely to be optimal.

The other significant results obtained from the study relate to the place of multilines among other modes of use of disease resistance genes. Since the main mechanism inducing optimality is the interference of the simpler races on the crop (except where stabilising selection is present), we can deduce that patterns of simultaneous use of multiline components as pure lines in separate regions will be less effective than multilines. The regions planted to the superline component will be less protected by the pathogen interaction because of the lack of nearby plants of the simpler lines, while farmers in the regions planted to the simpler lines will suffer excessive losses. The asynchronous use of pure crops of the simpler lines appears likely to suppress the simpler races more than the superrace (in the absence of adaptation effects - section (2.3.4)), and so would be expected to be less effective than the use of a multiline.

Finally, the essence of dealing with disease problems in crops of the type simulated is to reduce as far as possible the probability that the race which currently induces the greatest equilibrium yield loss will disperse successfully to susceptible host tissue. If this

cannot be done by the introduction of a variety with a high level of general resistance, then fallows, rotation with another crop or intercropping seem likely to be more effective than the use of multilines. In this connection it is worth remembering that the controversy about the dwindling variability of crop gene pools is not only a controversy about decreasing variability within, for example, the wheat crop, but about reliance on a very limited number of staple crops (Day, 1973). If the production of newer, more resistant version of these few crops is indeed a road that has a dead end, then the alternative routes that must be taken involve drastic changes in agribusiness practice and in food habits (and even population size) that will accommodate rotations that cause shortages in some years, intercroppings that are difficult to farm mechanically, and the development of little-used species as staple crops. On the evidence of this study, multilines can only be an adjunct to these processes.

7.3. Limitations of the Study

The limitations of this study are to some extent voluntary ones, motivated by the wish to remain in contact where possible with the current level of debate on resistance gene use. To a much larger extent the limitations are those of time, computation, and lack of relevant information. They can be listed under three headings as follows.

7.3.1. Lack of Experimentally Determined Parameter Values

While the lack of information on the real values of the parameters used is to be expected in an exploratory theoretical study, it means that it is still not known whether multilines are optimal in even one real crop. At present only a tiny minority of the parameters used can

be estimated from the literature, and in order to achieve useful results it was necessary to adopt the approach of tuning the model so that the disease free state was one of abundance and the baseline epidemic represented a serious disease problem. The difficulty is simply that physiological and pathological studies are not conducted with whole plant or crop modelling in mind; it is hard to see how they could be in the absence of a specific objective at the whole plant or crop level. In section (7.4) a systematic approach to this problem will be suggested for the case where the objective is to discover the optimal mode of gene use for a particular crop.

7.3.2. Lack of Realism in Modelling

This limitation was partly dictated by the current state of debate, and also by the impossibility of computing optimal controls for any epidemic that took much longer to simulate than the one used, and by the lack of published ideas on how to model the interaction of a plant and its pathogen. The realism that was introduced by the use of a simple growth model is only a start. The next candidate for inclusion should be a differential deposition of spores between the host plant and other plants as in Kiyosawa and Shiyomi (1972). Since the model used here lacks this feature it will tend to overestimate the disease escape of susceptible plants surrounded by resistant ones, and thus distort the development of the simpler races.

7.3.3. Computational Deficiencies

The method used in this study is both computationally laborious and, because of the use of linear interpolation, prone to various kinds of error which have already been discussed. In retrospect a more efficient way to proceed might have been to construct a large, multi-

dimensional look-up table of the epidemic input/output relations for each experiment. This would have been, for the type of epidemic simulated here, an eight dimensional table (3 state variables + 4 control variables + an output vector) and would have required special handling even in a large computer. In the long term, however, the number of function evaluations needed to find an optimal control might be handled faster by a look-up and interpolation routine than by repeated simulation.

It is also desirable, though perhaps not practicable, that linear interpolation be replaced by a process in which interpolation functions are developed in parallel with the convergence of the control, so that the interpolation method conforms to the emerging shape of the control law and the cost function. It would obviously be difficult to do this for the complex functions discovered during the experiments on the agribusiness criterion. For a cost function like that of the subsistence criterion such an approach should improve the results of the control and the accuracy of estimation of the expected cost. In any case the intelligent choice of an interpolation method demands some familiarity with the results of computation in a specific case before a good functional form can be chosen.

7.4. Directions for Further Research

Some of the possible directions of further research into the subject of optimal gene use are explicit in the discussion in the last section of the limitations of this study. There remain some important topics that deserve separate consideration. They may be classified as theoretical, biological and practical.

7.4.1. Theoretical Tasks

The principal theoretical problem that has been left unsettled by this study is that of defining what biological properties of a crop/pathogen system induce not only the optimal use of multilines for short periods but their continued use for long periods after (or before) the arrival of the superrace. This study suggests that the existence of stabilising selection leads to the optimality of prolonged use of multilines, yet none of the results described confirmed this, despite the advantages offered in terms of equilibrium yield. This may have been the result of a deficiency in the optimiser, but there was also the disadvantage of a more variable yield in the particular form of stabilising selection modelled here.

The first step in attacking this theoretical problem by the methods of this study would be to consider a series of systems in which the natural saturation level of the superrace corresponded to successively greater yield depressions. For the subsistence criterion at least it appears that for an intermediate range of these depressions, it would be both important and feasible for the farmers to influence this saturation level by the use of multilines. Whether this would occur only in the presence of stabilising selection is not deducible from the present results. Similarly it is not clear whether the special form of stabilising selection modelled here encourages the prolonged use of multilines because of some feature which is not representative of real instances of stabilising selection. For instance, the inclusion of saturation of the wild hosts seems to play an important role, and it may be that a simple differential survival rate of the superrace and the simpler races would not cause the long term yield advantage of multilines over the superline which is the most interesting feature of stabilising selection in this study. While it is important to note

that in many cases in the study the introduction of suboptimal multilines did little harm, it is still the lesson of the study that the use of multilines was transitional and made only small cost reductions. Until there is evidence that persistent use of multilines will bring consistently measurable cost reductions, it would be premature to devote resources to their commercial development.

7.4.2. Biological Tasks

The principal biological problem that is relevant to the understanding of the optimal use of resistance genes is that of investigating the nature of general resistance. This is important because of the role which the level of general resistance plays in determining the optimal use of major genes and, in particular, the optimality of multilines. It is even more important because the possibility of breeding for general resistance (as it has been defined in this study, following Caldwell (1968)), is itself in doubt. While every plant has its own level of general resistance, it was pointed out in Chapter 2 that there are no good explanations of the way in which a single gene or a group of genes will sometimes apparently raise this level. (Such effects must be described as "apparent" because the presence of a form of resistance permanently invulnerable to pathogen mutation cannot be proved conclusively.) If there is to be any confidence that general resistance can be adjusted so as to harmonise with the use of major resistance genes in multilines, its mechanism must be understood. We need to know why, in the pregnant words of I.A. Watson (1970b) some combinations of genes "place before the fungus barriers which are more difficult to negotiate".

7.4.3. Practical Tasks

The principal practical problem in multiline culture is, from the results of this study, to decide whether multilines are optimal for a given crop. Although both those in favour of multilines and those who are sceptical can draw support from the results, the important practical point is that the decision is a quantitative one, and that making it is likely to be very complex.

In a practical situation the simplifications of section (3.5.2) cannot usually be made. Environmental variability will have to be taken into account. There will be difficulty in establishing a correspondence between the infected area of the model which homogeneously produces spores and which does not grow vegetatively and the lesions on the real plant. There is the possibility that not one but several pathogens may be involved. And there is the probability that the pathogen will affect the plant in other ways than simply occupying photosynthetic tissue.

If the task of determining multiline optimality or suboptimality is to be attempted, the most promising line of attack is to start by establishing a yield/superrace output curve for a replacement series experiment, like the ones simulated for different parameter values in figures 6.6.1-3. It would not be possible at first to measure the input to the next crop because even a large experiment like those of Cournoyer (1970) would not have much effect on the general population level of a rapidly dispersing pathogen. However if a substantial ratio of superrace depression to increase in yield loss were found in the multiline plots, then this would encourage the taking of a second step. ("Substantial" in this context would mean "statistically significant at several experimental sites in several years".)

The second step would be to attempt to force the epidemic dynamics of the crop/pathogen system into the mould of an epidemic model of about the complexity used here. The growth rate of the pathogen and the disease cover/yield loss relationship would already be known, and comparison of the pathogen growth rates in the mixtures of the replacement series might allow the estimation of the probability of successful spore dispersal as a function of host density. If these pieces of information were combined with what was known about off-season dynamics and with an economic model appropriate to the crop's usual market, a trial optimal control could be computed. If the optimal policy for this crude model suggested persistent use of multilines, then refinements and improvements of the type suggested in this chapter should be made until either the plant breeder (and his employers) have confidence in the result or until computational resources are exhausted. Confidence in the result would be enhanced by the absence of multiple optima in the policy, because the chance that a maximum cost is close to a minimum in control space is increased by the presence of multiple optima. In any case, if the pure superline were locally optimal it would be very risky to use multilines because suboptimal results would be expected while they were being introduced, as a result of the previous predominance of the superline. Confidence would also be increased if the use of multilines were found to be robust in the minimax sense with respect to as many of the more uncertain parameters as possible. Finally, nothing can increase confidence as much as the prospect of substantial cost decreases from the use of multilines.

7.5. A Perspective on the Use of Crop Disease Resistance Genes

The previous debate about how to use resistance genes has centered on ideas of stability, even though the different points of view in the debate have corresponded to different definitions of stability. Much trouble has been caused by the fact that there are so many possible definitions of stability, and there is no likelihood that a single definition can be agreed on. The traditional view of plant breeding - as instanced in Chapter 1 by Stanton et al. (1934) - sought a situation that was stable in that the pathogen would not exist. The crop/pathogen system would degenerate into the single disease free state.

The multiline point of view as expressed by Browning and Frey (1969) appears to conceive of a crop/pathogen system in which state variables corresponding to all pathogen races take non-zero values and an equilibrium state exists. This concept is an essentially deterministic one, and it is worth mentioning that even in this kind of system, which resembles the standard subject matter of topological dynamics, there are many alternative definitions of stability (Siberskii, 1975). However the differences between these definitions may not have any significance in practice. It may be more important to remember that crop/pathogen systems are stochastic and that definitions of stability in stochastic systems are more elaborate and require the user to think about the expected distribution of system states rather than about single states (Kushner, 1971).

Current crop breeding practice, by contrast with the traditional and multiline alternatives is not explicitly concerned with questions of the stability of the pathogen population, but instead seems to work to produce a state analogous to equilibrium in an irreversible

thermodynamic process (Glansdorf and Prigogine, 1971). The continual tendency of the pathogen to produce new races (analogous to the loss of heat from a resistance element) is countered by the production of new resistant varieties (analogous to the inflow of current). Because of the small number of introductions compared to the number of electrons in a macroscopic resistor, the existence of an equilibrium state in plant breeding (constant and equal rates of irruptions of new races and introductions of new varieties) is hard to test. Whether such a state could be stable, as the analogous state in irreversible thermodynamics is, would be hard to discover. The burden of the points discussed in section (1.2.2) is not that the state is unstable but that the power source is exhaustible.

While it is possible to describe the use of multilines in terms of stabilising the pathogen population as though other plant breeding methods provoke instability, this approach is misleading. The type of crop/pathogen modelled here is inherently stable by most definitions, including stochastic ones. As was suggested in section (3.2.2.2), this type of system resembles a Markov process with a final absorbing state - defined by the appearance of the superrace - which is stable if the superrace cannot be driven to extinction by the other races. On a more detailed level of description, there are only two possibilities for the final state of the chain. Either the simpler races eventually disappear, or some of them will persist. Either possibility represents a stable state. The problem is that the stable state may have economically undesirable properties, as may the trajectories that lead to it.

What should be recognised in deciding on the best use of disease resistance genes is that it is not stability that is important,

but utility. The thought of discovering stability in a pathogen population is attractive to the disinterested researcher, but those who depend on the crop have other priorities. This study has suggested that looking at the use of resistance genes from the point of view of optimal control may clarify debate and result in choices that are sensible as seen by the users of crops. A number of theoretical questions remain unanswered, but these are less important than the biological question of discovering the basis of lasting resistance and the practical question of whether enough information can be gathered and processed so that the quantitatively optimal choice can be made with confidence.

Appendix A - Listing of program OPTIPLANT.

(The parameter values in this listing are those of a baseline-strength epidemic on the growing plant, using the subsistence criterion.)

```
CKM801*OPTIPLANT(1).MAIN
      C   ************************************************************************
      C   *                                                                      *
      C   *  MAIN PROGRAM.  THIS (1) SETS THE MAIN EXPERIMENTAL PARAMETERS,      *
      C   *  (2) SETS MODEL TYPE, (3) SETS THE OPTIMALITY CRITERION,             *
      C   *  (4) DEFINES THE STATE SPACE GRID (RANGE), (5) CONSTRAINS THE        *
      C   *  CONTROL VARIABLES TO THE REQUIRED SUBSPACE BY READING  DAVJST       *
      C   *  (6) SETS WHETHER AN ORDINARY CONTROL OR A MINIMAX IS TO BE          *
      C   *  FOUND (IENAB), AND (7) OPTIMISES BY SUCCESSIVELY CALLING            *
      C   *  "DAVMIN"  FROM EACH POINT OF THE GRID UNTIL THE CHI-SQUARED         *
      C   *  DIVERGENCE BETWEEN ALL THE COSTS OF SUCCESSIVE ITERATIONS           *
      C   *  FALLS BELOW  ERR.  IT THEN INVENTORIES THE FINAL POLICY TO          *
      C   *  SEE IF IT INCLUDES ANY MULTILINES                                   *
      C   *                                                                      *
      C   ************************************************************************
      C
            DIMENSION RANGE(4,5),F(3,3),IVPOS(4)
            DIMENSION REC(625,2),IEPARA(25),EPARAM(25)
            DOUBLE PRECISION DAVJST(10,10),UMIRR(20),SLIM(10,4)
      C
            COMMON VOPT(5,5,5,5,20)
      C
            CALL CKSET (0,500)
      C
            IUPDAT=0
            ISUCCS=0
            READ(1,149),RANGE
        149 FORMAT(4E8.3)
            READ(1,150),IEPARA
        150 FORMAT(/8I6,17(//I6))
            READ(1,155),EPARAM
        155 FORMAT(/3E12.5,22(//E12.5))
            READ(1,160),UMIRR
        160 FORMAT(/D12.5)
            READ(1,165),SLIM
        165 FORMAT(5D12.5)
      C  NOW AN ARRANGEMENT FOR THE INITIAL ESTIMATE OF VOPT
            IENAB=IEPARA(9)
            IBYP=IEPARA(19)
            IGSTEP=IEPARA(20)
            IF(IENAB)175,175,180
        175 PRINT 177
        177 FORMAT(///36H RUN OF OPTIPLANT AS CONTROL PROBLEM///)
            GOTO185
        180 PRINT181
        181 FORMAT(///35H RUN OF OPTIPLANT AS A MINIMAX GAME///)
        185 PRINT186,IEPARA,EPARAM,UMIRR
        186 FORMAT(15H PARAMETER LIST,//7H IEPARA/10I12/10I12/5I12
       1  /7H EPARAM/10E12.5/10E12.5/5E12.5/6H UMIRR/(10D12.5))
            IF(IBYP)5,5,10
          5 DO1I1=1,5
            DO1I2=1,5
            DO1I3=1,5
            DO1I4=1,5
            DO2I=1,16
          2 VOPT(I1,I2,I3,I4,I)=UMIRR(I)
            VOPT(I1,I2,I3,I4,I1)=50000.*(I2+10*I3+I4+(I1-3)**2)
      C  GUESS THAT COST IS INCREASED MAINLY BY SUPERRACE, AND LOWERED BY
      C  POSSESSION OF AMOUNTS OF STORED FOOD
```

```
C
      1 CONTINUE
        REWIND 23
        WRITE(23),VOPT
        END FILE 23
        REWIND 23
     10 CONTINUE
C
        FO=EPARAM(4)
        H=EPARAM(5)
        DO11I=1,3
        DO11J=1,3
     11 F(I,J)=FO
        F(1,2)=H
        F(1,3)=H
        F(3,1)=H
        F(3,2)=H
C  READ CURRENT VOPT
        REWIND 23
        READ(23),VOPT
        REWIND 23
C
C  NOW SET UP GAUSS-SEIDEL ITERATION
        DO15I=1,625
        DO15J=1,2
     15 REC(I,J)=0
        ERR=10
        IGSITR=0
C BEGIN ITERATION BY COMPARING PRESENT AND PREVIOUS COSTS
    100 J=0
        DO20I1=1,5
        DO20I2=1,5
        DO20I3=1,5
        DO20I4=1,5
        J=J+1
     20 REC(J,1)=VOPT(I1,I2,I3,I4,11)
        DIVER=0
        DO25I=1,625
     25 DIVER=DIVER+(REC(I,1)-REC(I,2))**2/REC(I,1)
        IGSITR=IGSITR+1
        IF(IGSITR.LE.10)GOTO29
        ISUCCS=-1
        CALL UPDATE
        GOTO300
     29 PRINT30,IGSITR,DIVER
     30 FORMAT(//30H GAUSS-SEIDEL ITERATION NUMBER,I8,
       1 35H CHI SQUARE FOR (PRESENT-LAST) COST,E14.8)
        IF(DIVER.GT.ERR)GOTO31
        ISUCCS=1
        CALL UPDATE
        GOTO300
     31 CONTINUE
C  IF NOT ITERATE ONCE MORE OVER ALL GRID VALUES
        DO40I1=1,5,IGSTEP
        IVPOS(1)=I1
        DO40I2=1,5,IGSTEP
        IVPOS(2)=I2
        DO40I3=1,5,IGSTEP
        IVPOS(3)=I3
```

```
      DO40I4=1,I2,IGSTEP
      IVPOS(4)=I4
      REWIND 16
      READ(16,170),DAVJST
  170 FORMAT(5D12.5)
C
      CALL DAVMIN(IVPOS,DAVJST,IEPARA,EPARAM,F,RANGE,SLIM)
C
      IF(I2.EQ.I4)GOTO39
      DO32I=1,20
   32 UMIRR(I)=VOPT(I1,I2,I3,I4,I)
      TEM=UMIRR(1)
      UMIRR(1)=UMIRR(3)
      UMIRR(3)=TEM
      TEM=UMIRR(4)
      UMIRR(4)=UMIRR(6)
      UMIRR(6)=TEM
      TEM=UMIRR(14)
      UMIRR(14)=UMIRR(16)
      UMIRR(16)=TEM
      DO33I=1,20
   33 VOPT(I1,I4,I3,I2,I)=UMIRR(I)
   39 CONTINUE
      ITMLFT=JOBSUP(1)
      IF(IUPDAT.LT.ITMLFT)IUPDAT=ITMLFT
      IF((IUPDAT - ITMLFT).GT.300000)CALL UPDATE
      PRINT190,ITMLFT
  190 FORMAT(10H TIME LEFT,I8//)
   40 CONTINUE
      CALL UPDATE
      DO60I=1,625
   60 REC(I,2)=REC(I,1)
      GOTO100
C
C  OUTPUT AND STOP
  300 DO 310I1=1,5,IGSTEP
      DO310I2=1,5,IGSTEP
      DO310I3=1,5,IGSTEP
      DO310I4=1,I2,IGSTEP
      APROD=VOPT(I1,I2,I3,I4,2)
      IF(APROD.LT..99)GOTO320
  310 CONTINUE
      PRINT315
  315 FORMAT(25H NO MULTILINE POINT FOUND//)
      GOTO350
  320 PRINT325
  325 FORMAT(36H AT LEAST ONE MULTILINE POINT EXISTS/
     1 20H INITIAL CONDITIONS;/ 11H   STORAGE ,9X,2HAA,11X,
     2 4HAABB,12X,2HBB)
      PRINT 326,RANGE(1,I1),RANGE(2,I2),RANGE(3,I3),RANGE(4,I4)
  326 FORMAT(1H ,4E14.8)
      PRINT327
  327 FORMAT (8H POLICY;/1H  ,6X,1HA,13X,2HAB,13X,1HB,
     1 12X,2HAA,11X,4HAABB,11X,2HBB,8X,11HDAYS WORKED)
      PRINT328(VOPT(I1,I2,I3,I4,I),I=1,7)
  328 FORMAT(1H ,7E14.8)
      PRINT329
  329 FORMAT(//8H OUTPUT;/8H  COST,12X,5HYIELD,9X,7HSTORAGE,
     1 10X,2HAA,11X,4HAABB,11X,2HBB)
```

```
      177              PRINT330(VOPT(I1,I2,I3,I4,I),I=11,16)
      178          330 FORMAT(1H 6E14.8)
      179              PRINT 332,VOPT(I1,I2,I3,I4,18)
      180          332 FORMAT(20H VARIANCE OF COST IS,E14.8)
      181          350 STOP
      182              SUBROUTINE UPDATE
      183        C
      184              REWIND 24
      185              WRITE(24),VOPT
      186              END FILE 24
      187              REWIND 24
      188              ITMLFT=JOBSUP(1)
      189              PRINT949,ITMLFT,IGSITR
      190          949 FORMAT(//13H FILE UPDATED,I10,28H MILLISECONDS FROM TIMELIMIT
      191             1 ,17H ON G-S ITERATION,I6)
      192              IF(ISUCCS)942,943,946
      193          942 PRINT 941,IGSITR
      194          941 FORMAT(26H DEFAULT STOP ON ITERATION,I6)
      195              RETURN
      196          943 PRINT945
      197          945 FORMAT(20H RUNSTREAM DIVERSION)
      198              IUPDAT=ITMLFT
      199              RETURN
      200          946 PRINT 944
      201          944 FORMAT(16H SUCCESSFUL EXIT)
      202              RETURN
      203              END

@PRT,S OPTIPLANT.DAVMIN
```

```
OKM801*OPTIPLANT(1).DAVMIN
      1         C     ************************************************************************
      2         C     *  THIS SUBROUTINE OPTIMISES THE COST OF A POLICY, WHICH COST          *
      3         C     *  IT GETS FROM SUBROUTINE "FUNCOS", BY USE OF THE METHOD OF           *
      4         C     *  DAVIDON (1975), WITH DIRECT GRADIENT ESTIMATES AFTER BIRTA          *
      5         C     *  1976.  THE SEARCH FOR A MINIMUM IS CONSTRAINED TO A SUBSPACE        *
      6         C     *  OF CONTROL POLICIES BY THE CHOICE OF ** DAVJST** .  WHEN THE        *
      7         C     *  SEARCH ENCOUNTERS AN ACTIVE CONSTRAINT, THE GRADIENT AT THAT        *
      8         C     *  POINT IS PROJECTED INTO THE CONSTRAINT, WITHOUT FURTHER             *
      9         C     *  COLLAPSING THE SUBSPACE.  IF  IENAB = 1 , ALTERNATE                *
     10         C     *  ITERATIONS MINIMISE AND MAXIMISE THE COST WITH RESPECT TO           *
     11         C     *  THE CURRENT SET OF CONTROL VARIABLES DEFINED BY  NUA  AND           *
     12         C     *  NUB.  OPTIMISATION CEASES WHEN THE ITERATION NUMBER EXCEEDS         *
     13         C     *  ITLIM  OR WHEN ONE OF THE OTHER CRITERIA DESCRIBED IN THE           *
     14         C     *  DOCUMENTATION IS MET.                                              *
     15         C     *                                                                     *
     16         C     ************************************************************************
     17         C
     18         C
     19               SUBROUTINE DAVMIN(IVPOS,DAVJST,IEPARA,EPARAM,F,RANGE,SLIM)
     20         C
     21         C
     22               DIMENSION IVPOS(4),F(3,3),RANGE(4,5),XSTAR(5)
     23               DIMENSION DELTA(10),IEPARA(25),EPARAM(25)
     24               DIMENSION IPEC(10,2),IDREC(2)
     25               DOUBLE PRECISION X(4,2),U(10),UTRY(10),UDASH(10),EGRAD(10)
     26               DOUBLE PRECISION E2GRAD(1),EG(10),DAVM(10),ESS(10),P(10)
     27               DOUBLE PRECISION Q(10),WUB(10),TEM(10),DAVJST(10,10),SLIM(10,4)
     28               DOUBLE PRECISION ALMIN,AMOD,RAT,STESCA,ABBA,COSTX,COST
     29               DOUBLE PRECISION ENVAL,EXCOS,FG,FO,FODASH,F2DASH,GHEK,YIELD
     30               DOUBLE PRECISION ALF,AYE,B,BO,CEE,DELT,EMEW,EMSQ,EMWUB
     31               DOUBLE PRECISION ENSQ,GAM,UGH,USQ,UTS,VEE,TEMP,VARRUS
     32               DOUBLE PRECISION DAVREF(10,10)
     33         C
     34               COMMON VOPT(5,5,5,5,20)
     35         C
     36               REAL LAMBDA
     37         C
     38               DEFINE DIREC(I,YIELD)=YIELD*(1-2*I)
     39         C
     40               IENAB=IEPARA(9)
     41               RHO=EPARAM(14)
     42               ISAD=0
     43               ISADPR=1
     44         C
     45               IPRINT=IEPARA(21)
     46               ITLIM=IEPARA(22)
     47               EXCOS=0
     48               YIELD=0
     49               VARRUS =0
     50               I1=IVPOS(1)
     51               I2=IVPOS(2)
     52               I3=IVPOS(3)
     53               I4=IVPOS(4)
     54               X(1,1)=RANGE(1,I1)
     55               X(2,1)=RANGE(2,I2)
     56               X(3,1)=RANGE(3,I3)
     57               X(4,1)=RANGE(4,I4)
     58               DO3 I=1,10
     65
```

```
      U(I)=VOPT(I1,I2,I3,I4,I)
      DELTA(I)=1
      DO 3 J = 1,10
    3 DAVREF(I,J) = DAVJST(I,J)
      ENVAL = 0.
C
  900 CONTINUE
      CALL FUNCOS(X,U,YIELD,RANGE,F,EXCOS,IEPARA,EPARAM,VARRUS,
     1 COSTX,COST)
C
      FO=EXCOS
      IF(ABS(FO-ENVAL).LT.FO*1E-6)GOTO901
      ENVAL=FO
      VOPT(I1,I2,I3,I4,11) = FO
      VOPT(I1,I4,I3,I2,11) = FO
      GO TO 900
  901 ENVAL = FO
      DO 102 I = 1,4
  102 XSTAR(I) = X(I,2)
      XSTAR(5) = YIELD
      VARTEM = VARRUS
      CSTXTM = COSTX
      COSTEM = COST
      VOPT(I1,I2,I3,I4,19) =  COSTX
      VOPT(I1,I2,I3,I4,20) = COST
      DO 2 I = 1,4
      I9 = I+12
    2 VOPT(I1,I2,I3,I4,I9)=X(I,2)
      VOPT(I1,I2,I3,I4,12)=YIELD
      IF(IPRINT.LT.1)GOTO800
      ITER = 0
      PRINT 600,FO,YIELD,I1,I2,I3,I4
  600 FORMAT(//11H START: FO=,D15.8,9H YIELD IS,D15.8,9H POINT IS,4I5/)
      PRINT 601,U
  601 FORMAT(17H INITIAL DECISION//1H ,10D13.6/)
      PRINT 618, VARRUS,COSTX,COST
  618 FORMAT(10H VAR(E(V)),D15.8,5H K(I),D15.8,6H V(X1),D15.8/)
  800 CONTINUE
      IRENAC = 0
      EPS = 1
C
C STEP 0    EVALUATE EUCLIDEAN GRADIENT AT START POINT
      IBLOCK=0
    5 CONTINUE
      NUA=IEPARA(ISAD+1)
      NUB=IEPARA(ISAD+3)
      NUC=IEPARA(ISAD+5)
      NUD=IEPARA(ISAD+7)
      DO10I=NUA,NUB
      STESCA=0
      AMOD=0
      DO12J=1,10
      TEM(J)=DELTA(J)*DAVJST(J,I)
      RAT=ABS(TEM(J)/SLIM(J,2))
      IF(RAT.LE.STESCA)GOTO12
      STESCA=RAT
   12 CONTINUE
      DO14J=1,10
      TEM(J)=TEM(J)*1E-4/STESCA
```

```
      AMOD=AMOD+TEM(J)**2
   14 UDASH(J)=U(J)+TEM(J)
      AMOD=DSQRT(AMOD)
C
   19 CALL FUNCOS(X,UDASH,YIELD,RANGE,F,EXCOS,IEPARA,EPARAM,VARRUS,
     1 COSTX,COST)
C
      FG=EXCOS
      GHEK=FG-ENVAL
      IF(IPRINT.LT.3)GOTO801
      PRINT602,I,GHEK,AMOD
  602 FORMAT(5H GROP,I6,2D20.10)
  801 CONTINUE
   15 EG(I)=DIREC(ISAD,GHEK)/AMOD
   10 CONTINUE
      IF(ISLOCK.EQ.1)GOTO89
   22 DO11 I=NUA,NUB
      EGRAD(I)=EG(I)
      WUB(I)=EGRAD(I)
   11 CONTINUE
C
C STEP ONE              BEGINNING OF ITERATION
   20 FODASH=0
      IF(IPRINT.LT.3)GOTO802
      PRINT603,EGRAD
  603 FORMAT(6H EGRAD,10D12.5)
  802 CONTINUE
      ITER=ITER+1
      IF(ITER.GT.ITLIM)GOTO300
      IF(IPRINT.LT.2)GOTO803
      IF(ISAD.EQ.0)PRINT604,ITER,ITLIM,FO
  604 FORMAT(//27H STEP ONE: ITERATION NUMBER,I6,3H OF,I6,
     1 8H COST IS,D15.6,15H AND MINIMISING)
      IF(ISAD.EQ.1)PRINT605,ITER,ITLIM,FO
  605 FORMAT(27H STEP ONE: ITERATION NUMBER,I8,3H OF,I8,
     1 8H COST IS,D15.6,15H AND MAXIMISING)
  803 CONTINUE
      DO25I=NUA,NUB
      ESS(I)=-EGRAD(I)
      FODASH=FODASH+EGRAD(I)*ESS(I)
   25 CONTINUE
      IF(IPRINT.LT.3)GOTO804
      PRINT606,FODASH
  606 FORMAT(7H FODASH,D15.7)
  804 CONTINUE
      LAMBDA=2
C
C STEP TWO              TAKE STEP AND (PERHAPS) STOP
   30 CONTINUE
      IDIV=0
      DO32I=1,10
      TEM(I)=0
      DO34J=NUA,NUB
   34 TEM(I)=TEM(I)+DAVJST(I,J)*ESS(J)
   32 CONTINUE
      IPROJ=0
   38 IFLAG=0
      IMOVE=0
      RAT=0.
```

```
      STESCA=0.
      DO 36 I =1,10
      RAT=ABS(TEM(I)/SLIM(I,2))
      IF(RAT.LE.STESCA)GOTO36
      STESCA=RAT
   36 CONTINUE
      STESCA=STESCA*20+1
      IF(IPRINT.LT.3)GOTO805
      PRINT 607,STESCA,TEM,(ESS(I),I=NUA,NUB)
  607 FORMAT(7H STESCA,D12.5/4H TEM,10D12.5/4H ESS,10D12.5)
  805 CONTINUE
      DO42I=1,10
      UTRY(I)=U(I)+TEM(I)/STESCA
C  NOW LOOK  FOR TRANSGRESSIONS.  FIRST ELIMINATE THOSE THAT  HAVENT MOVED
      DO42J=3,4
   42 SLIM(I,J)=0.
C  ASSUMES NO CONSTRAINTS ACTIVE
      DO44I9=1,10
      DO46J=1,2
   46 IREC(I9,J)=1
      ABBA= DABS(U(I9)-UTRY(I9))
      IF(ABBA.LT.SLIM(I9,2)*1E-20)GOTO49
      IMOVE=1
      DO48J=1,2
      J1=J+2
      ALMIN=(SLIM(I9,J)-UTRY(I9))/(U(I9)-UTRY(I9))
      IF(.99999999.LE.ALMIN.AND.ALMIN.LE.1)GOTO50
      IF(0.GT.ALMIN.OR.ALMIN.GT.1)GOTO48
C LEFT ARE ORDINARY INTERPOLATORS
      SLIM(I9,J1)=ALMIN
      GOTO48
   50 IF(J.EQ.1.AND.UTRY(I9).GE.SLIM(I9,1))GOTO48
      IF(J.EQ.2.AND.UTRY(I9).LE.SLIM(I9,2))GOTO48
      IREC(I9,J)=0
   48 CONTINUE
      GO TO 44
   49 UTRY(I9) = U(I9)
   44 CONTINUE
      IF(ISAD.EQ.1)GOTO69
      NUE=4
C  NOW PROJECT ACTIVE PLANT CONSTRAINTS
      ICNT=0
      DO52J=1,3
      IF(IREC(J,1).EQ.0)ICNT=ICNT+1
   52 CONTINUE
      IF(ICNT-1)54,56,58
   56 IFLAG=1
      IF(IPROJ.EQ.1)GOTO58
      DO60J=1,3
      IF(IREC(J,1))60,60,62
   62 IREC(J,1)=-1
      GOTO57
   60 CONTINUE
   57 TEMP=0
      DO64J=1,3
   64 TEMP=TEMP+TEM(J)*IREC(J,1)
      DO66J=1,3
   66 TEM(J)=TEMP*IREC(J,1)
      IPROJ=IPROJ+1
```

```
      GOTO54
C  RECYCLES THE PROJECTION
   58 IFLAG=1
      DO67J=1,3
   67 TEM(J)=0
      GOTO54
   69 NUE=NUC
   54 DO68J=NUE,NUD
      DO68J1=1,2
      IF(IREC(J,J1))55,55,68
   55 IFLAG=1
      TEM(J)=0
   68 CONTINUE
      IF(IFLAG.EQ.1)GOTO38
C NOW SELECT THE LARGEST ALMIN (NEAREST TO PREVIOUS POINT)
      DO37I=NUA,NUB
   37 ESS(I)=ESS(I)/STESCA
      FODASH=FODASH/STESCA
      ALMIN=0
      DO70I=1,10
      DO70J=3,4
      IF(SLIM(I,J).LE.ALMIN)GOTO70
      ALMIN=SLIM(I,J)
   70 CONTINUE
      IF(IPRINT.LT.3)GOTO806
      PRINT608,ALMIN
  608 FORMAT(6H ALMIN,D13.6)
  806 CONTINUE
      DO72I=1,10
      UTRY(I)=ALMIN*U(I)+(1.-ALMIN)*UTRY(I)
      DO72J=1,2
      ABBA = DABS(UTRY(I)-SLIM(I,J))
      IF(ABBA.LT.SLIM(I,2)*1E-5)UTRY(I)=SLIM(I,J)
   72 CONTINUE
      DO 73 I=4,10
      IF(UTRY(I).LT.SLIM(I,1))UTRY(I)=SLIM(I,1)
   73 IF(UTRY(I).GT.SLIM(I,2))UTRY(I)=SLIM(I,2)
C NOW DECIDE WHETHER TO CONTINUE
      IF(IMOVE)85,85,74
   74 IF(-FODASH.LT.EPS)GOTO85
C 300 IS STOP AND RETURN
      IF(IPRINT.LT.3)GOTO810
      PRINT612,UTRY
  612 FORMAT(8H PRESTEP,10D12.5)
  810 CONTINUE
C
   76 CALL FUNCOS(X,UTRY,YIELD,RANGE,F,EXCOS,IEPARA,EPARAM,VARRUS,
     1 COSTX,COST)
C
      ENVAL=EXCOS
      IF(IPRINT.LT.2)GOTO807
      PRINT609,FO,ENVAL
  609 FORMAT(8H OLD,NEW,2D22.12)
  807 CONTINUE
      IF(DIREC(ISAD,FO-ENVAL).GT.FO*1E-6)GOTO78
      DO80I=1,10
      UTRY(I) = (U(I) + UTRY(I))/2
      IF(UTRY(I).LT.SLIM(I,1))UTRY(I)=SLIM(I,1)
   80 IF(UTRY(I).GT.SLIM(I,2))UTRY(I)=SLIM(I,2)
```

```
      IDIV=IDIV+1
      IF(IDIV.LE.5)GOTO82
      IF(IPRINT.LT.3)GOTO811
      PRINT613
  613 FORMAT(22H INFINITE SUBDIVISIONS)
  811 CONTINUE
   85 IF(IENAB.EQ.1)GOTO334
      IF(IMOVE.EQ.0.OR.IRENAC.EQ.0)GOTO300
      ENVAL = F0
      IRENAC = 0
      IBLOCK = 0
      GOTO 5
  334 CONTINUE
      IF(ISADPR.EQ.2)GOTO300
      ISADPR=ISADPR+1
      IDREC(ISAD+1)=0
      GOTO275
   82 CONTINUE
      FODASH=FODASH/2
      LAMBDA=.5
      GOTO76
C GO BACK AND TAKE A SHORTER STEP
   78 CONTINUE
      ISADPR=1
      IDREC(ISAD+1)=1
C AN IMPROVEMENT HAS BEEN FOUND
C
C STEP THREE    UPDATES
      DO86I=1,10
   86 U(I)=UTRY(I)
      F2DASH=0
      IBLOCK=1
      GOTO5
   89 CONTINUE
      DO90I=NUA,NUB
C
      E2GRAD(I)=EG(I)
      F2DASH=F2DASH+E2GRAD(I)*ESS(I)
   90 CONTINUE
      DO100I=NUA,NUB
      DAVM(I)=ESS(I)+EGRAD(I)-E2GRAD(I)
  100 EGRAD(I)=E2GRAD(I)
      B0=F2DASH-FODASH
C NEW VALUES
      F0=ENVAL
      VARTEM = VARRUS
      CSTXTM = COSTX
      COSTEM = COST
      DO 101 I = 1,4
  101 XSTAR(I)=X(I,2)
      XSTAR(5) = YIELD
C
      FODASH=F2DASH
      IF(B0-EPS)105,120,120
  105 IRENAC = 1
      DO107I=NUA,NUB
  107 ESS(I)=ESS(I)*LAMBDA
      FODASH=FODASH*LAMBDA
      GOTO30
```

```
C RETURN TO STEP TWO
  120 CONTINUE
      EMSQ=0
      EMWUB=0
      VEE=0
      DO125I=NUA,NUB
      EMSQ=EMSQ+DAVM(I)**2
      EMWUB=EMWUB+DAVM(I)*WUB(I)
  125 VEE=VEE+DAVM(I)*ESS(I)
      IF(EMSQ-EPS)20,130,130
  130 EMEW=VEE-EMSQ
      DO135I=NUA,NUB
  135 WUB(I)=WUB(I)-DAVM(I)*EMWUB/EMSQ
      UGH=0
      USQ=0
      UTS=0
      DO140I=NUA,NUB
      UGH=UGH+DAVM(I)*WUB(I)
      UTS=UTS+WUB(I)*ESS(I)
  140 USQ=USQ+WUB(I)*WUB(I)
      IF(UGH*UGH*1E6-EMSQ*USQ)160,145,145
  145 DO147I=NUA,NUB
  147 WUB(I)=0
      ENSQ=0
      GOTO180
C STEP 4A TYPE OF UPDATE
  160 DO165I=NUA,NUB
  165 WUB(I)=WUB(I)*UTS/USQ
      ENSQ=UTS*UTS/USQ
C STEP FIVE
  180 B=ENSQ+EMEW*VEE/EMSQ
      IF(B-EPS)185,200,200
  185 DO190I=NUA,NUB
  190 WUB(I)=ESS(I)-DAVM(I)*VEE/EMSQ
      ENSQ=BQ-EMEW*VEE/EMSQ
      B=BQ
C
C STEP 6
  200 CONTINUE
      IF(EMEW*VEE-EMSQ*ENSQ)210,205,205
  205 GAM=0
      DELT=SQRT(VEE/EMEW)
      GOTO220
C STEP 6A
  210 AYE=B-EMEW
      CEE=B*VEE
      GAM=SQRT((1-EMEW*VEE/EMSQ/ENSQ)/AYE/B)
      DELT=SQRT(CEE/AYE)
      IF(CEE-AYE)215,220,220
  215 GAM=-GAM
C
C STEP 7
  220 ALF=VEE+EMEW*DELT+EMSQ*ENSQ*GAM
      TEMP=0
      DO225I=NUA,NUB
      P(I)=DAVM(I)*(DELT-ENSQ*GAM)+WUB(I)*GAM*VEE
      Q(I)=DAVM(I)*(1+ENSQ*GAM)/ALF-WUB(I)*GAM*EMEW/ALF
      WUB(I)=DAVM(I)*ENSQ*(1+GAM*EMEW*VEE/ALF)-
     1 WUB(I)*(1+DELT)*EMEW*VEE/ALF
```

```
      225 TEMP=TEMP+Q(I)*EGRAD(I)
          DO230I=NUA,NUB
      230 EGRAD(I)=EGRAD(I)+P(I)*TEMP
        , DO240I=1,10
          TEM(I)=0
          DO245J=NUA,NUB
      245 TEM(I)=TEM(I)+DAVJST(I,J)*Q(J)
      240 CONTINUE
          DO250I=1,10
          DO250J=NUA,NUB
      250 DAVJST(I,J)=DAVJST(I,J)+TEM(I)*P(J)
          IRENAC = 1
          IF(IPRINT.LT.3)GOTO808
          PRINT610,((DAVJST(I,J),I=1,10),J=NUA,NUB)
      610 FORMAT(1H ,10D12.4)
      808 CONTINUE
    C FINAL TESTS AFTER UPDATING
          IF(ENSQ)255,255,275
      255 DO260I=NUA,NUB
      260 WUB(I)=EGRAD(I)
      275 ITEM = 0
          IF(IENAB.EQ.1.AND.(ITER.EQ.1.OR.IDREC(ISAD+1).EQ.1))ITEM=1
          IF(IENAB.EQ.1)ISAD=ISAD+1
          IF(ISAD.EQ.2)ISAD=0
          NUA=IEPARA(ISAD+1)
          NUB=IEPARA(ISAD+3)
          NUC=IEPARA(ISAD+5)
          NUD=IEPARA(ISAD+7)
          IF(ITEM.EQ.0)GOTO20
          IBLOCK=0
          GOTO5
    C NOW FOR STOPPING AND RETURNING
      300 CONTINUE
          DO310I=1,10
      310 VOPT(I1,I2,I3,I4,I)=U(I)
          DO312I=1,4
          I9=I+12
      312 VOPT(I1,I2,I3,I4,I9)=XSTAR(I)
          VOPT(I1,I2,I3,I4,12)=XSTAR(5)
          VOPT(I1,I2,I3,I4,11) = FO
          VOPT(I1,I2,I3,I4,18) = VARTEM
          VOPT(I1,I2,I3,I4,19) = CSTXTM
          VOPT(I1,I2,I3,I4,20) = COSTEM
          IF(IPRINT.LT.1)GOTO809
          PRINT611,ITER,(VOPT(I1,I2,I3,I4,I),I=1,11),(VOPT(I1,I2,I3,I4,J),
         1 J=18,20)
      611 FORMAT(19H CEASE ON ITERATION,I6/
         1 13H FINAL CHOICE,/1H ,10E13.5/5H COST,E14.7,
         2 20H VARIANCE OF COST IS,E14.7,7H   K(I),E14.7,8H   V(X1),E14.7)
          PRINT 619,XSTAR
      619 FORMAT(23H EXPECTED: STATE OUTPUT,4E14.7,6H YIELD,E14.7)
      809 CONTINUE
          RETURN
          END

@PRT,S OPTIPLANT.FUNCOS
```

```
DKM801*OPTIPLANT(1).FUNCOS
      1       C    ************************************************************************
      2       C    *                                                                      *
      3       C    *  THIS SUBROUTINE TAKES IN A STATE VALUE,  X , A CONTROL ACTION       *
      4       C    *  U , AND THE CURRENT ARRAY OF COSTS (IN  VOPT ), TOGETHER WITH       *
      5       C    *  THE GIVEN MODEL AND CRITERION TYPES, AND PROBABILITIES OF           *
      6       C    *  MUTATION AND IRRUPTION OF DISEASES.  IT THEN SETS UP THE EIGHT      *
      7       C    *  POSSIBLE IRRUPTIONS USING  SPEXT , AND TRANSMITS THEM TO            *
      8       C    *  "EPIMUL", WHICH RETURNS EACH POSSIBLE YIELD AND STATE OUTPUT.       *
      9       C    *  THE OUTPUTS ARE THEN TRANSMITTED TO  "COST", WHICH RETURNS THE      *
     10       C    *  COST OF EACH EPIDEMIC AND THE RESULTING STORAGE.  THESE ARE THEN    *
     11       C    *  USED BY  "FUNCOS" TO CALCULATE THE EXPECTED COST AND STATE          *
     12       C    *  OUTPUT, AND THE COMPONENTS OF THE COST AND ITS VARIANCE.            *
     13       C    *                                                                      *
     14       C    ************************************************************************
     15       C
     16       C
     17             SUBROUTINE FUNCOS(X,U,YIELD,RANGE,F,EXCOS,IEPARA,EPARAM,VARRUS,
     18            1 COSTX,COSTM)
     19       C
     20       C
     21             DIMENSION F(3,3)
     22             DIMENSION EPARAM(25),MUTMAT(3,3)
     23             DIMENSION RANGE(4,5),IEPARA(25)
     24             DOUBLE PRECISION X(4,2),U(10),UBAS(6),XSTAR(4),P(2),SPAD(3)
     25             DOUBLE PRECISION ECU(4),COSMIN,YIELD,PROB
     26             DOUBLE PRECISION VARRUS,EXCOS,PEXT,REALMU,COSTX,COSTM,PRES,FUT
     27       C
     28             COMMON VOPT(5,5,5,5,20)
     29       C
     30             INTEGER VIR(3,3)
     31             DATA((MUTMAT(I,J),J=1,3),I=1,3)/0,1,0,1,0,1,0,1,0/
     32             DATA((VIR(I,J),J=1,3),I=1,3)/1,0,0,1,1,1,0,0,1/
     33       C
     34       C
     35             PEXT=U(9)
     36             REALMU=10.**(-U(10))
     37             RHO = EPARAM(14)
     38             YIELD=0
     39             VARRUS =0
     40             PRES = 0.
     41             FUT = 0.
     42             DO 162 I = 1,4
     43         162 XSTAR(I)= 0.
     44       C
     45             DO75I=1,6
     46          75 UBAS(I)=U(I)
     47             DO76I=1,3
     48             I1=I+6
     49          76 ECU(I)=U(I1)
     50             ECU(4)=U(4)+U(5)+U(6)
     51             COSMIN = 0
     52             P(1) = PEXT
     53             P(2) = 1 - PEXT
     54             SPEXT = EPARAM(15)
     55             DO 120 I1 = 1,2
     56             SPAD(1) = SPEXT*(I1-1)
     57             DO 120 I2 = 1,2
     58             SPAD(2) = SPEXT*(I2-1)
     65
```

```
      DO 120 I3 = 1,2
      PROB = P(I1)*P(I2)*P(I3)
      SPAD(3) = SPEXT*(I3-1)
C
  100 CALL EPIMUL(X,UBAS,F,VIR,SPAD,REALMU,IEPARA,EPARAM,MUTMAT)
C
      YIELD=YIELD+X(1,2)*PROB
      CALL COST(X,ECU,EXCOS,RANGE,IEPARA,EPARAM,COSTX,COSTM)
C
      DO 160I7=1,4
  160 XSTAR(I7)=XSTAR(I7)+X(I7,2)*PROB
      VARRUS = VARRUS + PROB*EXCOS**2
      COSMIN = COSMIN + EXCOS*PROB
      PRES = PRES + PROB*COSTX
      FUT = FUT + PROB*COSTM
  120 CONTINUE
C
      DO 161 I7= 1,4
  161 X(I7,2)= XSTAR(I7)
      EXCOS = COSMIN
      VARRUS = VARRUS - (EXCOS)**2
      COSTX = PRES
      COSTM = FUT
      RETURN
      END

@PRT,S OPTIPLANT.EPIMUL
```

```
DKM801*OPTIPLANT(I).EPIMUL
      C    ************************************************************************
      C    *                                                                      *
      C    *  THIS SUBROUTINE TAKES IN PARAMETER VALUES RELATING                  *
      C    *  TO THE LENGTH AND SEVERITY OF THE EPIDEMIC AND THE SPORE INPUT      *
      C    *  GIVEN BY THE STATE VARIABLE  X , AND RETURNS A YIELD CALCULATED     *
      C    *  BY THE RELATIONSHIP OF ROMIG AND CALPOUZOS (1970) FROM THE          *
      C    *  DISEASE COVER ESTIMATED BY SIMULATING EITHER THE MULTIVARIABLE      *
      C    *  FORM OF THE EQUATION OF VAN DER PLANK (1975), OR A MORE REFINED     *
      C    *  MODEL INCORPORATING PLANT GROWTH EFFECTS, WITH OR WITHOUT           *
      C    *  ADDED CROSSPROTECTION.                                              *
      C    *                                                                      *
      C    ************************************************************************
      C
            SUBROUTINE EPIMUL(X,UBAS,F,VIR,SPAD,REALMU,IEPARA,EPARAM,MUTMAT)
      C
      C
            DIMENSION EPARAM(25),MUTMAT(3,3)
            DIMENSION F(3,3),VIGORS(3)
            DIMENSION IEPARA(25)
            INTEGER VIR(3,3)
            DOUBLE PRECISION Y(4,3,70),S(3),UBAS(6),X(4,2),ALOSS(3),TEM,TEM1
            DOUBLE PRECISION SPOSUB(3),PROP(3),YHOLD(3),SPAD(3),REALMU
            DOUBLE PRECISION PTOT,SPORAV,OCCUP,FRAC,FRACM,Q,DELY,SPLAT,DENPRO
            DOUBLE PRECISION DELYT,YG,WINTOT
            COMMON VOPT(5,5,5,5,20)
      C
      C UNPACK PARAMETERS FROM EPARAM
            N=IEPARA(10)
            N1=N+1
            IMOD=IEPARA(11)
            ISTART=IEPARA(13)
            LATEP=IEPARA(15)
            INFECP=IEPARA(16)
            ICPTIM=IEPARA(17)
            VIGORS(1)=EPARAM(1)
            VIGORS(2)=EPARAM(2)
            VIGORS(3)=EPARAM(3)
            GAIN=EPARAM(22)
            YPAM1=EPARAM(6)
            FACT = EPARAM(23)
            YPAM2=EPARAM(7)
            IYCRIT=IEPARA(18)
      C
      C
      C FILLUP FIRST LOT OF S AND Y AND ADD MUTANT SPORES
      C
            DO180I=1,3
            S(I)=0
            ALOSS(I)=0
            SPOSUB(I)=0
        180 CONTINUE
            DO 7 I =1,4
            X(I,2)=0.
            DO 7 J = 1,3
            DO7K=1,N1
            Y(I,J,K)=0
          7 CONTINUE
            GOTO(4,5),IMOD
```

```
      4 DO9J=1,3
        YHOLD(J)=UBAS(J)
        DO9K=1,N1
      9 Y(1,J,K)=UBAS(J)
        DELYT=1.
        SPLAT = 1.
        GOTO10
C *********************  5 WOULD GIVE COTYLEDONS TO THE REFINED MODEL
      5 SEED=EPARAM(16)
        GROW=EPARAM(17)
        FRACK=EPARAM(18)
        DO 200 J = 1,3
  200 Y(1,J,I)=SEED*UBAS(J)
C ************************
C START CROP CYCLE
     10 DO 20I1=1,N
        I2=I1+1
        IF (II.NE.ISTART)GOTO14
        DO 15 I = 1,3
     15 S(I) = X(I+1,1) + UBAS(I+3) + SPAD(I)
     14 CONTINUE
        PTOT=S(1)+S(2)+S(3)
        IF(PTOT.LT.1E-10)GOTO26
        IDEP=1
        GOTO27
     26 IDEP=0
     27 IF(IDEP.EQ.0.AND.IMOD.EQ.1)GOTO12
        IF(IDEP.EQ.0.AND.IMOD.EQ.2)GOTO40
C NO DEPOSITION IF NO SPORES
        DO25I=1,3
     25 PROP(I)=S(I)/PTOT
C DENUMERATOR OF HIT PROBABILITY
        GOTO(35,40),IMOD
     35 DENPRO =EPARAM(8)
        GOTO45
C *************** REFINED MODEL REFERS TO LAI HERE
     40 SPLAT=0
        DO 210 J = 1,3
        SPLAT = SPLAT + ALOSS(J)
        DO 210 I = 1,4
    210 SPLAT = SPLAT + Y(I,J,I1)
        DENPRO = (EPARAM(8)-1.+SPLAT)/SPLAT
        IF(IDEP.EQ.1)GOTO45
        DO 211 J = 1,3
        Y(1,J,I2)=Y(1,J,I1)
    211 CONTINUE
        GO TO 12
C*********************
     45 DO 50J=1,3
C  J INDEXES PLANTS.  DEPOSITION OF SUCCESSIVE SPORES
        SPORAV=0
        DO55K=1,3
     55 SPORAV=SPORAV+PROP(K)*F(K,J)
        OCCUP=SPORAV*PTOT/DENPRO
        FRAC=DEXP(-OCCUP)
        FRACM=1-FRAC
        DO60I=2,4
        IF(SPORAV.LT.1E-35)GOTO65
        I3 = I-1
```

```
      Y(I,J,I2)=Y(1,J,I1)*F(I3,J)/SPORAV*PROP(I3)*FRACM
      GO TO 60
   65 Y(I,J,I2)=0.
   60 CONTINUE
      Y(1,J,I2)=Y(1,J,I1)*FRAC
C Y(1, , ) IS UNOCCUPIED AREA.  NOW CHECK FOR DEFUNCT C.P.
      I3=I1-ISTART-ICPTIM
      IF(I3)50,70,70
   70 I3=I3+ISTART
      Q=0
      DO75I=1,3
      TEM = VIR(I,J) - 1
      Q=Q+Y(I+1,J,I3)*TEM
   75 CONTINUE
      Y(1,J,I2)=Y(1,J,I2)+Q
   50 CONTINUE
C NOW GROW AND PRODUCE SPORES
   12 DO 61 I = 1,3
   61 S(I) = 0.
      I3=I1-LATEP
      I4=I3-INFECP
      IF(I3-ISTART)85,85,105
  105 DO90I=1,3
      I5=I+1
      DO95J=1,3
      TEM = VIR(I,J)
   97 SPOSUB(I)=SPOSUB(I)+(Y(I5,J,I3)-Y(I5,J,I4))*TEM
   95 CONTINUE
      S(I)= SPOSUB(I)*VIGORS(I)
   90 CONTINUE
      DO 99 I = 1,3
      PROP(I) = 0.
      DO 99 J = 1,3
      PROP(I) = PROP(I) + REALMU*S(J)*MUTMAT(I,J)
   99 CONTINUE
      DO 101 I = 1,3
      S(I) = S(I) + PROP(I)
      IF(S(I).LT.1E-10)S(I)=0
  101 CONTINUE
C SPORES EQUAL SPORES OF PREVIOUS PERIOD + NEW PRODUCERS - THE AGED
   85 GO TO(110,86),IMOD
C ************************** 86 IS THE START OF REFINED GROWTH
   86 TEM = I1
      TEM1 = IYCRIT
      TEM1 = 1. -TEM/TEM1
      DO 220 J = 1,3
      YG = Y(1,J,I2)
      DO 224 I = 2,4
      I9 = I-1
  224 YG = YG + Y(I,J,I2)*(1-VIR(I9,J))
      TEM = 0.
      DO 225 I = 1,4
  225 TEM= TEM + Y(I,J,I1)
      YHOLD(J) = TEM
C  NOW CALCULATE FEEDBACK ELEMENT OF GROWTH
C
      DELYT = (YHOLD(J)/(UBAS(J)+1E-25)/SPLAT)
C
C NOW THE MAIN GROWTH FUNCTION
```

```
      DELY = GROW*TEM1*YG*DELYT**FBACK
  230 Y(1,J,I2)=Y(1,J,I2)+DELY
  220 CONTINUE
C *************************
C YIELD CALCULATIONS
  110 IF(II.GT.IYCRIT)GOTO125
      DO130J=1,3
      Q=0
      DO135I=2,4
      TEM = VIR(I-1,J)
      Q=Q+Y(I,J,I1)*TEM
  135 CONTINUE
      ALOSS(J)=ALOSS(J)+Q
      IF(II.LT.IYCRIT)GOTO130
      TEM = ALOSS(J)
      Q= (TEM*100/(YHOLD(J)+1E-25)/.37)
      IF(Q.LT.1E-10)Q=1.
      Q=DLOG(Q)
      TEM=YPAM1+YPAM2*Q
      IF(TEM.LT.0)TEM=0
      IF(TEM.GT.100)TEM=100
      DELYT=YHOLD(J)/(UBAS(J)+1E-25)/SPLAT
      TEM1 = 400*(1.-TEM/100)*UBAS(J)*DELYT**GAIN
  131 X(1,2)=X(1,2)+TEM1
  130 CONTINUE
C NOW FILL UP STATE VARIABLES AFTER CLOSING SEASON
  125 CONTINUE
   20 CONTINUE
      WINTOT = 0.
      DO 139 I = 1,3
      IF(I.EQ.2)SFACT=FACT
      IF(I.NE.2)SFACT=1.
  139 WINTOT = WINTOT + S(I)*SFACT
      DO 140I=2,4
      IF(I.EQ.3)SFACT=FACT
      IF(I.NE.3)SFACT=1.
  140 X(I,2)= SFACT*S(I-1)*1000./(WINTOT + 500.)
C
C
      RETURN
      END

$PRT,S OPTIPLANT.COST
```

```
DKM80I*OPTIPLANT(1).COST
      1       C  ************************************************************************
      2       C  *                                                                      *
      3       C  * THIS SUBROUTINE TAKES IN  X , AND ECONOMIC CONTROL VARIABLES        *
      4       C  * STORED IN  ECU , AND SOME PARAMETERS AND RETURNS THE COST  EXCOS , *
      5       C  * AND ITS COMPONENTS  COSTX  AND  COST.  COSTX IS CALCULATED FROM    *
      6       C  * THE RELEVANT ECONOMIC MODEL.  COST IS FOUND BY INTERPOLATING IN THE*
      7       C  * COST ARRAY STORED IN  VOPT , USING THE METHOD OF PIECEWISE LINEAR  *
      8       C  * INTERPOLATION (SCHULTZ, 1973).                                     *
      9       C  *                                                                    *
     10       C  ************************************************************************
     11       C
     12             SUBROUTINE COST(X,ECU,EXCOS,RANGE,IEPARA,EPARAM,COSTX,COST)
     13       C
     14       C
     15             DIMENSION NNB(2,4),RANGE(4,5)
     16             DIMENSION IEPARA(25),EPARAM(25)
     17             DOUBLE PRECISION XINT(4),FUNT,ECU(4)
     18             DOUBLE PRECISION EXCOS,X(4,2),RLOSS,COSTX,COST
     19             DOUBLE PRECISION WORK,SPOADD,DEB,SUM
     20       C
     21             COMMON VOPT(5,5,5,5,20)
     22       C
     23       C
     24             NUMCRT=IEPARA(12)
     25             RHO = EPARAM(14)
     26             GO TO(10,50),NUMCRT
     27       C    ***************************************************
     28       C     START OF CALCULATIONS FOR THE SUBSISTENCE FARMING CRITERION
     29       C
     30       C NOW UNPACK SOME USEFUL PARAMETERS
     31          10 WORK=ECU(1)
     32             SPOADD=ECU(4)
     33             FIXCOS=EPARAM(9)
     34             WORMUL=EPARAM(10)
     35             HAFWOR=EPARAM(11)
     36             CONRAT=EPARAM(12)
     37             SPOCOS=EPARAM(13)
     38       C LOSS FORMULA FOR SUBSISTENCE FARMING
     39             DEB = 0.
     40             RLOSS=X(1,1)+X(1,2)*WORK*WORMUL/(WORK+HAFWOR)
     41            1 -FIXCOS-CONRAT*WORK-SPOCOS*SPOADD
     42             IF(RLOSS.GE.-300.AND.RLOSS.LE.300)GOTO11
     43             DEB = DABS(RLOSS) -300
     44             IF(RLOSS)12,12,13
     45          12 RLOSS = -300
     46             GO TO 11
     47          13 RLOSS = 300
     48          11 COSTX=RLOSS*RLOSS +  600*DEB/.08
     49             X(1,2)=RLOSS
     50       C NOW GO TO 100 TO INTERPOLATE IN VOPT FOR EOV
     51             GO TO 100
     52       C ***********************************************
     53       C     START OF CALCULATIONS FOR THE AGRIBUSINESS CRITERION
     54       C
     55          50 BASRAT = EPARAM(19)
     56             FERCOS = EPARAM(20)
     57             QOPT   = EPARAM(21)
     58             FIXCOS = EPARAM(9)
     65
```

```
      SPOCOS = EPARAM(13)
      PROPSO = ECU(1)
      AREA   = ECU(2)
      SUM= X(1,1) + AREA*X(1,2)
      IF(SUM.GT.500)SUM=500
      QUANT  = PROPSO*SUM
      RETN   = QUANT*BASRAT*DEXP(-QUANT/QOPT) -FIXCOS-FERCOS*AREA
     1 -AREA*SPOCOS*ECU(4)
      X(1,2) = SUM - QUANT
      IF(X(1,2).LT.0.)X(1,2)=0.
      DO 51 I = 2,4
   51 X(I,2)= X(I,2)*AREA
      COSTX= - RETN
C ***************************************************
C     START OF THE INTERPOLATION PROCEDURE
C
  100 DO 110 I = 1,4
      IF(X(I,2).LE.RANGE(I,1))GOTO115
      IF(X(I,2).GT.RANGE(I,5))GOTO120
      DO 125 J = 1,4
      JM = J+1
      IF(X(I,2).GT.RANGE(I,JM))GOTO125
      NNB(1,I) = J
      NNB(2,I) = JM
      GO TO 110
  125 CONTINUE
      GOTO110
  115 NNB(2,I)=2
      NNB(1,I) = 1
      GO TO 110
  120 NNB(1,I)=4
      NNB(2,I) = 5
  110 CONTINUE
C  NOW FORM THE INTERPOLATING FUNCTIONS
      COST = 0.
      DO 130 I1 = 1,2
      J1=NNB(I1,1)
      XINT(1)  = (-1)**I1*(X(1,2)-RANGE(1,J1))/(RANGE(1,2)-RANGE(1,1))
      DO 130I2 = 1,2
      J2 = NNB(I2,2)
      XINT(2) = (-1)**I2*(X(2,2)-RANGE(2,J2))/(RANGE(2,2)-RANGE(2,1))
      DO 130I3=1,2
      J3 = NNB(I3,3)
      XINT(3) = (-1)**I3*(X(3,2)-RANGE(3,J3))/(RANGE(3,2)-RANGE(3,1))
      DO 130 I4 = 1,2
      J4 = NNB(I4,4)
      XINT(4) = (-1)**I4*(X(4,2)-RANGE(4,J4))/(RANGE(4,2)-RANGE(4,1))
      FUNT= VOPT(J1,J2,J3,J4,11)
      DO 135 J= 1,4
  135 FUNT = FUNT*(1.+XINT(J))
      COST = COST +FUNT
  130 CONTINUE
C NOW UPDATE E(V)
      EXCOS = RHO*(COSTX + COST)
C
C
      RETURN
      END
```

```
DKM80I*OPTIPLANT(1).DATA
      1        -.300+03 .000+00 .000+00 .000+00
      2        -.150+03 .250+03 .250+03 .250+03
      3         .000+00 .500+03 .500+03 .500+03
      4         .150+03 .750+03 .750+03 .750+03
      5         .300+03 .100+04 .100+04 .100+04
      6              NUMAT
      7            5     9     5    10     7     9     7    10
      8             IENAB
      9            0
     10             N
     11           68
     12             IMOD
     13            2
     14             NUMCRT
     15            1
     16             ISTART
     17           20
     18             ISTEP
     19           24
     20             LATEP
     21            7
     22             INFECP
     23           20
     24             ICPTIM
     25            7
     26             IYCRIT
     27           60
     28             IBYP
     29            1
     30             IGSTEP
     31            1
     32             IPRINT
     33            1
     34             ITLIM
     35           10
     36             VACANY
     37            0
     38             VACANT
     39            0
     40             VACANT
     41            0
     42             VIGORS
     43       +00.10000+06+00.10000+06+00.10000+06
     44              F
     45       +00.50000-05
     46              H
     47       +00.00000+00
     48              YPAM1
     49       -00.25530+02
     50              YPAM2
     51       +00.27170+02
     52              DENPRO
     53       +00.13000+01
     54              FIXCOS
     55       +00.36525+03
     56              WORMUL
     57       +00.50000+01
     58              HAFWOR
```

```
     59        +00.10000+03
     60              CONRAT
     61        +00.15000+01
     62              SPOCOS
     63        +00.10000+00
     64              RHO
     65        +00.92000+00
     66              SPEXT
     67        +00.50000+02
     68              SEED
     69        +00.11200-01
     70              GROW
     71        +00.23000+00
     72              FBACK
     73        +00.30000+00
     74              BASRAT
     75        +00.10000+03
     76              FERCOS
     77        +00.10000+04
     78              QOPT
     79        +00.30000+03
     80              GAIN
     81        +00.10000+01
     82              FACT
     83        +00.10000+01
     84              VACANT
     85        +00.00000+00
     86              VACANT
     87        +00.00000+00
     88              A
     89        +0.50000-001
     90              AB
     91        +0.90000+000
     92              B
     93        +0.50000-001
     94              AA
     95        +0.10000+001
     96              AABB
     97        +0.00000+000
     98              BB
     99        +0.10000+001
    100              WORK/PROPSO
    101        +0.50000+002
    102              AREA
    103        +0.10000+001
    104              PEXT
    105        +0.90000+000
    106              REALMU
    107        +0.80000+001
    108              COST
    109        +0.00000+000
    110              YIELD
    111        +0.00000+000
    112              X1
    113        +0.00000+000
    114              X2
    115        +0.00000+000
    116              X3
    117        +0.00000+000
```

```
  118                X4
  119        +0.00000+000
  120              VACANT
  121        +0.00000+000
  122              VACANT
  123        +0.00000+000
  124              VACANT
  125        +0.00000+000
  126              VACANT
  127        +0.00000+000
  128        +0.00000+000+0.00000+000+0.00000+000+0.00000+000+0.00000+000
  129        +0.00000+000+0.00000+000+0.00000+000+0.00000+000+0.00000+000
  130        +0.10000+001+0.10000+001+0.10000+001+0.10000+004+0.10000+004
  131        +0.10000+004+0.36525+003+0.10000+001+0.10000+001+0.20000+002
  132        +0.00000+000+0.00000+000+0.00000+000+0.00000+000+0.00000+000
  133        +0.00000+000+0.00000+000+0.00000+000+0.00000+000+0.00000+000
  134        +0.00000+000+0.00000+000+0.00000+000+0.00000+000+0.00000+000
  135        +0.00000+000+0.00000+000+0.00000+000+0.00000+000+0.00000+000

&RESUME,P PR2
```

BIBLIOGRAPHY

Adams, M.W., Ellingboe, A.H., and Rossman, E.C. (1971).
Biological uniformity and disease awareness.
Bioscience 21, 1067-1070.

Aiyer, A.K.Y.N. (1949).
Mixed cropping in India.
Indian J. agric. Sci. 19, 439-543.

Arnold, G.W., and Bennett, D. (1975).
The problem of finding an optimal solution.
In 'Study of Agricultural Systems'.
(Ed. G.E. Dalton.)
(Applied Science Publishers: London.)

Balch, M. and Fishburn, P.C. (1974).
Subjective expected utility for conditional priorities.
In 'Essays on Economic Behaviour under Uncertainty'. (Ed. M. Balch, D. McFadden and S. Wu.)
(American Elsevier Publishing Co.: New York.)

Banks, H.T. (1975).
Modelling and control in the biomedical sciences.
Lecture Notes in Biomathematics, Vol. 6.
(Berlin: Springer-Verlag.)

Bellman, R. (1957).
'Dynamic Programming'.
(Oxford University Press: London.)

Bellman, R. (1971).
'Introduction to the Mathematical Theory of Control Processes, Vol. II'
(Academic Press: New York.)

Bertalanffy, L. von (1968).
'General System Theory: Foundations, Development, Applications'.
(G. Braziller: New York.)

Biffen, R.H. (1905).
Mendel's laws of inheritance and wheat breeding.
J. agric. Sci. 1, 4-48.

Biffen, R.H. (1931).
The cereal rusts and their control.
Presidential address. Trans. Br. mycol. Soc. 16, 19-37.

Birta, L.G. (1976).
Some investigations in function minimization.
IEEE Trans. SMC. 6, 186-197.

Borlaug, N.E. (1953).
New approach to the breeding of wheat varieties resistant to
Puccinia graminis tritici.
Phytopathology 43, 467 (abstract).

Borlaug, N.E. (1958).
The use of multilineal or composite varieties to control airborne epidemic diseases of self-pollinated crop plants.
First International Wheat Genetics Symposium 12-27.

Borlaug, N.E. (1965).
Wheat, rust and people.
Phytopathology 55, 1088-1098.

Bose, A., and Shaw, M. (1974).
In vitro growth of wheat and flax rust fungi on complex and chemically defined media.
Can. J. Bot. 52, 1183-1196.

Box, M.J. (1965).
A New method of constrained optimisation and a comparison with other methods.
Comput. J. 8, 42-52.

Brigg, F.N., and Knowles, P.F. (1967).
"Introduction to Plant Breeding",
(Reinhold Publishing Corporation: New York).

Browning, J.A., and Frey, K.J. (1969).
Multiline cultivars as a means of disease control.
A. rev. phytopathol. 7, 355-382.

Brown, J.F. (1975).
Factors affecting relative ability of strains of fungal pathogens to survive in populations.
J. Aust. Inst. agric. Sci. 41, 3-11.

Brown, J.F., and Sharp, E.L. (1970).
The relative survival ability of pathogenic types of Puccinia striiformis in mixtures.
Phytopathology 60, 529-533.

Burdon, J.J. and Chilvers, G.A. (1976).
Epidemiology of Pythium-induced damping off in mixed species seedling stands.
Ann. appl. Biol. 82, 233-240.

Caldwell, R.M. (1968).
Breeding for general and for specific plant disease resistance.
In Third International Wheat Genetics Symposium
(Ed. K.W. Shepherd and K.W. Finlay.)
(Butterworths: Sydney, 1969.)

Caldwell, R.M., Schaffer, J.F.,
Compton, L.E., and Patterson, F.J. (1958).
Tolerance to cereal leaf rusts.
Science 178, 714-715.

Caten, C.E. (1974).
Inter-racial variation in Phytophthora infestans adaptation to field resistance for potato blight.
Ann. appl. Biol. 77, 259.

Centro Internacional de Mejoramento de Maiz y Trigo (1974).
Report on Wheat Improvement.
(El Batan: Mexico.)

Chamberlain, A.C. (1974).
Mass transfer to bean leaves.
Boundary-layer Meteorol. 6, 477-486.

Cheung, D.S.M., and Barber, H.N. (1972).
Activation of resistance of wheat to stem rust.
Trans. Br. mycol. Soc. 58, 333-336.

Clark, C.G., and Haswell, M. (1970).
The Economics of Subsistence Agriculture. 4th Edition.
(Macmillan: London.)

Clifford, B.C. (1968).
Relations of disease resistance mechanisms to pathogen dynamics in oat crown rust epidemiology.
Ph.D. Thesis. Purdue University.

Cornell University (1970).
The Eighty-Second Annual Report of the New York State College of Agriculture at Cornell University, Ithaca, New York, and of the Cornell University Agricultural Experiment Station, 1969.

Cournoyer, B.M. (1970).
Crown rust epiphytology with emphasis on the quantity and periodicity of spore dispersal from heterogeneous oat cultivar-rust race populations.
Ph.D. Thesis. Iowa State University of Science and Technology.

Curtis, GJ., and Hornsey, K.G. (1972).
Competition and yield compensation in relation to breeding sugar beet.
J. agric. Sci. 79, 115-119.

Davidon, W.C. (1975).
Optimally conditioned optimization algorithms without line searches.
Math. Program. 9, 1-30.

Day, P.R. (1973).
Genetic variability of crops.
A. rev. phytopathol. 11, 293-312.

Day, P.R. (1974).
'Genetics of Host-parasite Interactions'.
(W.H. Freeman: San Francisco.)

De Wit, C.T. (1960).
On Competition.
Versl. Landbouwk. Onderz. Ned. 66.8, 1-82.

Dem'yanov, V.F., and Malozemov, V.N. (1974).
'Introduction to Minimax'. (trans. D. Louvish)
(John Wiley: New York.)

Eversmayer, M.G., Burleigh, J.R., and Roelfs, A.P. (1973).
Equations for predicting wheat stem rust development.
Phytopathology 63, 348-351.

Flor, H.H. (1955).
Host-parasite interactions in flax rust - its genetic and other implications.
Phytopathology 45, 680-685.

Flor, H.H. (1971).
Current status of the gene-for-gene hypothesis.
A. rev. phytopathol. 9, 275-296.

Frankel, O.H. (1970).
Evaluation and utilization - introductory remarks.
In "Genetic Resources in Plants - Their Exploration and Conservation".
(Ed. O.H. Frankel and E. Bennett.)
(IBP Handbook no. 11: London.)

Frankel, O.H. (1974).
Genetic conservation: our evolutionary responsibility.
Genetics, N.Y. 78, 53-65.

Frankel, O.H., and Bennett, E. (1970).
Genetic resources - introduction.
In "Genetic Resources in Plants - Their Exploration and Conservation".
(Ed. O.H. Frankel and E. Bennett.)
(IBP Handbook no. 11: London.)

Frey, K.J., Browning, J.A., and Grindeland, R.L. (1971).
Implementation of oat multiline cultivar breeding.
In "Mutation Breeding for Disease Resistance."
(IAEA PL 412, IAEA: Vienna.)

Glansdorf, P., and Progogine, I. (1971).
'Thermodynamic Theory of Structure, Stability and Fluctuations'.
(Wiley-Interscience: New York.)

Gregory, P.H. (1973).
'The Microbiology of the Atmosphere', 2nd ed.
(L. Hill: London.)

Halkin, H. (1964).
Optimal control for systems described by difference equations.
Adv. Contr. Syst. 1, 173-196.

Harlan, H.V., and Martini, M.L. (1936).
Problems and results in barley breeding.
U.S.D.A. Yearbook of Agriculture, 303-346.

Harlan, J.R. (1972).
Genetics of disaster.
J. Environ. Qual. 1, 212-215.

Hesketh, J.D., and Jones, J.W. (1976).
Some comments on computer simulators for plant growth.
Ecological Modelling 2, 235-247.

Hooker, A.L. (1972).
Southern leaf blight of corn - present status and future prospects.
J. Environ. Qual. 1, 244-249.

Hooker, A.L. (1967).
The genetics and expression of resistance in plants to rusts of the genus Puccinia.
A. rev. Phytopathol. 5, 163-182.

International Atomic Energy Agency (1971).
"Mutation Breeding for Disease Resistance."
(International Atomic Energy Agency: Vienna.)

James, W.C. (1974).
Assessment of plant diseases and losses.
A. rev. phytopathol. 12, 27-48.

Jensen, N.F. (1952).
Intraverietal diversification in oat breeding.
Agron. J. 44, 30-34.

Jensen, N.F. (1965).
Multiline superiority in cereals.
Crop Sci. 5, 566-568.

Jensen, N.F. (1966).
Broadbase hybrid wheat.
Crop Sci. 6, 376-377.

Johnson, R., and Allen, D.J. (1975).
Induced resistance to most diseases and its potential role in the resistance of multiline varieties.
Ann. appl. Biol. 80, 359-363.

Johnson, T. (1961).
Man guided evolution in plant rusts.
Science 133, 357-362.

Jordan, B.W., and Polak, E. (1964).
Theory of a class of discrete optimal control systems.
J. Electron. Control 17, 697-711.

Jorgensen, J.H. (1971).
Comparison of induced mutant genes with spontaneous genes in barley conditioning resistance to powdery mildew.
In "Mutation Breeding for Disease Resistance."
(IAEA PL-412: Vienna.)

Jowett, D., Browning, J.A., and Haning, B.C. (1974).
Non-linear disease progress curves.
In 'Epidemics of Plant Diseases: Mathematical Analysis and Modelling'.
(Ed. J. Kranz.)
(Springer Verlag: Berlin.)

Justesen, S.H., and Tammes, P.M.L. (1960).
Studies of yield loss. I. The self-limiting effect of injurious yield or competitive organism on crop yield.
Tijdschr. Plziekt. 66, 281-287.

Kalman, R.E. (1963).
Mathematical descriptions of linear dynamical systems.
Siam J. Contr. 1, 152-192.

Kiyosawa, S. (1972).
Theoretical comparison between mixture and rotation cultivation of disease-resistant varieties.
Ann. phytopath. Soc. Japan. 38, 52-59.

Kiyosawa, S. (1976).
Pathogenic variations of Pyricularia oryzae and their use in genetic and breeding studies.
Sabrao Jnl. 8, 53-67.

Kiyosawa, S., and Shiyomi, M. (1972).
A theoretical evaluation of the effect of mixing resistant variety with susceptible variety for controlling plant diseases.
Ann. phytopath. Soc. Japan. 38, 41-51.

Knott, D.R. (1971).
Can losses from wheat stem rust be eliminated in North America?
Crop Sci. 11, 97-99.

Kochman, J.K., and Brown, J.F. (1975).
Studies on the mechanism of cross protection in cereal rusts.
Physiol. Plant Pathol. 6, 19-27.

Kowalik, J., and Osborne, M.R. (1968).
'Methods for Unconstrained Optimization'.
(American Elsevier Publishing Co.: New York.)

Kuc, J. (1972).
Phytoalexins,
A. rev. phytopathol. 10, 207-232.

Kushner, H. (1971).
'Introduction to Stochastic Control'.
(Holt, Rinehart and Winston: New York.)

Keftwich, R.H. (1970).
"The Price System and Resource Allocation".
(The Dryden Press: Hingdale, Illinois.)

Leonard, K.J. (1969a).
Factors affecting rates of stem rust increase in mixed plantings of susceptible and resistant varieties.
Phytopathology 59, 1845-1850.

Leonard, K.J. (1969b).
Selection in heterogeneous populations of Puccinia graminis f. sp. avenae.
Phytopathology 59, 1851-1857.

Leonard, K.J. (1969c).
Genetic equilibria in host-pathogen systems.
Phytopathology 59, 1858-1863.

Levins, R. (1968).
'Evolution in Changing Environments'.
(Princeton University Press: Princeton.)

Littlefield, L.J. (1969).
Flax rust resistance induced by prior inoculation with a virulent race of Melampsora lini.
Phytopathology 59, 1323-1328.

Luce, R.D., and Raiffa, H. (1957).
'Games and Decisions: Introduction and Critical Survey'.
(John Wiley: New York.)

Luig, N.H., and Watson, I.A. (1970).
The effect of complex genetic resistance in wheat on the variability of Puccinia graminis f.sp. triciti.
Proc. Linn. Soc. N.S.W. 95, 22-45.

Luke, H.H., Chapman, W.H., and Barnett, R.D. (1972).
Horizontal resistance of red rustproof oats to crown rust.
Phytopathology 62, 414-417.

Luke, H.H., Barnett, R.D., and Pfahler, P.L. (1975).
Inheritance of horizontal resistance to crown rust in oats.
Phytopathology 65, 631-632.

Lupton, F.G.H. (1972).
Plant breeding for disease resistance.
In "Biology in Pest and Disease Control."
(Ed. D. Price Jones and M.E. Solomon.)
(Blackwell Scientific Publications: Oxford.)

Macer, R.C.F. (1972).
The resistance of cereals to yellow rust and its exploitation by plant breeding.
Proc. R. Soc. B. 181, 281-301.

Marshall, D.R., and Brown, A.H.D. (1973).
Stability of performance of mixtures and multilines.
Euphytica 27, 405-417.

Martens, J.W. (1973).
Competitive ability of oat stem rust races in mixtures.
Can. J. Bot. 51, 2233-2236.

McReynolds, S.R. (1972).
On the open-loop optimal control of Markov processes.
Proceedings of the 5th Hawaii Annual Conference on Systems Sciences.
327-330.

McAusland, I. (1969).
'Introduction to Optimal Control'.
(John Wiley and Sons, Inc.: New York.)

Meredith, D.S. (1973).
Significance of spore release and dispersal mechanisms in plant disease epidemiology.
A. rev. phytopathol. 11, 313-342.

Mode, L.J. (1958).
A mathematical model for the co-evolution of obligate parasites and their hosts.
Evolution 12, 158-165.

Melchers, L.E., and Parker, J.H. (1922).
Rust resistance in winter wheat varieties.
(U.S. Dept. Agriculture: Bulletin No. 1046.)

National Academy of Sciences (1972).
"Genetic Vulnerability of Major Crops".
(National Academy of Sciences: Washington, D.C.)

Nelson, R.R. (1972).
Stabilising racial populations of plant pathogens by use of resistance genes.
J. Environ. Qual. 1, 220-227.

Nemhauser, G.L. (1966).
'Introduction to Dynamic Programming'.
(John Wiley: New York.)

Ou, S.H. (1971).
Disease resistance in rice.
In "Mutation Breeding for Disease Resistance".
(IAEA PL 412, IAEA: Vienna.)

Owen, G. (1968).
'Game Theory'.
W.B. Saunders Co.: Philadelphia.)

Person, C., and Sidhu, G. (1971).
Genetics of host-parasite inter-relationships.
In "Mutation Breeding for Disease Resistance".
(IAEA: Vienna.)

Pesek, J. (1974).
Resistance of crops to diseases: A game-theory model.
In "Mathematical Models in Biology and Medicine"
(Ed. N.T.J. Bailey, Bl. Sendov, R. Tsanev)
(North-Holland Publishing Company, Amsterdam.)

Pontryagin, L.S., Boltyanskii, V.G., Gamkrelidze, R.V., and Mischenko, E.F. (1962).
'The Mathematical Theory of Optimal Processes'.
(John Wiley and Sons, Inc.: New York.)

Porter, B. (1969).
'Synthesis of Dynamical Systems'.
(Thomas Nelson and Sons Ltd.: London.)

Rawson, H.M., and Hofstra, G. (1969).
Translocation of ^{14}C assimilated at different stages by each leaf of the wheat plant.
Aust. J. Biol. Sci. 22, 321-331.

Roane, C.W. (1973).
Trends in breeding for disease resistance in crops.
A. rev. phytopathol. 11, 463-486.

Robinson, R.A. (1968).
The concept of vertical and horizontal resistance as illustrated by bacterial wilt of potatoes.
In "Phytopathological Papers, no. 10."
(Commonwealth Mycological Institute: Kew, Surrey, England.)

Rockefeller Foundation, The (1965).
'Program in the Agricultural Sciences'.
Annual Report, 1964-65.

Rockefeller Foundation, The (1966).
'Program in the Agricultural Sciences'.
Annual Report, 1965-66.

Roelfs, A.P. (1972).
Gradients in horizontal dispersal of cereal rust uredospores.
Phytopathology 62, 70-76.

Romig, R.W., and Calpouzos, L. (1970).
The relationship between stem rust and loss in yield of spring wheat.
Phytopathology 60, 1801-1805.

Rosen, H.R. (1949).
Oat parentage and procedures for combining resistance to crown rust, including race 45, and Helminthosporium blight.
Phytopathology 39, 20.

Rosen, J.B. (1960).
The gradient projection method for nonlinear programming.
Part i. Linear constraints.
SIAM. J. Appl. Math. 8, 181-217.

Sage, A.P. (1966).
'Optimum Systems Control'.
(Prentice-Hall, Inc.: Englewood Cliffs, New Jersey.)

Schafer, J.B. (1971).
Tolerance to plant disease.
A. rev. phytopathol. 9, 235-252.

Schrodter, H. (1960).
Dispersal by air and water - the flight and landing.
In 'Plant Pathology: an Advance Treatise', Vol 3.
(Ed. J.G. Horsfall and A.E. Dimond.)
(Academic Press: New York.)

Schultz, M.H. (1973).
'Spline Analysis'.
(Prentice-Hall Inc.: Englewood Cliffs, New Jersey.)

Shaner, G., Finney, R.E., and Patterson, F.L. (1975).
Expression and effectiveness of resistance in wheat to Septoria leaf blotch.
Phytopathology 65, 761-766.

Shoemaker, C. (1973a).
Optimization of agricultural pest management. I. Biological and mathematical background.
Math. Biosci. 16, 143-175.

Shoemaker, C. (1973b).
Optimization of agricultural pest management. 2.
Formulation of a control model.
Math. Biosci. 17, 357-365.

Shoemaker, C. (1973c).
Optimization of agricultural pest management. 3.
Results and extension of a model.
Math. Biosci. 18, 1-22.

Shubik, M. (1975).
'Games for Society, Business and War'.
(Elsevier: New York.)

Siberskii, K.S. (1975).
'Introduction to Topological Dynamics.'
(trans. L.F. Boron)
(Noordhoff: Leyden.)

Sibma, L., Kort, J., and de Wit, C.T. (1964).
Experiments on competition as a means of detecting possible damage by nematodes.
Jarrb. Inst. Biolog. Scheik. Onderz. Landbouwk., Wageningen. 119-124.

Simmonds, N.W. (1962).
Variability in crop plants, its use and conservation.
Biol. Rev. 37, 422-465.

Simons, M.D. (1972).
Polygenic resistance to plant disease and its use in breeding resistant cultivars.
J. Environ. Qual. 1, 232-240.

Simons, M.D. (1975).
Heritability of field resistance to the oat crown rust fungus.
Phytopathology 65, 324-328.

Slesinski, R.S., and Ellingboe, A.H. (1969).
The genetic control of primary infection of wheat by Erysiphe graminis f.sp. triciti.
Phytopathology 59, 1833-1837.

Smith, H.M. (1971).
Broadening the base of genetic variability in plants.
Theoret. app. Genetics 62, 265-276.

Sokolnikoff, I.S., and Redheffer, R.M. (1966).
'Mathematics of Physics and Modern Engineering'.
(McGraw-Hill Book Company: New York.)

Stakman, E.C., and Christensen, J.J. (1960).
The problem of breeding resistant varieties.
In "Plant Pathology: an Advanced Treatise" Vol. 3.
(Ed. J.G. Horsfall and A.E. Dimond.)
(Academic Press: New York.)

Stakman, E.C., and Harrar, J.G. (1957).
'Principles of Plant Pathology'.
(The Ronald Press Company: New York.)

Stakman, E.C., Parker, J.H., and Piemeisel, F.J. (1918).
Can biologic forms of stem rust on wheat change rapidly enough to interfere with breeding for rust resistance?
Jnl. Agr. Res. 14, 111-123.
(quoted in Stakman and Christensen, 1960.)

Stanton, T.R., Murphy, H.C., Coffman, F.A. and Humphrey, H.B. (1934).
Development of oats resistant to smuts and rusts.
Phytopathology 24, 165-167.

Stevens, N.E. (1939).
Disease, damage and pollination types in "grains".
Science 89, 339-340.

Stevens, N.E. (1942).
How plant breeding programs complicate plant disease problems.
Science 95, 313-316.

Stevens, R.B. (1949).
Replanting "discarded" varieties as a means of disease control.
Science 110, 49. (letter).

Stuckey, R.E., and Ellingboe, A.H. (1974).
Elongation of secondary Hyphae of Erysiphe graminis f.sp. triciti on wheat with compatible and incompatible parasite/host genotypes.
Phytopathology 64, 530-533.

Sumner, D.R., and Littrell, R.H. (1974).
Influence of tillage, planting date, inoculum survival, and mixed populations on epidemiology of southern corn leaf blight.
Phytopathology 64, 168-173.

Turnbull, C. (1973).
'The Mountain People'.
(Jonathan Cape: London.)

Thom, R.J. (1975).
'Structural Stability and Morphogenesis'.
(W.A. Benjamin: Reading, Mass.)

Trenbath, B.R. (1974).
Biomass productivity of mixtures.
Adv. Agron. 26, 177-210.

Trenbath, B.R. (1977).
Interactions among diverse hosts and diverse parasites.
Proceedings of a conference: The genetic basis of epidemics in agriculture. Ann. N.Y. Acad. Sci. (in press).

Van der Plank, J.E. (1960).
Analysis of epidemics.
In "Plant Pathology: an Advanced Treatise."
(Ed. J.G. Horsfall and A.E. Dimond.)
(Academic Press: New York)

Van der Plank, J.E. (1968).
'Disease Resistance in Plants'.
(Academic Press: New York.)

Van der Plank, J.E. (1975).
'Principles of Plant Infection'.
(Academic Press: New York.)

Van der Plank, J.E. (1963).
'Plant Diseases: Epidemics and Control'.
(Academic Press: New York.)

Van Dyne, G.M., and Abramsky,Z. (1975).
Agricultural systems models and modelling: an overview.
In "Study of Agricultural Systems"
(Ed. G.E. Dalton.)
(Applied Science Publishers: London.)

Volin, R.B., and Sharp, E.L. (1973).
Physiologic specialization and pathogen aggressiveness in stripe rust.
Phytopathology 63, 699-703.

Waggoner, P.E., and Horsfall, J.G. (1969).
EPIDEM A Simulator of Plant Disease Written for a Computer.
(Bulletin No. 698: The Connecticut Agricultural Experiment Station.)

Waggoner, P.E., Horsfall, J.G., and Lukens, R.J. (1972).
EPIMAY A Simulator of Southern Corn Leaf Blight.
(Bulletin No. 729: The Connecticut Agricultural Experiment Station.)

Waggoner, P.E., and Parlange, J.Y. (1974).
Mathematical model for spore germination at changing temperatures.
Phytopathology 64, 605-610.

Walsh, G.R. (1975).
'Methods of Optimization'.
(Wiley: London.)

Walton, P.D. (1969).
Multiline strains of spring wheat.
J. Hered. 60, 205.

Watson, I.A. (1970a).
The utilization of wild species in the breeding of cultivated crops resistant to plant pathogens.
In "Genetic Resources in Plants: their Exploration and Conservation."
(Ed. O.H. Frankel and E. Bennett.)
(IBP Handbook no. 11: London.)

Watson, I.A. (1970b).
Changes in virulence and population shifts in plant pathogens.
A. rev. phytopathol. 8, 209-230.

Watson, I.A., and Luig, N.H. (1968).
The ecology and genetics of host-pathogen relationships in wheat rusts in Australia.
Third International Wheat Genetics Symposium (Canberra.) 227-238.

Webster, R.K. (1974).
Recent advances in the genetics of plant pathogen fungi.
A. rev. phytopathol. 12, 331-353.

Wheeler, H. (1975).
"Plant Pathogenesis".
(Springer-Verlag: Berlin.)

Williams, P.H. (1975).
Genetics of resistance in plants.
Genetics, N.Y. 79, 409-419.

Windeknecht, T.W. (1971).
'General Dynamical Processes A Mathematical Introduction'.
(Academic Press: New York.)

Wolfe, M.S. (1975).
Forced evolution of cereal rusts and mildews.
Ann. appl. Biol. 81, 443.

Wood, R.K.S. (1967).
"Physiological Plant Pathology".
(Blackwell Scientific Publications: Oxford.)

Woodger, J.H. (1939).
'The Technique of Theory Construction'.
(University of Chicago Press: Chicago, III.)

Zeeman, E.C. (1976).
Catastrophe theory.
Scientific American 234, 65-83.