A Comparison of Seven Allocation Rules for a Clinical Trial Model

by Gordon Simons*
University of North Carolina

SUMMARY

The risk in a trial to compare two medical treatments is borne by the patients who receive the inferior treatment during the experimental phase and by those remaining after the experiment who will all receive the inferior treatment if the results are misleading. An allocation rule's task is to balance these competing risks by deciding, during the course of the trial, when the experimental phase should be terminated. This paper compares the performance characteristics of various allocation rules that have been mentioned in the literature, and it seeks to correlate the structure of these and other allocation rules with their performance characteristics. Much use is made of computer-based graphical techniques.

AMS 1980 *subject classification.* Primary 62L05

Key words and phrases. Clinical trials, Bayes rules, markovian states

*This research was supported by the National Science Foundation Grant DMS-8400602

1. Introduction.

Several allocation rules have been proposed which attempt to assign the better of two competing treatments to as many patients as possible. A few have been touted for their large sample properties, others for being Bayes or minimax, and one- a truncated sequential probability ratio test - for its simplicity. The first objective here is to compare these and other rules, for a moderate sample size, among themselves and with a best-possible standard. The second objective is to illustrate some novel computer-based graphical techniques. These are used to compare the allocation rules, and to find other, in some ways better, allocation rules. The techniques can be used by Bayesians as well as frequentists. The third objective is to correlate the structure of allocation rules with their performance characteristics.

An allocation rule is a stopping rule; its only task is to decide when an initial "testing phase" is to be stopped. During the testing phase, the two competing treatments are randomly allocated to pairs of patients; after the testing phase, all of the remaining patients are given the (same) more promising treatment. It is assumed that the total number of patients, the "horizon", is known in advance. This clinical-trial formulation was proposed by Anscombe (1963) for the setting of normally distributed treatment responses. The same formulation is suitable for the success-failure responses considered here. The risk (function) for an allocation rule is the usual "expected successes lost"- lost because the superior treatment is unknown. It is equal to the product of $|p_1-p_2|$ and the expected number of patients assigned to the inferior treatment (during and after the testing phase), where p_1 and p_2 are the unknown probabilities of success for the two treatments. The inferior treatment is, of course, the one with the smaller success probability.

Only symmetric stopping rules will be considered - rules which are indifferent to the ordering of the treatments. It seems appropriate to use a symmetric stopping rule whenever there is little or no reason to suspect (initially) that a particular treatment is superior. In such a setting, it is reasonable to view the treatment which achieves the most successes during the testing phase as the "more promising treatment" (that is to be given to all of the remaining patients).

The risk function $R_\Delta(p_1,p_2)$ for a symmetric rule Δ is symmetric in p_1 and p_2: $R_\Delta(p_2,p_1) = R_\Delta(p_1,p_2)$ for $p_1,p_2 \in [0,1]$. Its value at a particular point $(p_1,p_2) = (a,b)$ is equal to the Bayes risk for the rule when the prior is symmetric and supported on the two points (a,b) and (b,a). Thus *the risk at* (p_1,p_2) *can be minimized among all symmetric stopping rules by choosing the Bayes stopping rule for the symmetric prior supported on* (a,b) *and* (b,a). This Bayes rule, denoted $\Delta(a,b)$, and the value B(a,b) of its minimal Bayes risk, can easily be obtained numerically by using the standard backward induction formulas. (The details for the current setting are given in [7].) The difference $R^*_\Delta(p_1,p_2) = R_\Delta(p_1,p_2) - B(p_1,p_2)$ is the "regret function" for the rule Δ. Since

$$\min_{\Delta} R^*(p_1,p_2) = 0 \qquad , \quad p_1,p_2 \in [0,1], \tag{1}$$

a symmetric rule is recommended by a regret function that is close to zero for most of the points (p_1,p_2) in the unit square. This perspective motivates a useful graphical technique: *the plotting of the contour lines of the regret function.*

The regret contour plots for seven allocation rules appear in Figures Ia-Ig. These include the "Anscombe", "LLRS" (Lai-Levin-Robbins-Siegmund (1980)), "minimax", "truncated SPRT" and "envelope" rules. The first four of these are described in [3], and the latter is described, wthout a name, in the introduction of [7].

It is easy - by mere inspection - to compare the shapes of the regret contours for pairs of rules and, thereby, to critically assess their relative strengths and weaknesses. (More detailed comparisons can be made, with the aid of an overhead projector, by overlaying pairs of plots.) All of the plots appearing in this paper assume a horizon of 100 patients (a moderate sample size) in order to facilitate comparisons.

Another revealing graphical description is provided by "risk region" plots: A region of the parameter space - of the unit square - will be called an $\alpha\%$ *risk region*, for a particular allocation rule Δ, if the risk at each point within the region exceeds the minimal possible risk by no more than $\alpha\%$. The largest possible region of this type is given by

$$\mathcal{R}_\Delta(\alpha) = \{(p_1,p_2): R_\Delta(p_1,p_2) \leq (1+\alpha/100)\ B(p_1,p_2)\} \qquad (2)$$

Clearly one should be looking for an allocation rule which has relatively large risk regions.

Risk region plots for the seven rules are displayed in Figures IIa-IIg. The Anscombe, minimax and truncated SPRT rules have fairly large 1% risk regions covering the central region of the parameter space. The "one per cent" rule has an even larger 1% risk region. It was discovered by using another graphical technique, described in Section 2. The "A75B25" rule - the Bayes rule $\Delta(.75,.25)$ - has a large, but very differently shaped, 1% risk region. The difference is elucidated in Section 2.

A practical Bayesian, who is willing to use a good suboptimal allocation rule, can effectively use regret contour plots and risk region plots. There are several reasons why a Bayesian might willing forgo the Bayes allocation rule:

(i) The Bayes rule is hard to compute for most priors.

(ii) The designing of a clinical trial is seldom the task of a single individual; conflicting interests must be resolved. One can imagine several Bayesians, with different priors, trying to find a single mutually acceptable allocation rule.

(iii) The amount of additional Bayes risk of a suboptimal rule can be assessed without a knowledge of the Bayes rule and without a knowledge of the minimal Bayes risk. This amount may be acceptably small. (See the example given below.)

For any prior distribution G and allocation rule Δ, the minimal Bayes risk is bounded below by $\int_0^1\int_0^1 B\,dG$, and bounded above by $\int_0^1\int_0^1 R_\Delta\,dG$. So the Bayes risk for Δ exceeds the minimal Bayes risk by no more than the integrated regret: $\int_0^1\int_0^1 R^*_\Delta\,dG$. A visual estimate of the size of the integrated regret can be obtained from the regret contour plot for Δ.

If the support of the prior G lies within an α% risk region $\mathcal{R}_\Delta(\alpha)$, then the Bayes risk for Δ can exceed the minimal Bayes risk by no more than α%. When $\alpha=1$, for instance, the rule Δ would probably be acceptable to most. If some of the support lies outside the region $\mathcal{R}_\Delta(\alpha)$, then the Bayes risk for Δ may be larger still by no more than $\gamma(1-G(\mathcal{R}_\Delta(\alpha)))$, where γ is the maximal regret outside the region. Several of the rules considered in this paper have the same value $\gamma=2$ when $\alpha=1$. The most attractive of these is the one per cent rule since it has the largest 1% risk region.

<u>An example</u>: Consider the uniform prior on the unit square. A quick examination of the regret contour plots (Figures Ia-Ig) suggests that the most promising candidate is the rule A75B25. It has a maximal regret of one, and a much smaller regret over much of the unit square. Its integrated regret is about .18. (Its nearest competitor is the one per cent rule with an integrated regret of .51.) This can be interpreted as follows: For the 100 patients of the clinical trial, the rule A75B25 (at worst) can be expected to yield .18 fewer successes than the (unknown) Bayes rule. To put this in perspective: Under a uniform prior, one can expect 64.16 successes if the A75B25 rule is used, one could expect 66.67 successes if one were told which treatment is superior, and, one could expect 50 successes if a treatment were chosen at random.

It may seem strange that an allocation rule (A25B25) that is optimal for a two-point symmetric prior performs so well for the diffuse uniform prior. A similar phenomenon for two-arm bandit allocation rules has been documented by Berry (1978). There is a particularly persuasive reason, in his setting, for replacing priors by two-point modifications. Feldman (1962) has shown (for priors supported on two symmetric points) that it is optimal to always choose the arm with the higher posterior probability of yielding a success. The Bayes rules for two-point symmetric priors are not so easily described for the present model. Nevertheless, they have a relatively simple structure that makes them easy to employ.

2. The Structure of Allocation Rules.

Several questions are raised by the plots appearing in Figures Ia-Ig and IIa-IIg: (a) Why do the Anscombe, minimax, truncated SPRT and one per cent rules have better regret contour and risk region plots than do the LLRS and envelope rules? (b) Why are the plots for the A75B25 rule so different from those of the other six rules? (c) Why do the Anscombe, minimax and one per cent rules have slightly better regret contour and risk region plots than does the truncated SPRT, particularly in the two corners where $|p_1-p_2|$ is small? (d) More generally, what is the right way to fashion an allocation rule from the sequence of success-failure data generated during the testing phase of the clinical trial? A useful technique for getting at questions of this sort is to study the Bayes rules. The class of (symmetric) Bayes rules is too complicated, in this setting, to obtain, but one can completely describe the class of Bayes rules for symmetric two-point priors. Four of the seven rules studied in this paper are in this class. The remaining three (the truncated SPRT, Anscombe and LLRS rules) have certain features in common with these Bayes rules. But the LLRS rule, and the Anscombe rule for a larger horizon, introduce a feature that is not common to these Bayes rules; this might be significant.

The Bayes (stopping) rules can be described in terms of Markovian states (n,r,s): Depending on the prior and the horizon N, certain states are *stopping states (points)*, and the remainder are *continuation states (points)*. Here, n denotes the current number of pairs of patients treated; r and s denote the current numbers of successfully treated patients by the two treatments. When the prior is symmetric and supported on two-points (a,b) and (b,a), the dependence of the Bayes rule $\Delta(a,b)$ on N and (n,r,s) can be described more simply in terms of bivariate Markovian states $(t,k) = (N-2n, |r-s|)$. Observe that t is the current number of untreated patients. For each pair a and b, there is an increasing sequence of nonnegative integers $T_0, T_1, T_2, \ldots$ The point (t,k) is a (Bayes) stopping point iff $t \leq T_k$. (See [3] and [7].)

This Bayes rule can be described without referring to the entire sequence of integers. For if $T_\ell \geq N-2\ell$, then stopping is guaranteed for every observable (t,k), $k \geq \ell$. And $\Delta(a,b)$ is fully identified by specifying the minimal possible value ℓ, called the *level* of the Bayes rule, and the partial sequence $T_0, \ldots, T_{\ell-1}$. The partial sequence can be obtained numerically from a straightforward induction-based algorithm. In addition, an excellent approximation is described in [1].

The values ℓ and $T_0, \ldots, T_{\ell-1}$ are not equally important: There is considerable numerical evidence that the risk at a particular point (p_1, p_2) is very sensitive to the choice of ℓ. Next in importance is $T_{\ell-1}$, then $T_{\ell-2}$, etc. The levels for $\Delta(a,b)$, $0 \leq a,b \leq 1$, are shown in Figure III. Notice that (for 100 patients) there are four possible levels: 1, 2, 3 and 4. The levels 2 and 3 together occupy 78% of the parameter area. Some values for $T_{\ell-1}$, are shown in Figure III, for ℓ=2, 3, 4. (For ℓ=1, T_0 is always 2.)

A level can be assigned to any allocation rule: The *level* is ℓ if ℓ is the least integer k for which stopping is guaranteed if $|r-s| \geq k$.

A second plot for the Bayes rules $\Delta(a,b)$ is given in Figure IV. This is a plot of the possible $(T_{\ell-2}, T_{\ell-1})$ pairs, where ℓ denotes the level of

$\Delta(a,b)$, ℓ=2, 3, 4. There is only one version of $\Delta(a,b)$ when ℓ=1 (defined by T_0=2), and it does not appear in the plot. The symbols plotted for ℓ=4 are coded so that the values of $T_{\ell-3}=T_1$ are specified as well. (Since T_0 is always 2, the plot completely describes the entire class of Bayes rules $\Delta(a,b)$.) In addition, the plot shows where the seven rules under study "fit into the picture": The minimax, one per cent, A75B25 and envelope rules are Bayes rules (corresponding to $\Delta(.594,.406)$, $\Delta(.66\ ,.34)$, $\Delta(.75\ ,.25)$, and $\Delta(a,b)$ for a and b near to .5, respectively); the truncated SPRT, Anscombe and LLRS rules are not. (Only a slightly modified version of the LLRS rule can be shown. This is discussed later.)

The plots in Figures I and II provide strong evidence for the importance of level. Roughly stated, *rules of like level have similar regret contour and risk region plots*. Striking evidence for this is provided by the minimax and truncated SPRT rules: While they have the same level 3 and similar plots, their values for (T_1,T_2) ((14, 47) and (0,0) respectively) are very different. Again roughly stated, *the regret for an allocation rule at a point* $(p_1,p_2)=(a,b)$ *is small if the allocation rule and the Bayes rule* $\Delta(a,b)$ *have the same level*. The converse is false near the line $p_1=p_2$, where all allocation rules (regardless of level) tend to have small regrets, but it is approximately correct elsewhere. This "explains" why there appears to be no allocation rule whose regret is small throughout the unit square. A compromise is needed. And the best compromises are to be found among the level-two and level-three rules. This is illustrated by the seven rules under study: The LLRS and envelope rules are level four and perform poorly away from the line $p_1=p_2$. The rest are level 2 or 3, and, in their various ways, perform much better.

There are other features besides level that matter. This fact is illustrated by the minimax and truncated SPRT rules. Their performance, though similar, are quite different in the corners near $(p_1,p_2)=(0,0)$ and (1,1). (See Figures I c,d and II c,d.) The right way to understand this appears to be as follows: Figure III shows that one wants a level 2 rule fairly near (1,1), and even a level 1 rule very near (1,1). While neither desire is met by either rule, the minimax rule more nearly approximates a level 2 rule through its control of T_2 (by making it large); and it more nearly approximates a level 1 rule through its control of T_1. The truncate SPRT's values $T_1=T_2=0$ are simply too small. Figure IV suggests that T_2 should be at least 43. The minimax rule has $T_2=47$. Appropriate values for T_1 seem less clear, but Figure IV suggests that it should be at least 14. The minimax rule has $T_1=14$.

A third plot for the Bayes rules $\Delta(a, b)$ is given in Figure V, a three dimensional plot of the minimal Bayes risk $B(a, b)$ as a function of a and b, $0 \leq a, b \leq 1$. Its maximum value - about 3.7 - occurs at a point of the form $(a_0, 1-a_0)$, where a_0 is .594 approximately. It is shown in [3] (with more details provided in [2]) that the Bayes rule $\Delta(a_0, 1-a_0)$ is, in fact, the minimax rule. The form of this rule is described in Figure IV.

The derivation of the one per cent rule began with a question: How much room for improvement is left by the Anscombe and minimax rules? For the sake of definiteness, a rule with a larger one per cent risk region was sought. The decision was made to look among the Bayes rules $\Delta(a,1-a)$. However, there are more than 100 such rules, and it is impractical to examine them all. A useful compromise is to consider the risk $R_{\Delta(a,1-a)}(p_1,1-p_1)$ as a function of a and p_1 ($a \geq .5$), and to examine the region defined by

$$\{(a,p_1):\ R_{\Delta(a,1-a)}(p_1,1-p_1) \leq 1.01 \times B(p_1,1-p_1)\}. \qquad (3)$$

(cf. (2).) The result appears in Figure VI. The location of the minimax, one per cent and A75B25 rules, and the location of the boundaries between levels are indicated by vertical lines. The vertical axis identifies the location of the envelope rule. The strong influence of "level" is strikingly apparent. The influence of other factors (such as the value of $T_{\ell-1}$) is greater in the higher levels. In various ways, the one per cent rule appears to be superior to the Anscombe and minimax rules. (Examine Figures I and II.) If one is not overly concerned about parameter values in the region slightly removed from the line $p_1=p_2$, then there are many, nearly equivalent level 2 rules - such as A75B25 - which are excellent competitors to the Anscombe, minimax, and one per cent rules. (See the example at the end of Section 1.)

The rules discussed in this paper have a common simplifying property: They decide on the basis of the current value of $(t,k) = (N-2n, |r-s|)$ whether it is an appropriate time to stop the testing phase. It is reasonable to ask how much is lost by using (t,k) instead of (n,r,s). The example at the end of Section 1 provides some evidence that the loss is modest, but the issue remains unresolved.

Most of the rules discussed in this paper have a further simplifying property: They are expressible in the form

$$\text{stop at } (t,k) \quad \text{if} \quad t \leq T_k \tag{4}$$

for some set of values $\{T_k\}$. The LLRS rule is an exception. When N=100, it stops for k=3 when t=94 and when $t \leq 78$, but continues for $78 < t < 94$. The modification of the LLRS rule, appearing in Figure IV, stops for k=3 when $t \leq 78$, and is expressible as in (4). The Anscombe rule exhibits the same kind of exceptional behavior for horizons somewhat larger than 100.

It is reasonable to ask whether exceptional behavior of the type just described can be worthwhile. There is some evidence that it can be: The LLRS rule has a substantially smaller risk than its modification when $|p_1-p_2|$ is large (and a slightly smaller risk when $|p_1-p_2|$ is small). When $|p_1-p_2|$ is large, one wants to stop quickly, to avoid unnecessary losses, and the LLRS rule strongly encourages this by allowing stopping at (t,k) = (94,3). Unfortunately, the LLRS rule is not a very good rule when N=100, and the author has failed to find any good rule whose form differs from (4). It may be that one can not benefit from the exceptionable behavior exhibited by the LLRS rule until the horizon is much larger than 100 - when, presumably, the level of a rule is not so important a feature.

3. The Graphics

All of the plots appearing in this paper were produced using SAS/GRAPH. And all were artistically enhanced. In most cases, the enhancements were simply cosmetic. In some cases, there was a need to "override" an inadequacy in the computer graphics package. Drawing smooth and accurate contour lines is not an easy task. SAS/GRAPH performed this task well most of the time, as well as could be expected at other times, and poorly in a few instances. An example of the latter appears in Figure IIg, where the boundary of the 1% risk region could have been smoother if an appropriate use of splines had been made.

A substantial effort was required to produce the smooth contours appearing in Figure III. A major difficulty is caused by the fact that the dependent variable is discrete (integer valued). This makes it impossible to use interpolation effectively. And a traditional use of splines is not going to help either. Initial plots based on a grid size of .01 (10,201 points) in the unit square were not very smooth.

A possible approach to this kind of problem would be to assume that the dependent variable is implicitly derived, by rounding, from another (unseen) dependent variable which *is* smoothly related to the two (continuous-valued) independent variables. The resolution, in the present setting, was to use the smoothly dependent approximations (to the T_k's) derived in Theorem 2 of [7]. These approximations introduced no error for the outer two contours appearing in Figure III except very near the line $a = b$. And the approach was helpful in smoothing out the inner contour.

References

[1] Anscombe, F.J. (1963). Sequential medical trials. J. Amer. Statist. Assoc., 58, 365-383.

[2] Bather, J.A. and Simons, G.D. (1984). The minimax risk for clinical trials. Institute of Statistics Mimeo Series #1565, University of North Carolina.

[3] Bather, J.A. and Simons, G.D. (1985). The minimax risk for clinical trials. To be published, JRSS B.

[4] Berry, D.A. (1978). Modified two-arm bandit strategies for certain clinical trials. J. Amer. Statist. Assoc. 73, 339-345.

[5] Feldman, D. (1962). Contributions to the two-armed bandit problem. Ann. Math. Statist. 33, 847-856.

[6] Lai, T.L., Levin, B., Robbins, H. and Siegmund, D. (1980). Sequential medical trials. Proc. Natl. Acad. Sci. USA, 77, 3135-3188.

[7] Simons, G.D. (1985). The Bayes rule for a clinical-trials model when the treatment responses are successes and failures. Institute of Statistics Mimeo Series #1551, University of North Carolina.

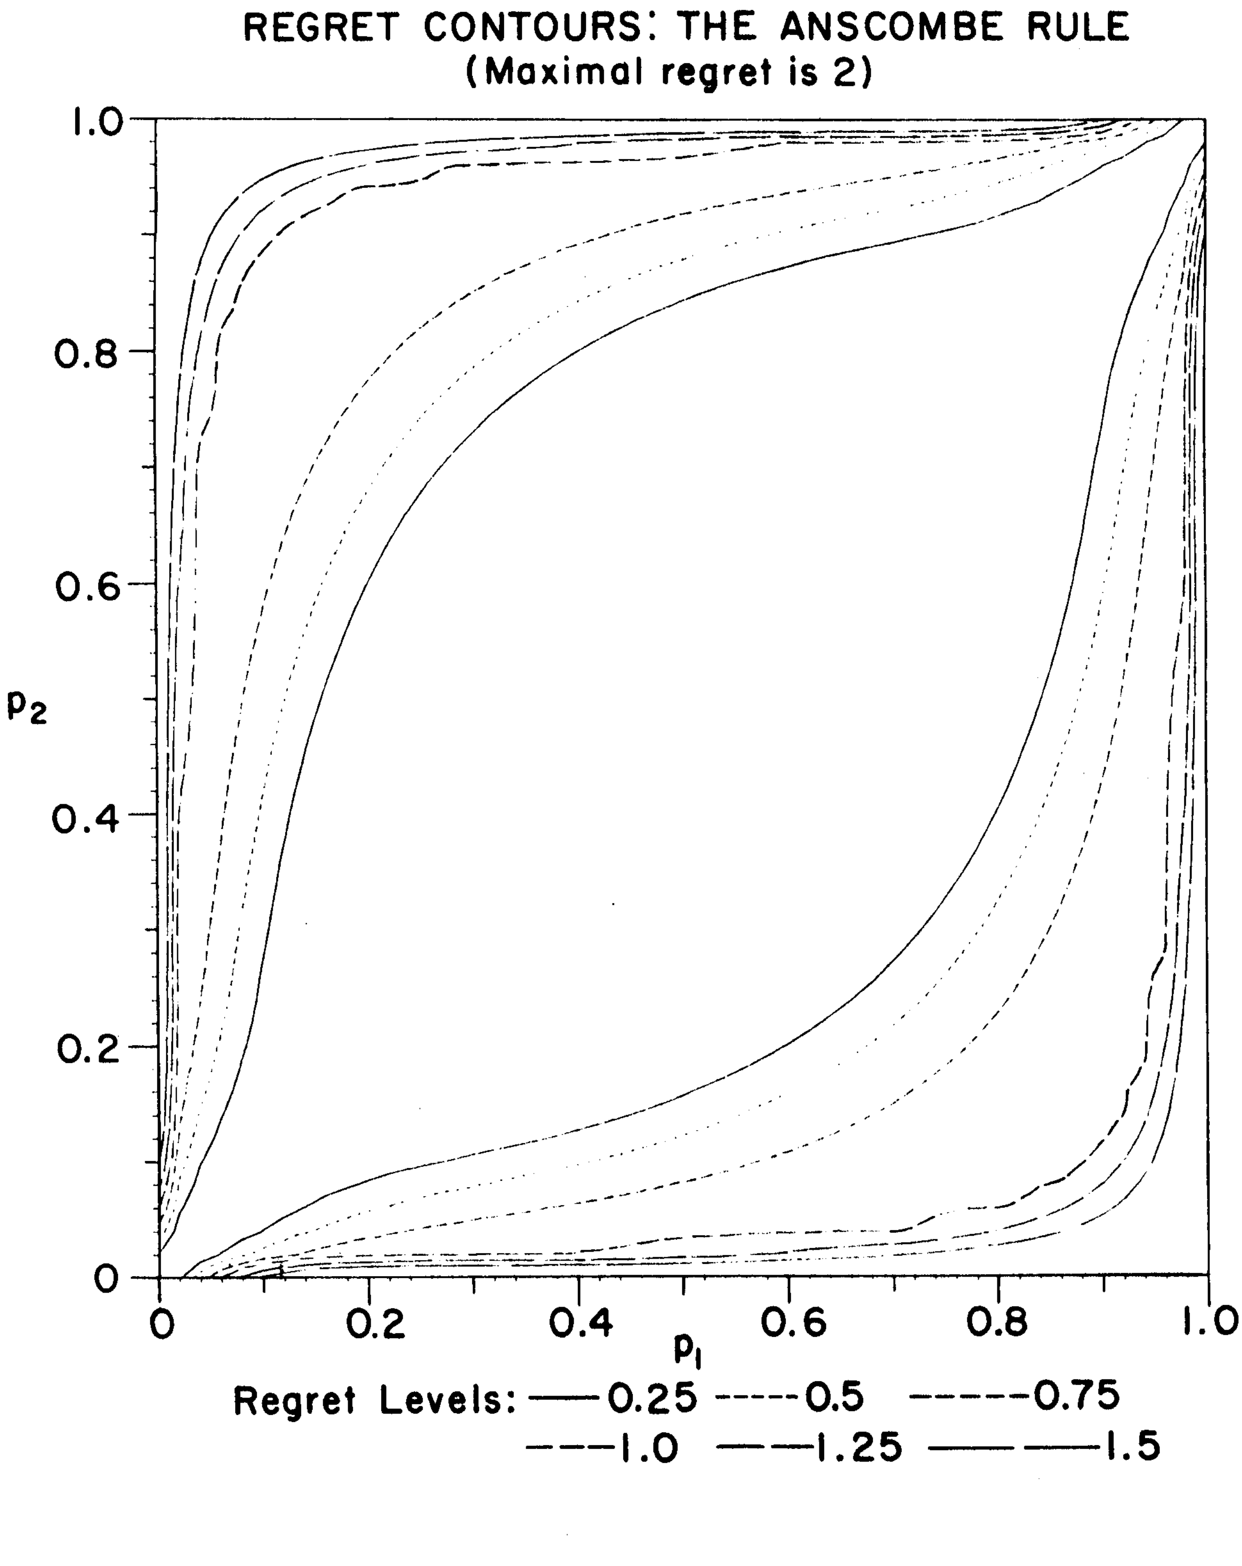
REGRET CONTOURS: THE ANSCOMBE RULE
(Maximal regret is 2)
1.0
0.8
0.6
0.4
0.2
0
p_2
0
0.2
0.4
0.6
0.8
1.0
p_1
Regret Levels: 0.25 0.5 0.75
1.0 1.25 1.5

Figures I a

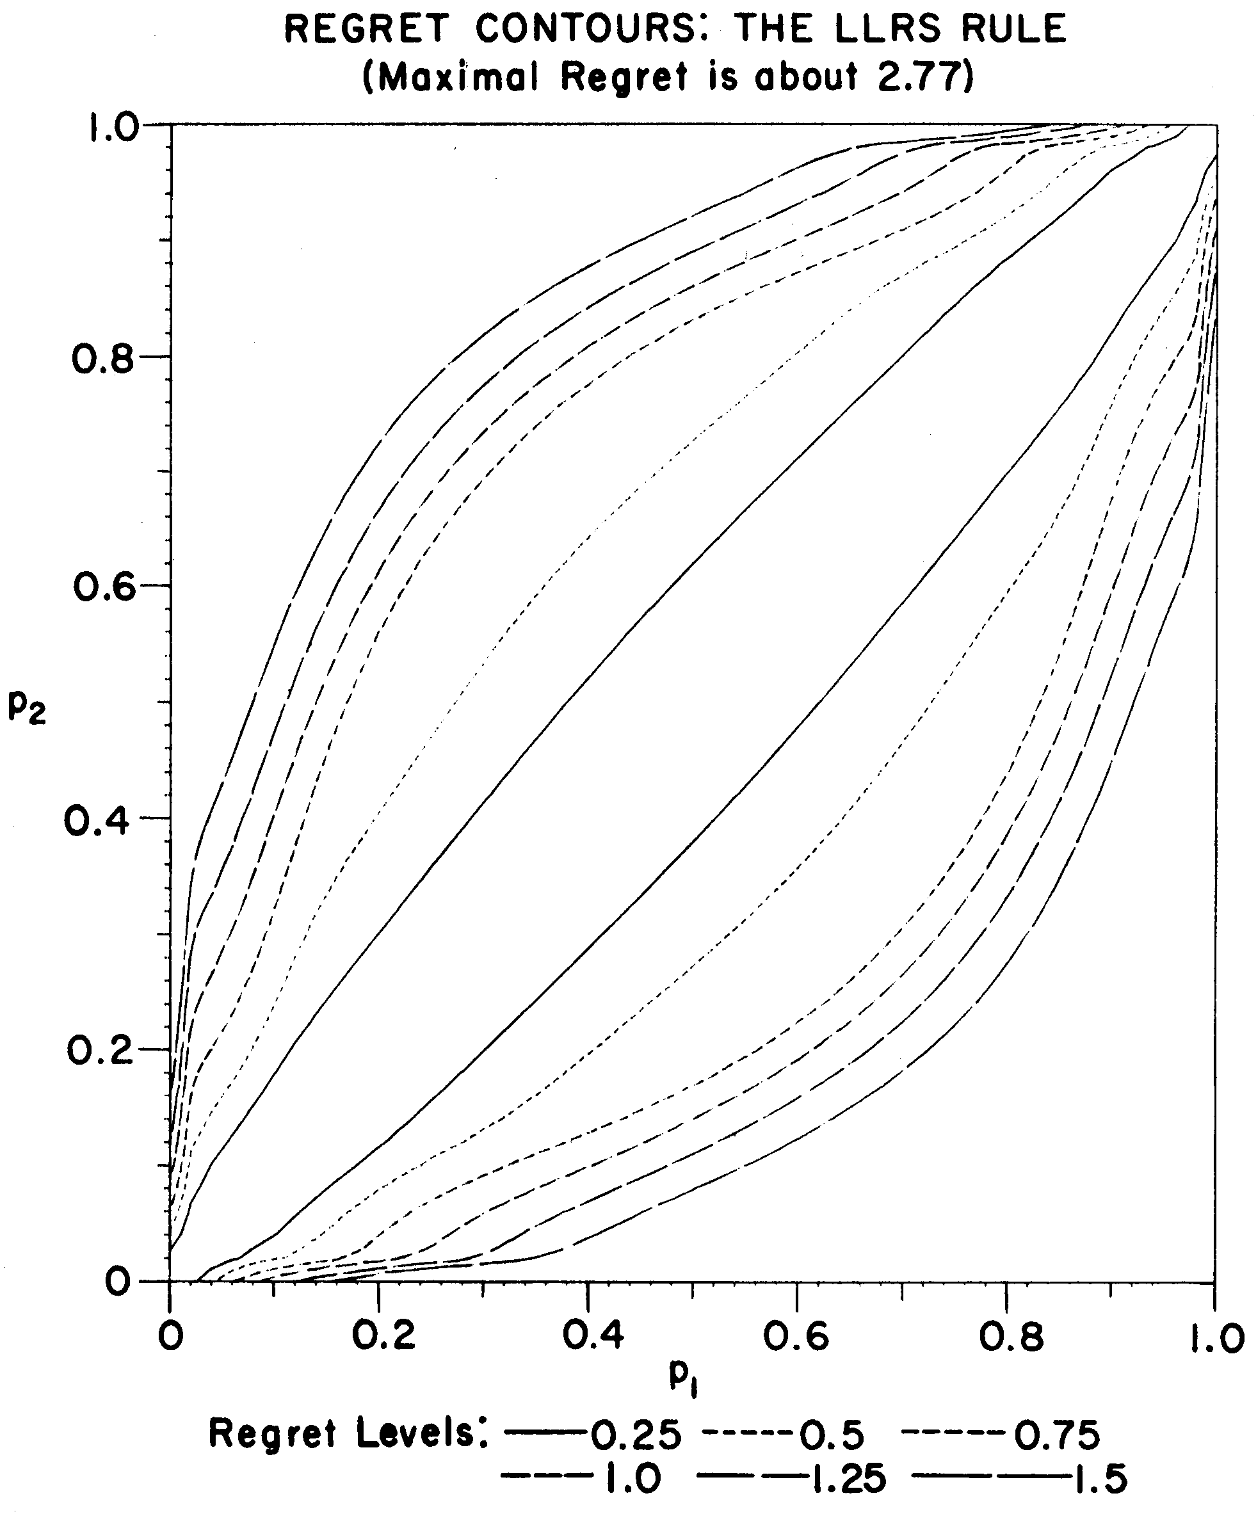
REGRET CONTOURS: THE LLRS RULE
(Maximal Regret is about 2.77)
1.0
0.8
0.6
0.4
0.2
0
p_2
0
0.2
0.4
0.6
0.8
1.0
p_1
Regret Levels: 0.25 0.5 0.75
1.0 1.25 1.5

Figure I b

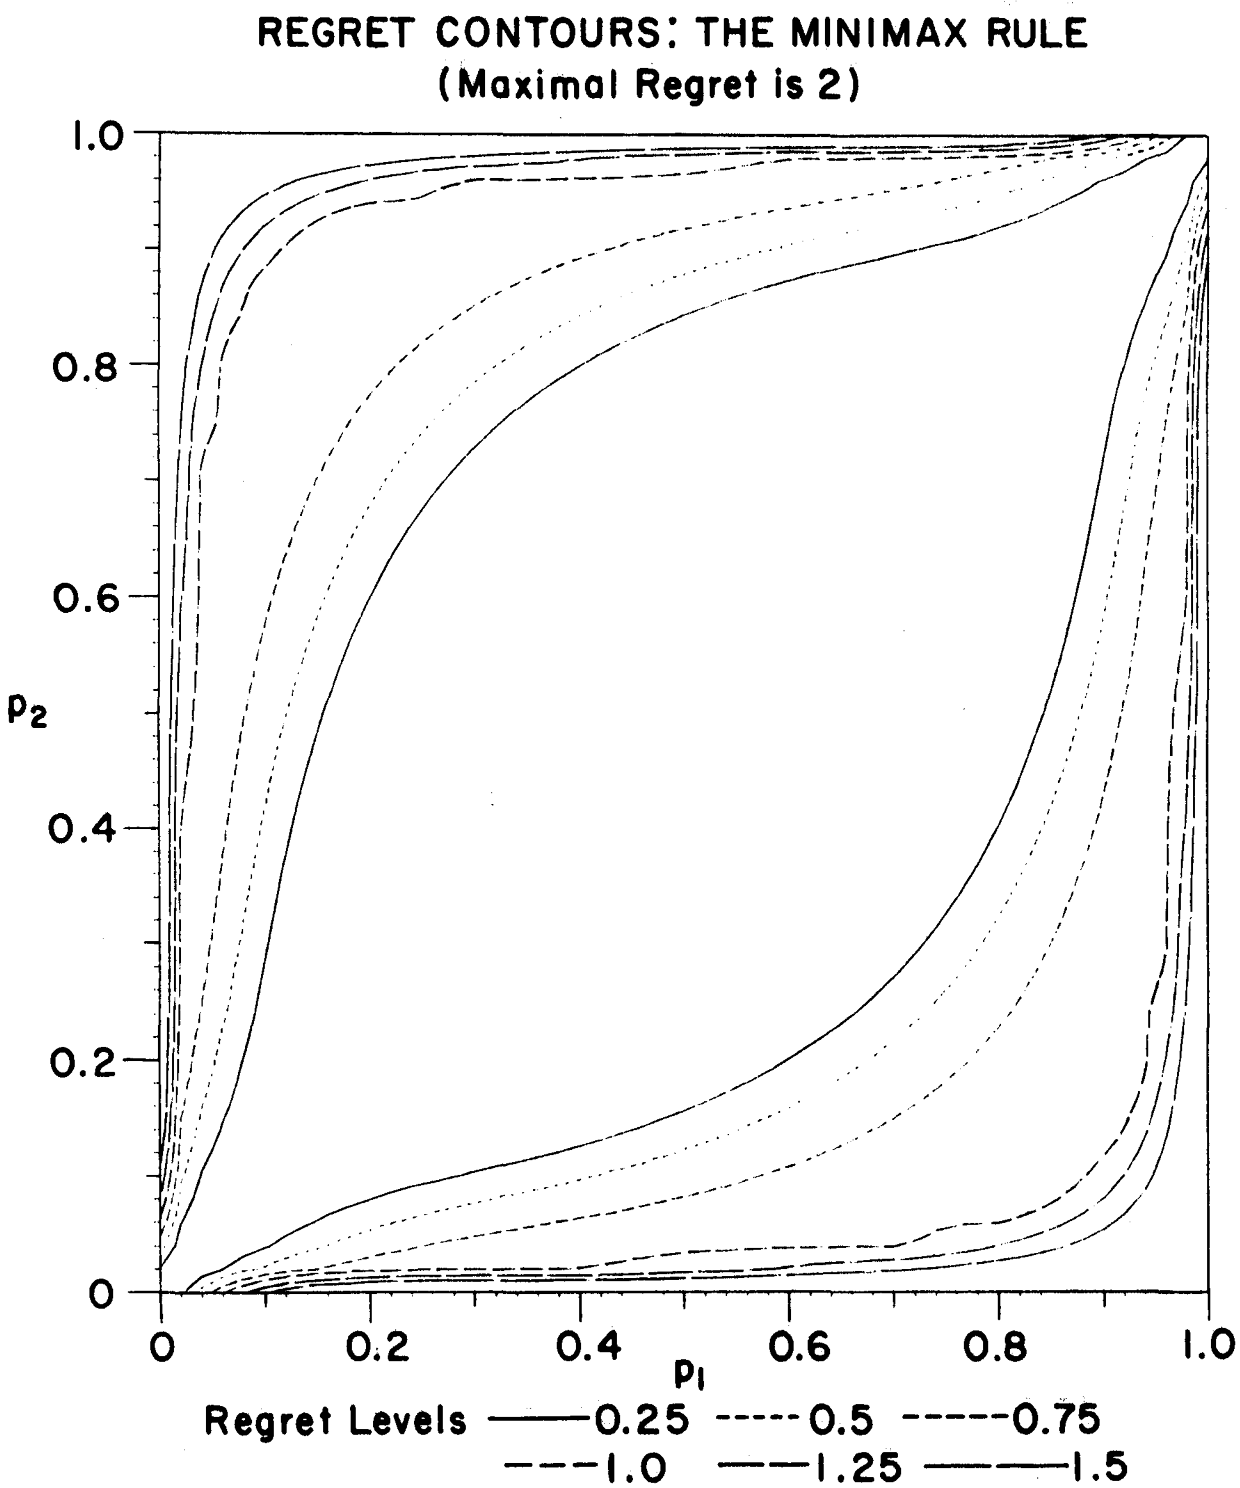
REGRET CONTOURS: THE MINIMAX RULE
(Maximal Regret is 2)
1.0
0.8
0.6
0.4
0.2
0
p2
0
0.2
0.4
0.6
0.8
1.0
p1
Regret Levels —0.25 -----0.5 ------0.75
---1.0 —1.25 ——1.5

Figure Ic

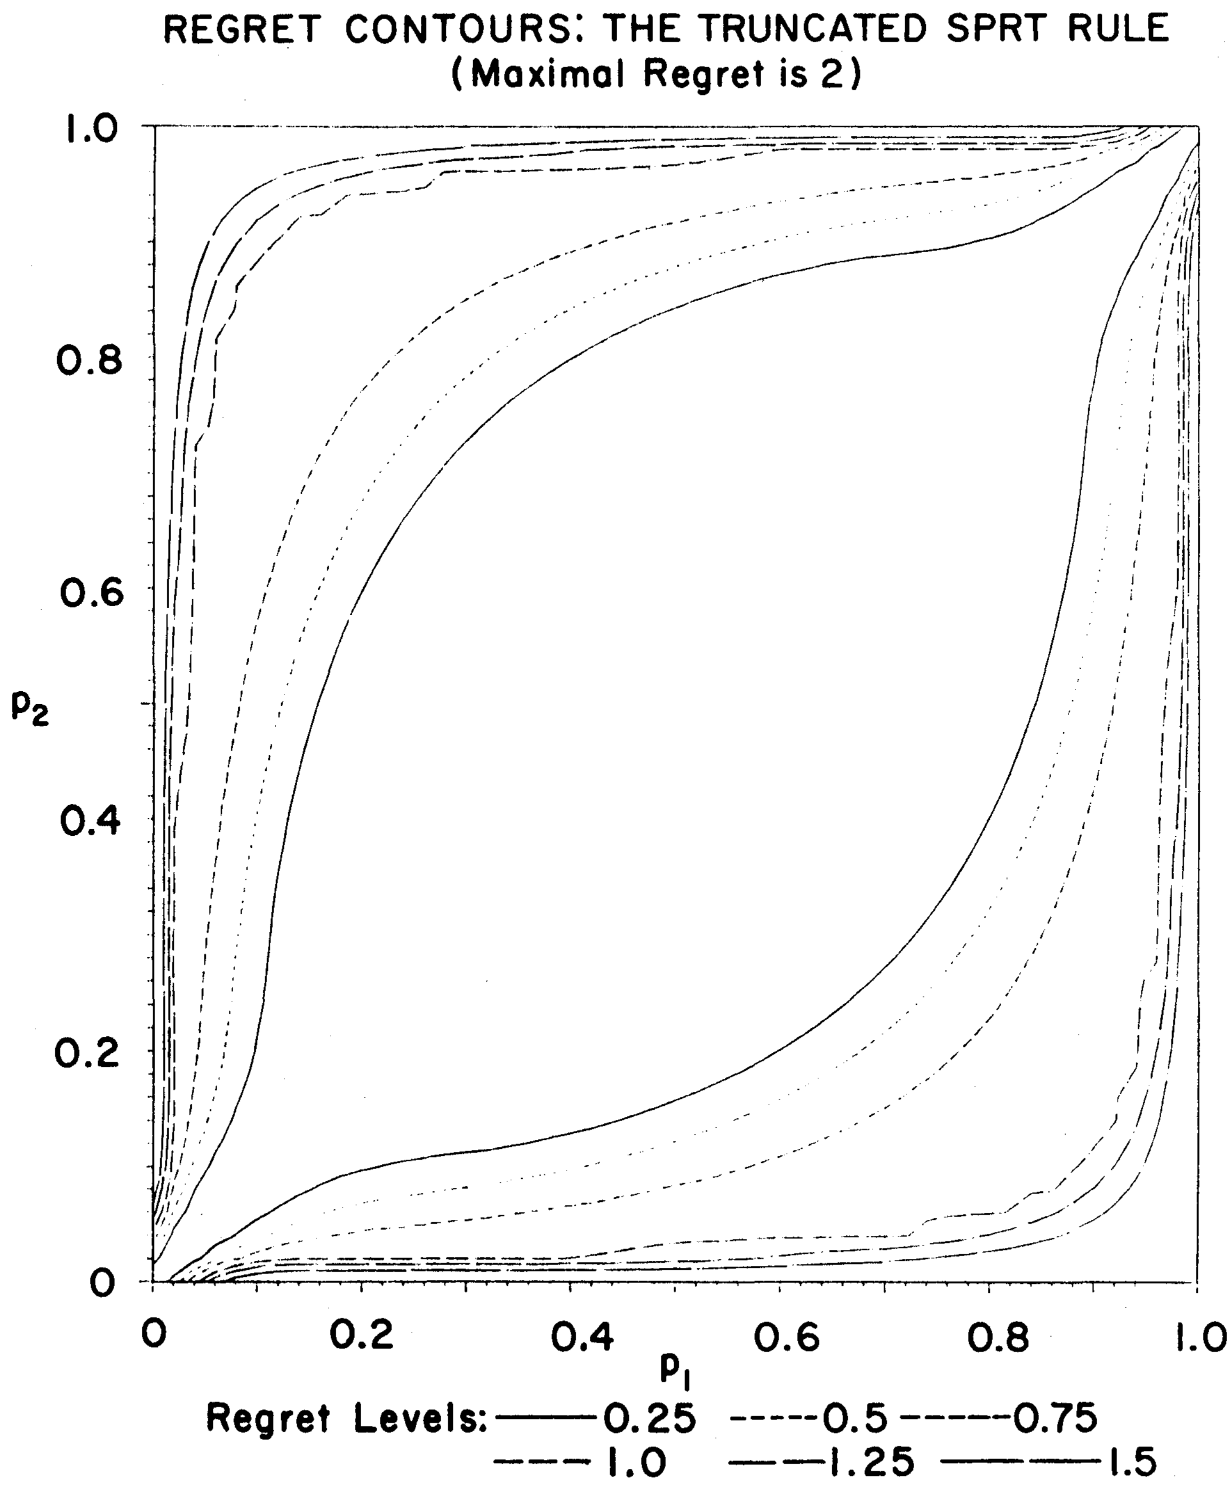
REGRET CONTOURS: THE TRUNCATED SPRT RULE
(Maximal Regret is 2)
1.0
0.8
0.6
0.4
0.2
0
p_2
0
0.2
0.4
0.6
0.8
1.0
p_1
Regret Levels: 0.25 0.5 0.75
1.0 1.25 1.5

Figure Id

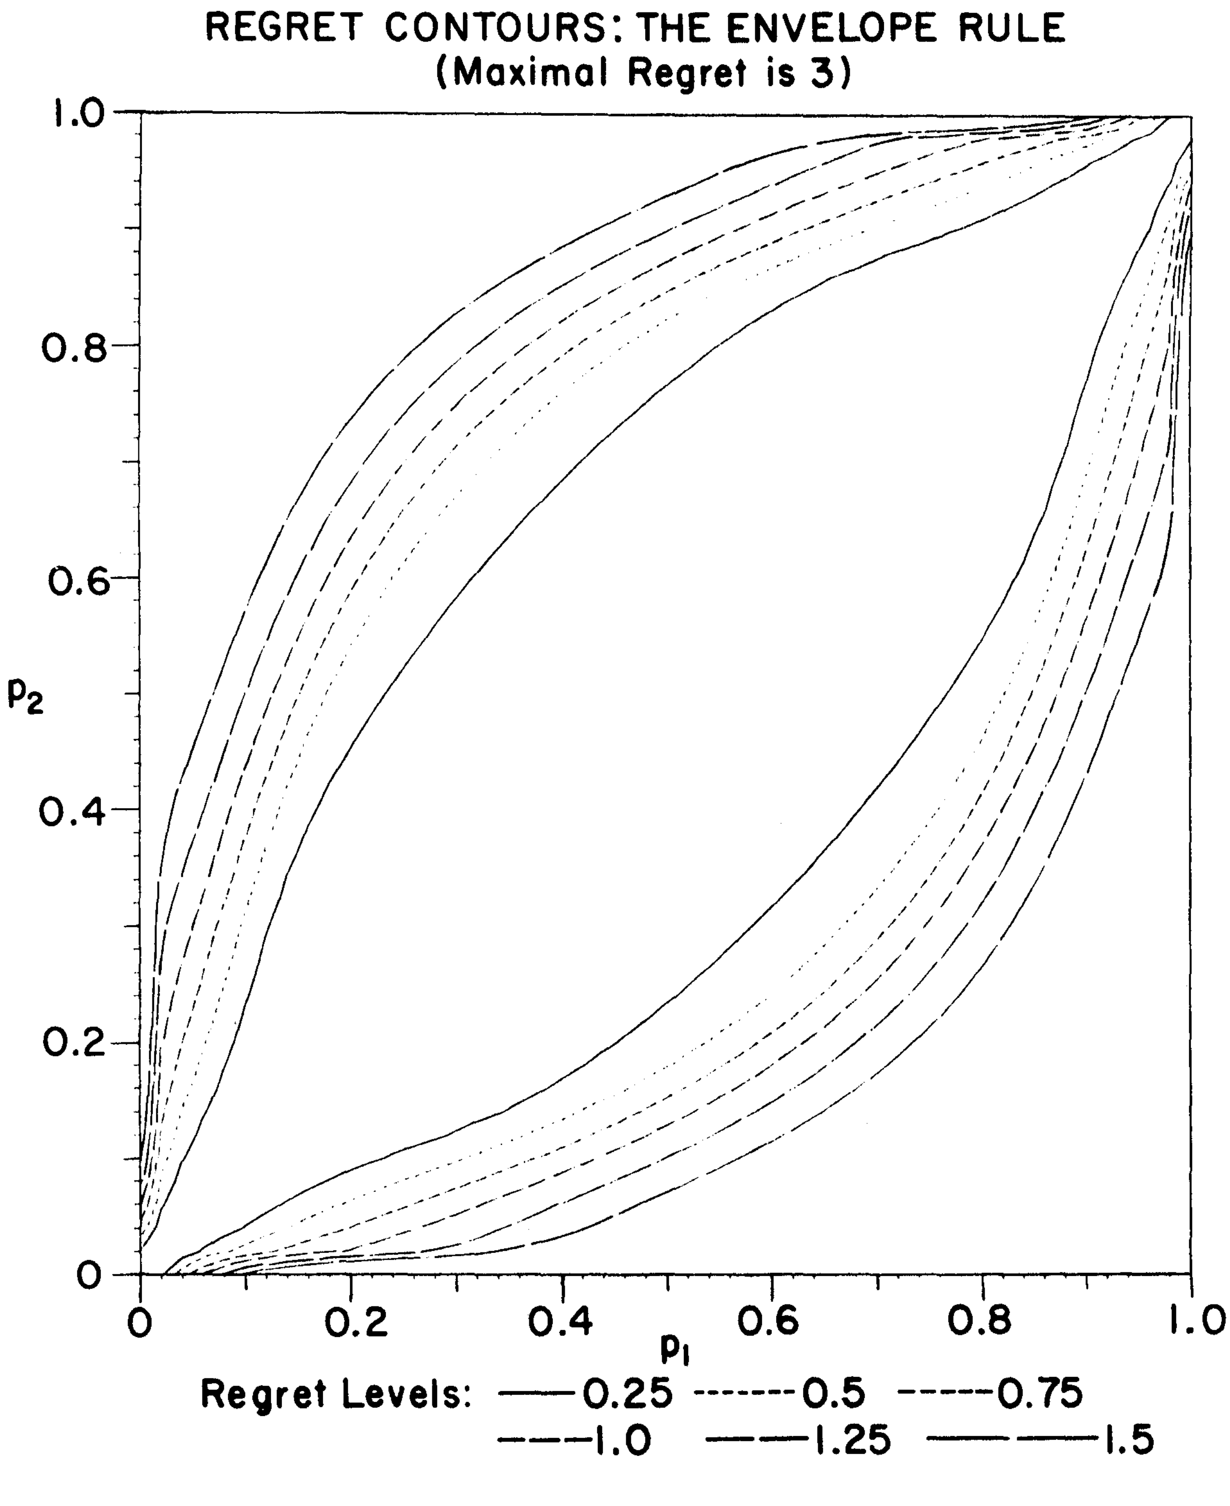
REGRET CONTOURS: THE ENVELOPE RULE
(Maximal Regret is 3)
1.0
0.8
0.6
0.4
0.2
0
p_2
0
0.2
0.4
0.6
0.8
1.0
p_1
Regret Levels: 0.25 0.5 0.75
1.0 1.25 1.5

Figure I e

REGRET CONTOURS: THE ONE PER CENT RULE
(Maximal Regret is 2)

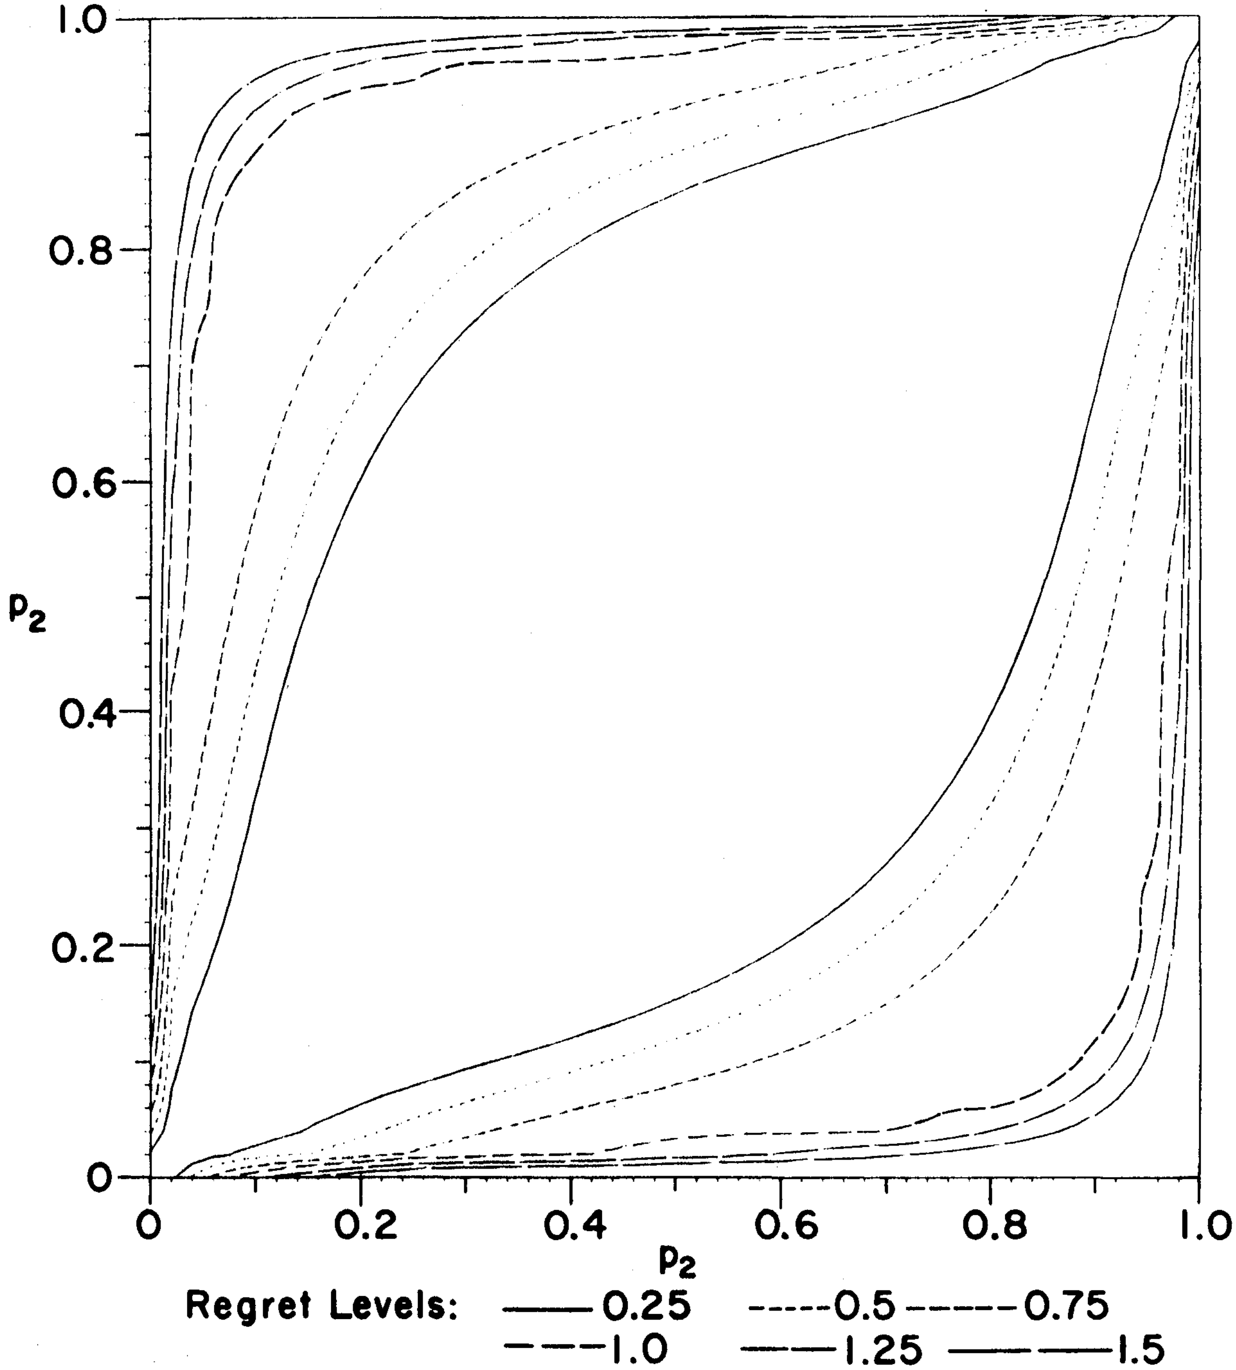

Figure I f

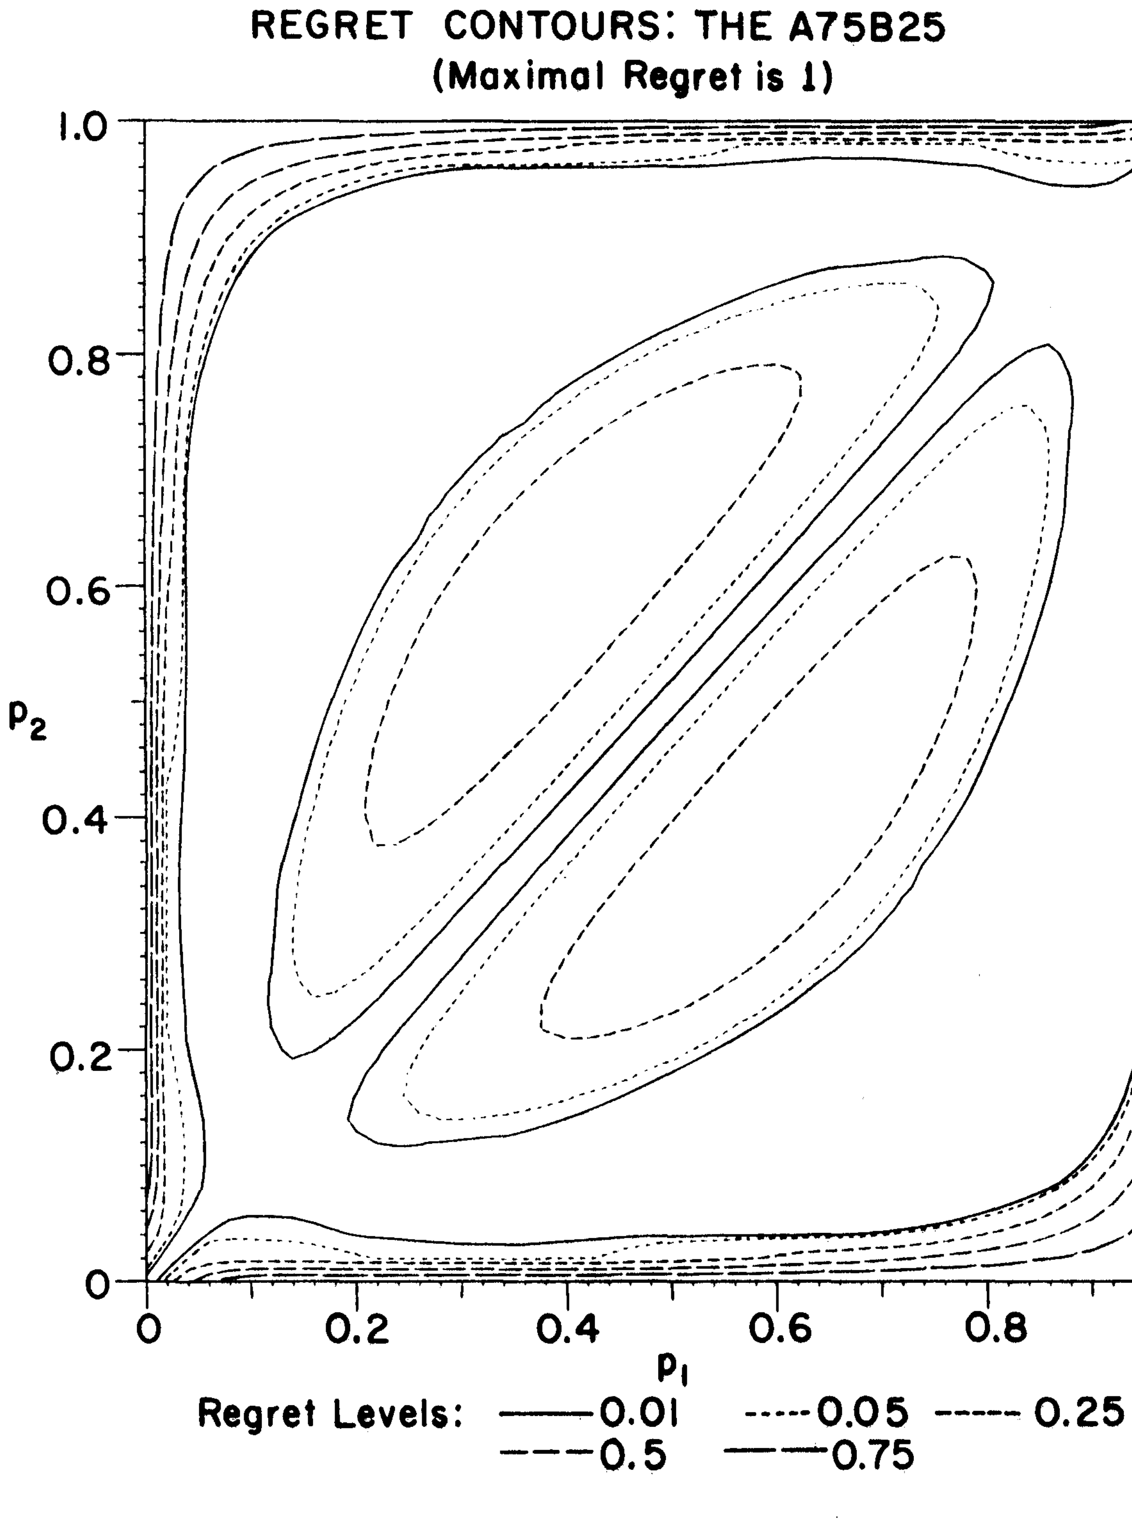
REGRET CONTOURS: THE A75B25
(Maximal Regret is 1)
1.0
0.8
0.6
0.4
0.2
0
p_2
0
0.2
0.4
0.6
0.8
1.0
p_1
Regret Levels: 0.01 0.05 0.25
0.5 0.75

Figure I g

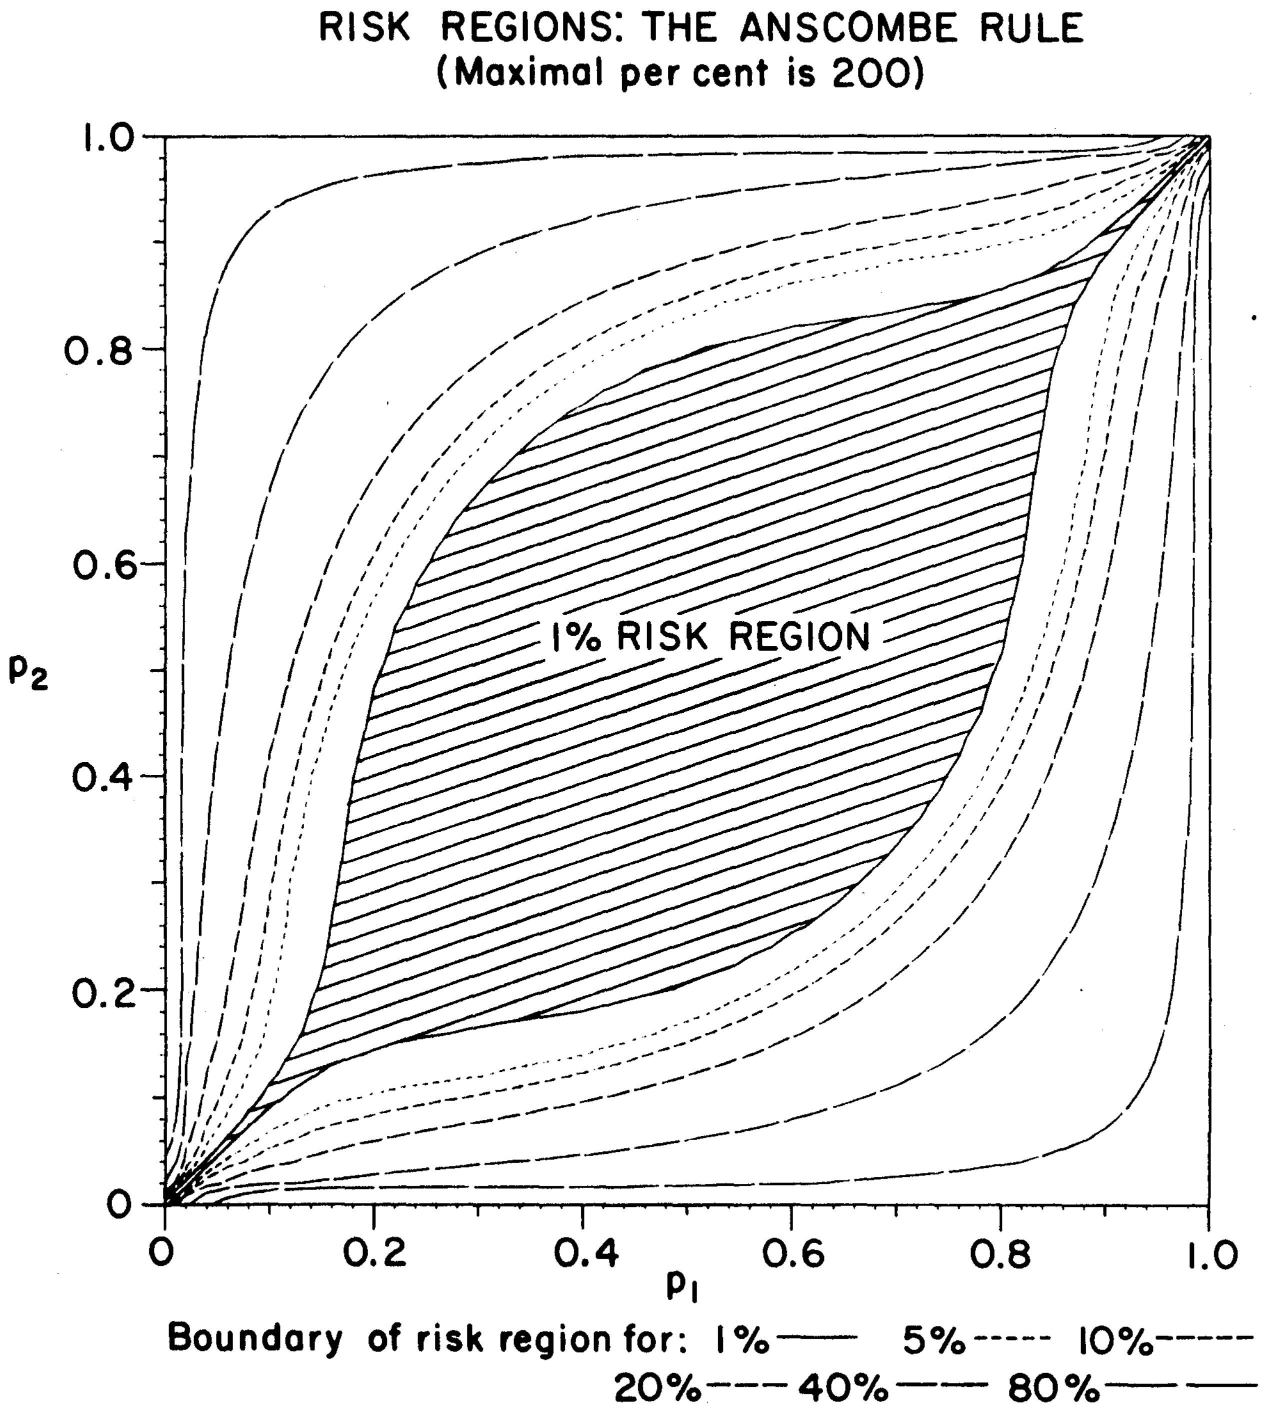
RISK REGIONS: THE ANSCOMBE RULE
(Maximal per cent is 200)
1.0
0.8
0.6
0.4
0.2
0
p_2
1% RISK REGION
0
0.2
0.4
0.6
0.8
1.0
p_1
Boundary of risk region for: 1% 5% 10%
20% 40% 80%

Figure II a

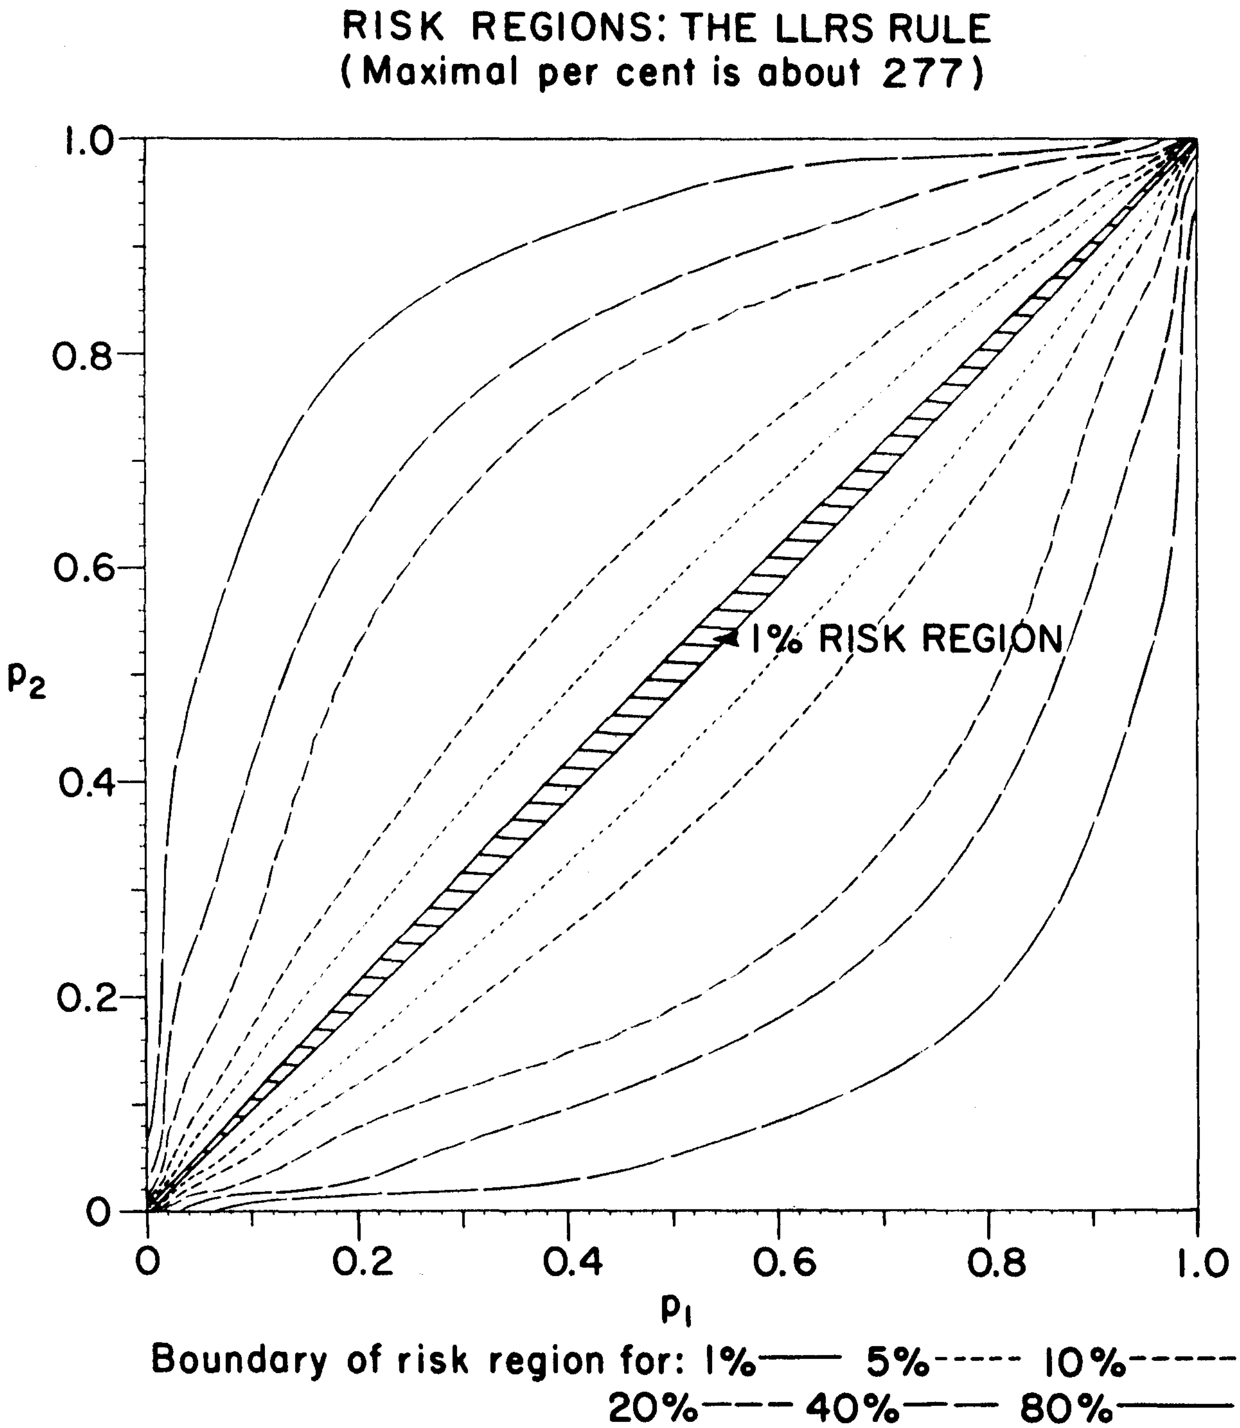
RISK REGIONS: THE LLRS RULE
(Maximal per cent is about 277)
1.0
0.8
0.6
0.4
0.2
0
p_2
1% RISK REGION
0
0.2
0.4
0.6
0.8
1.0
p_1
Boundary of risk region for: 1% 5% 10%
20% 40% 80%

Figure II b

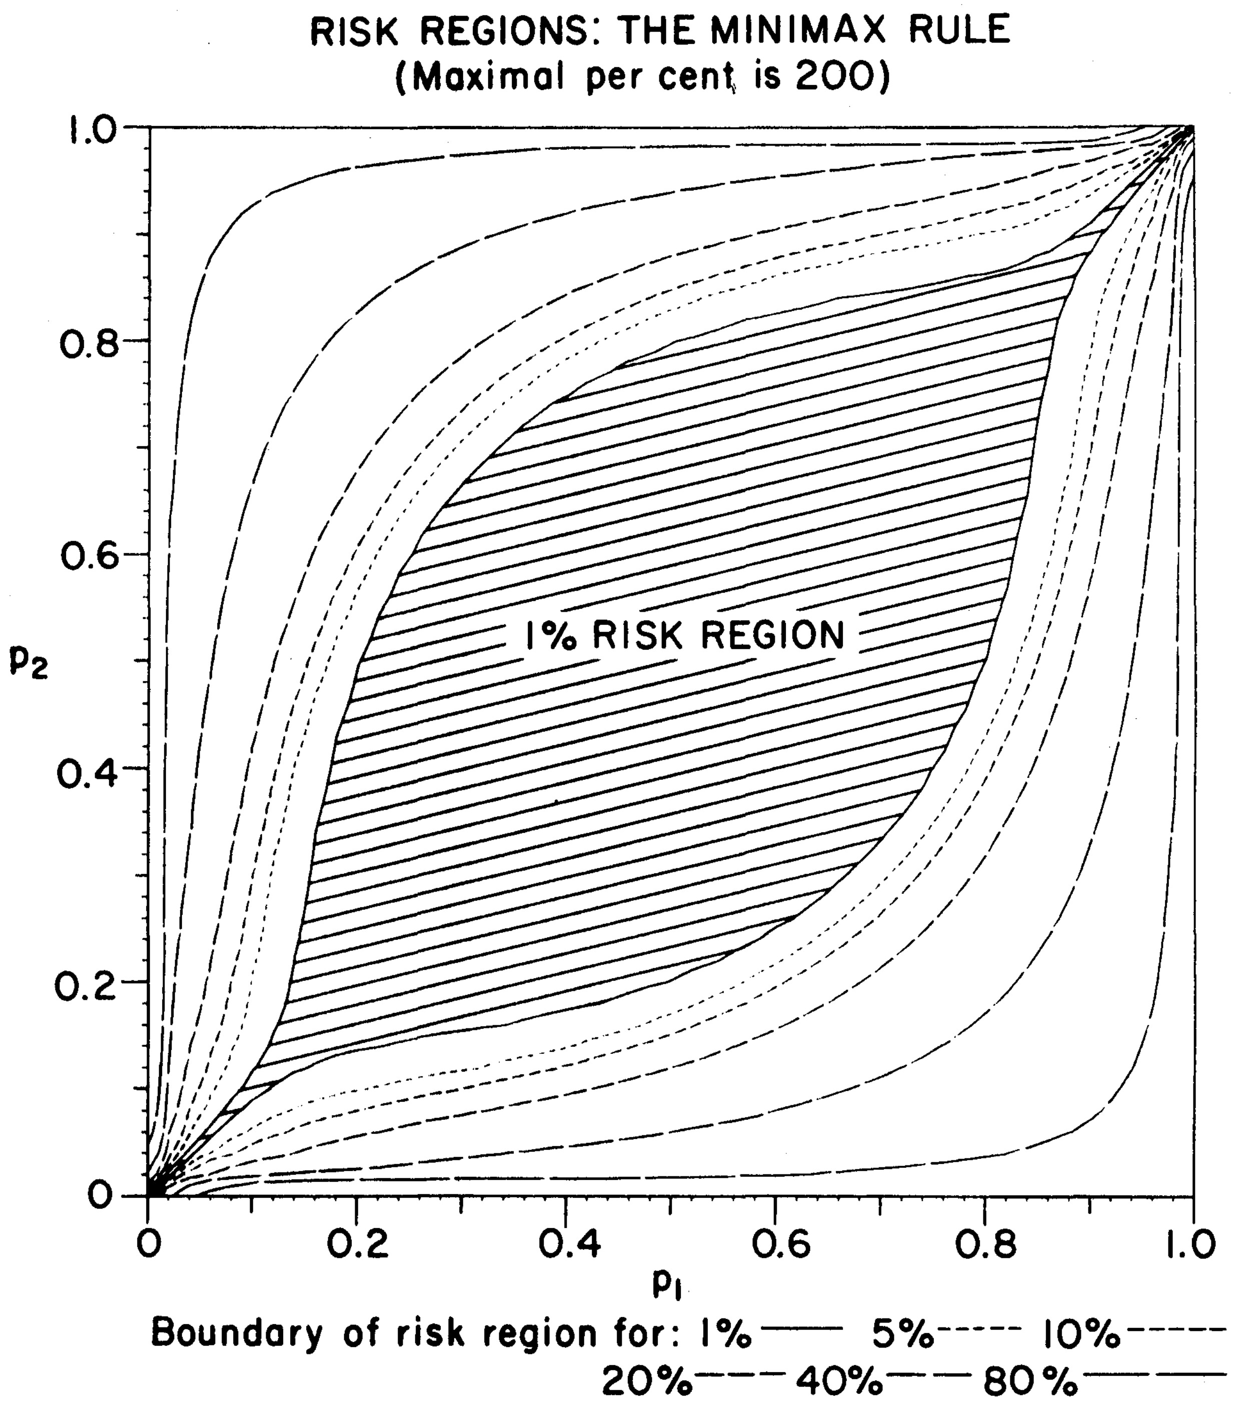
RISK REGIONS: THE MINIMAX RULE
(Maximal per cent is 200)
1.0
0.8
0.6
0.4
0.2
0
p2
1% RISK REGION
0
0.2
0.4
0.6
0.8
1.0
p1
Boundary of risk region for: 1% 5% 10%
20% 40% 80%

Figure II c

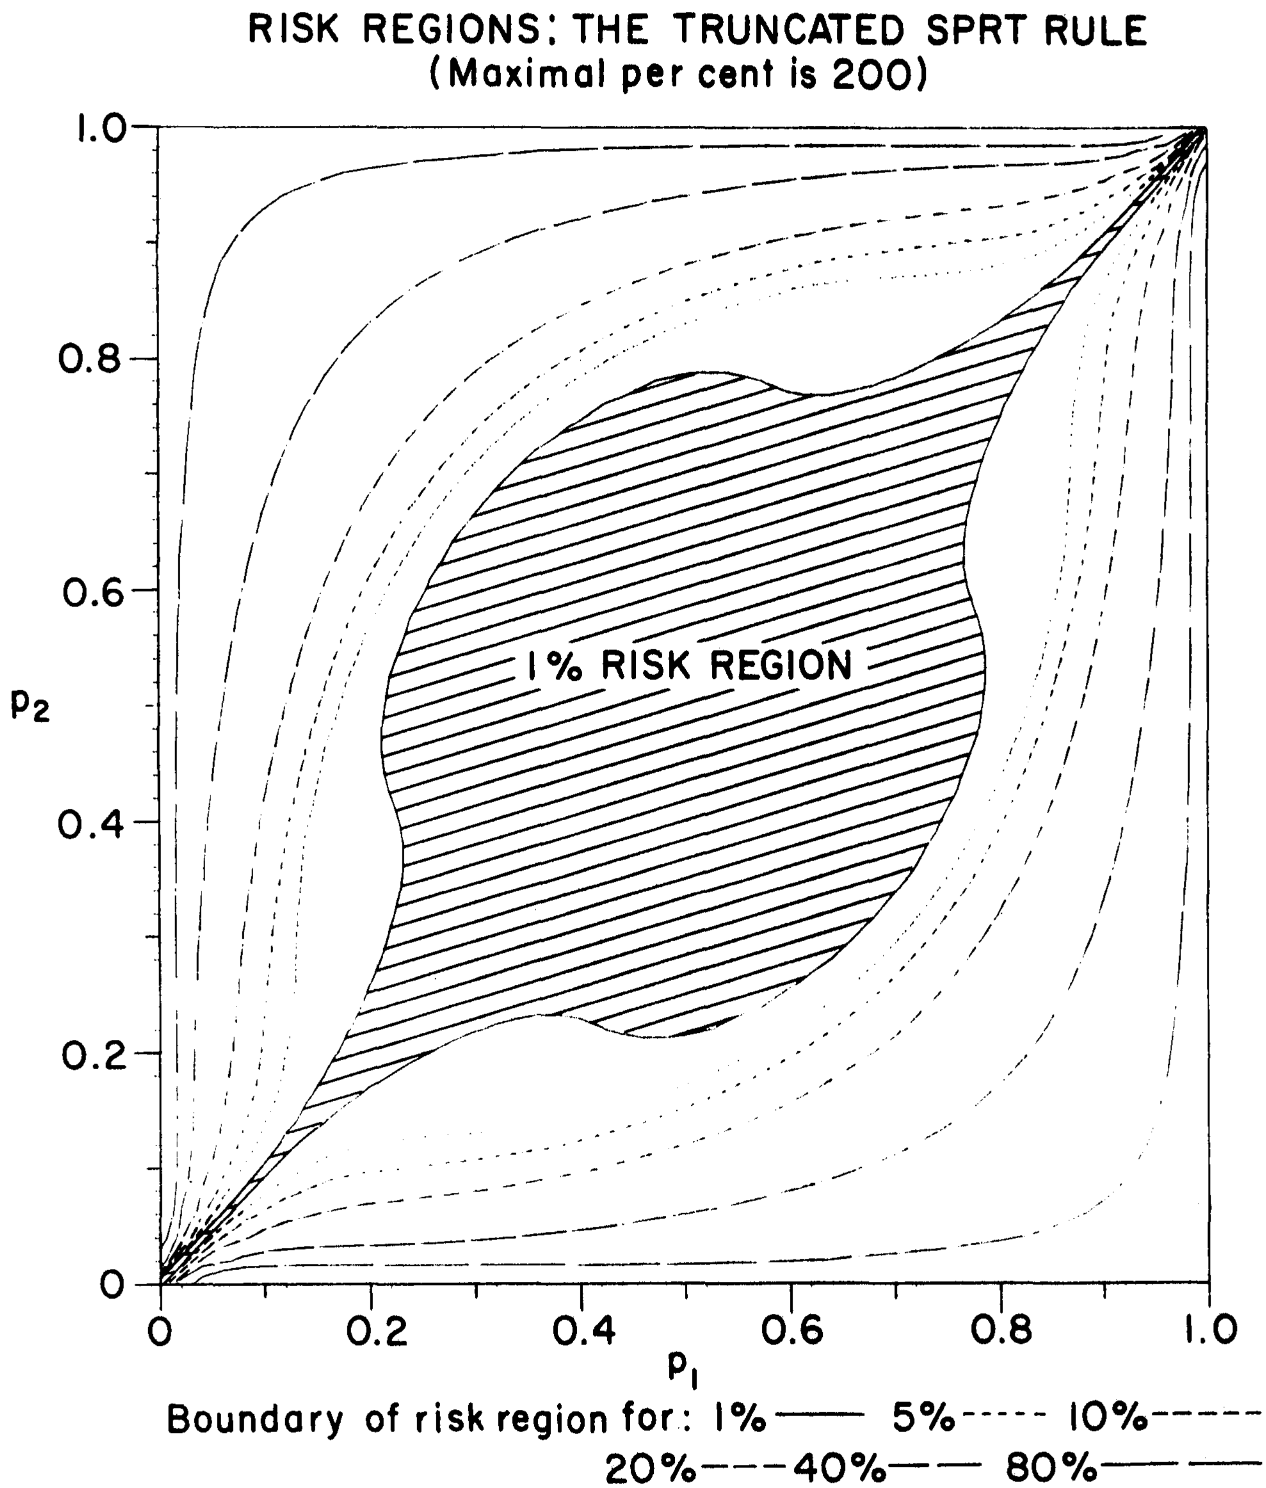
RISK REGIONS: THE TRUNCATED SPRT RULE
(Maximal per cent is 200)
1.0
0.8
0.6
0.4
0.2
0
p_2
1% RISK REGION
0
0.2
0.4
0.6
0.8
1.0
p_1
Boundary of risk region for: 1% 5% 10%
20% 40% 80%

Figure II d

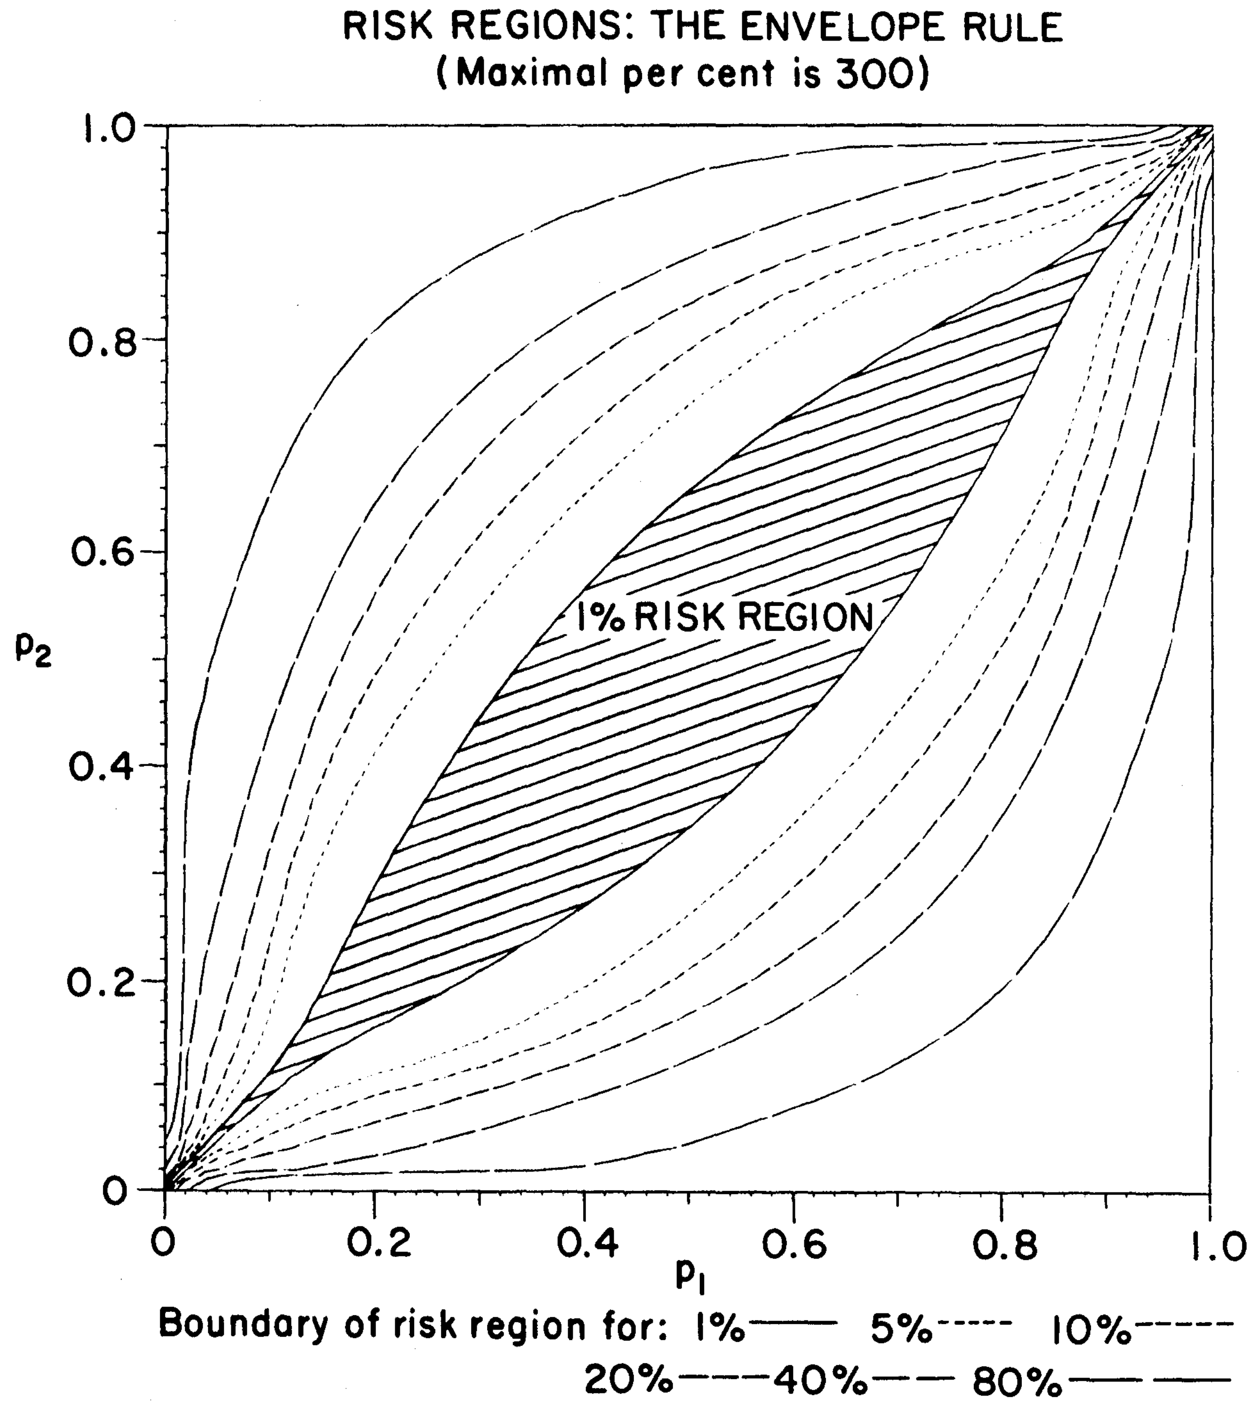
RISK REGIONS: THE ENVELOPE RULE
(Maximal per cent is 300)
1.0
0.8
0.6
0.4
0.2
0
p2
1% RISK REGION
0
0.2
0.4
0.6
0.8
1.0
p1
Boundary of risk region for: 1% 5% 10%
20% 40% 80%

Figure II e

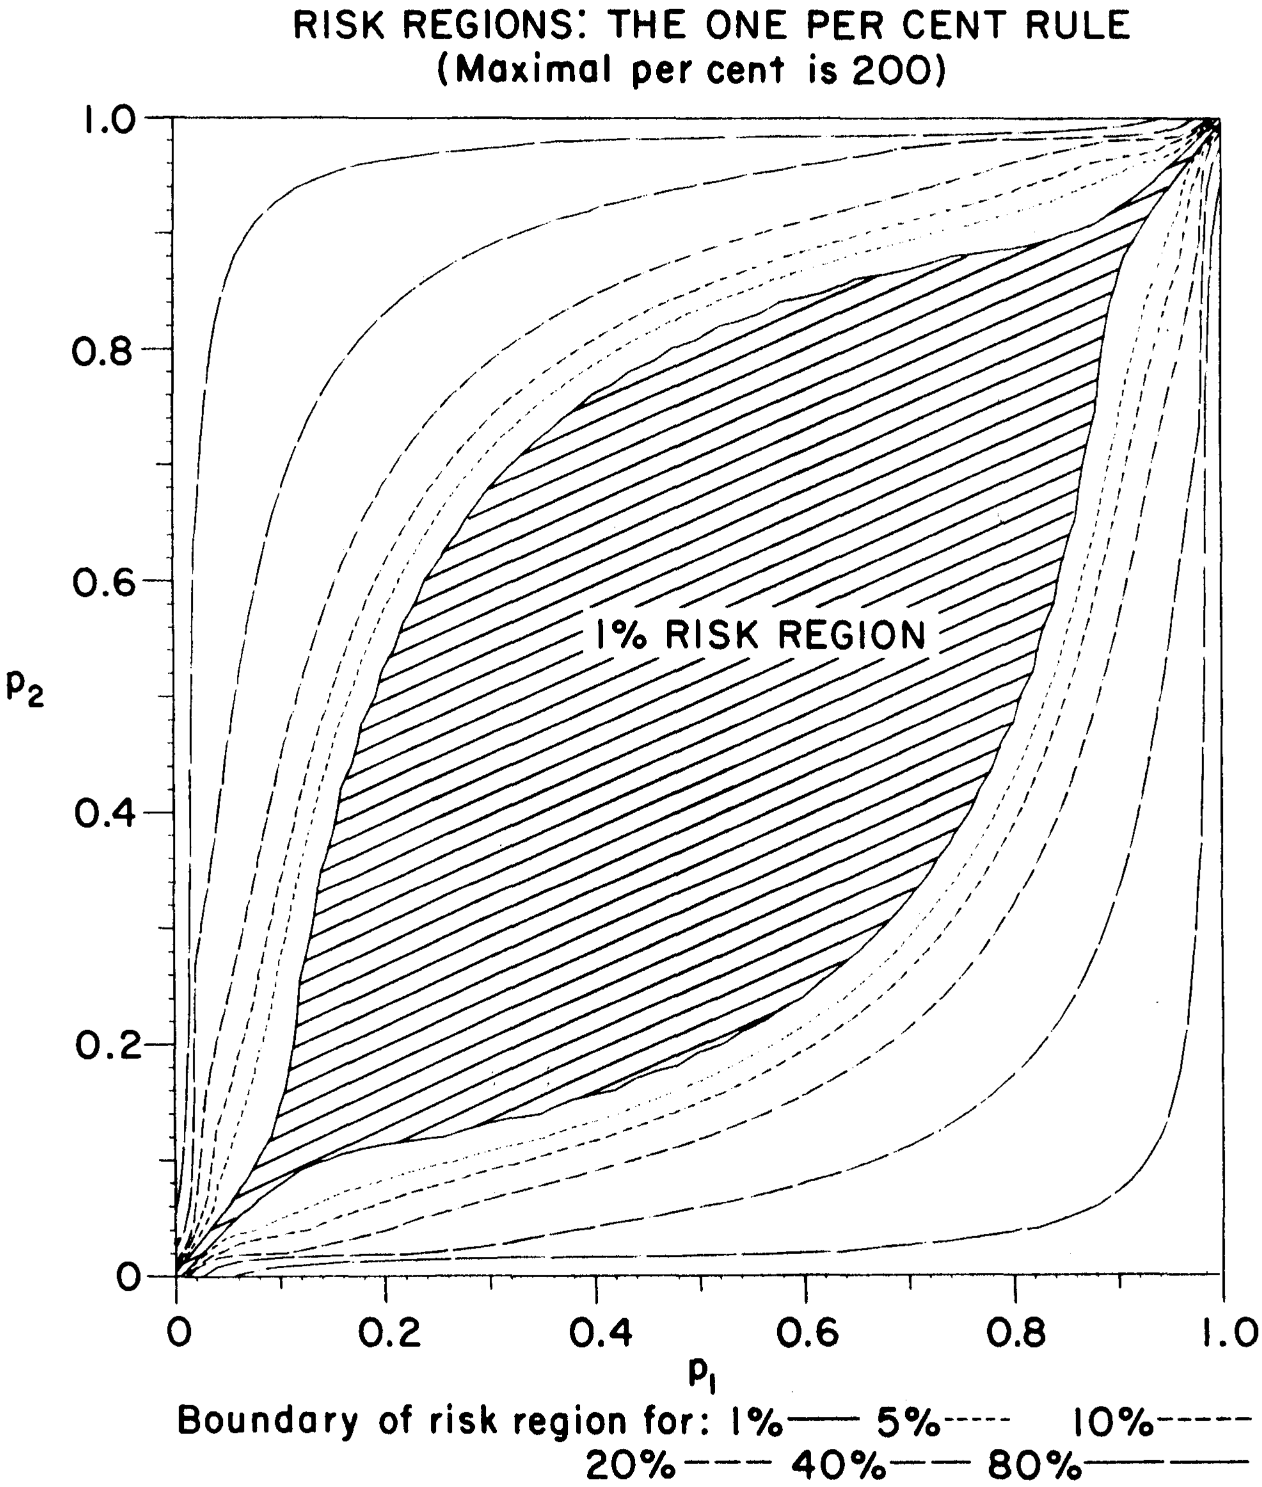
RISK REGIONS: THE ONE PER CENT RULE
(Maximal per cent is 200)
1.0
0.8
0.6
0.4
0.2
0
p_2
1% RISK REGION
0
0.2
0.4
0.6
0.8
1.0
p_1
Boundary of risk region for: 1% 5% 10%
20% 40% 80%

Figure II f

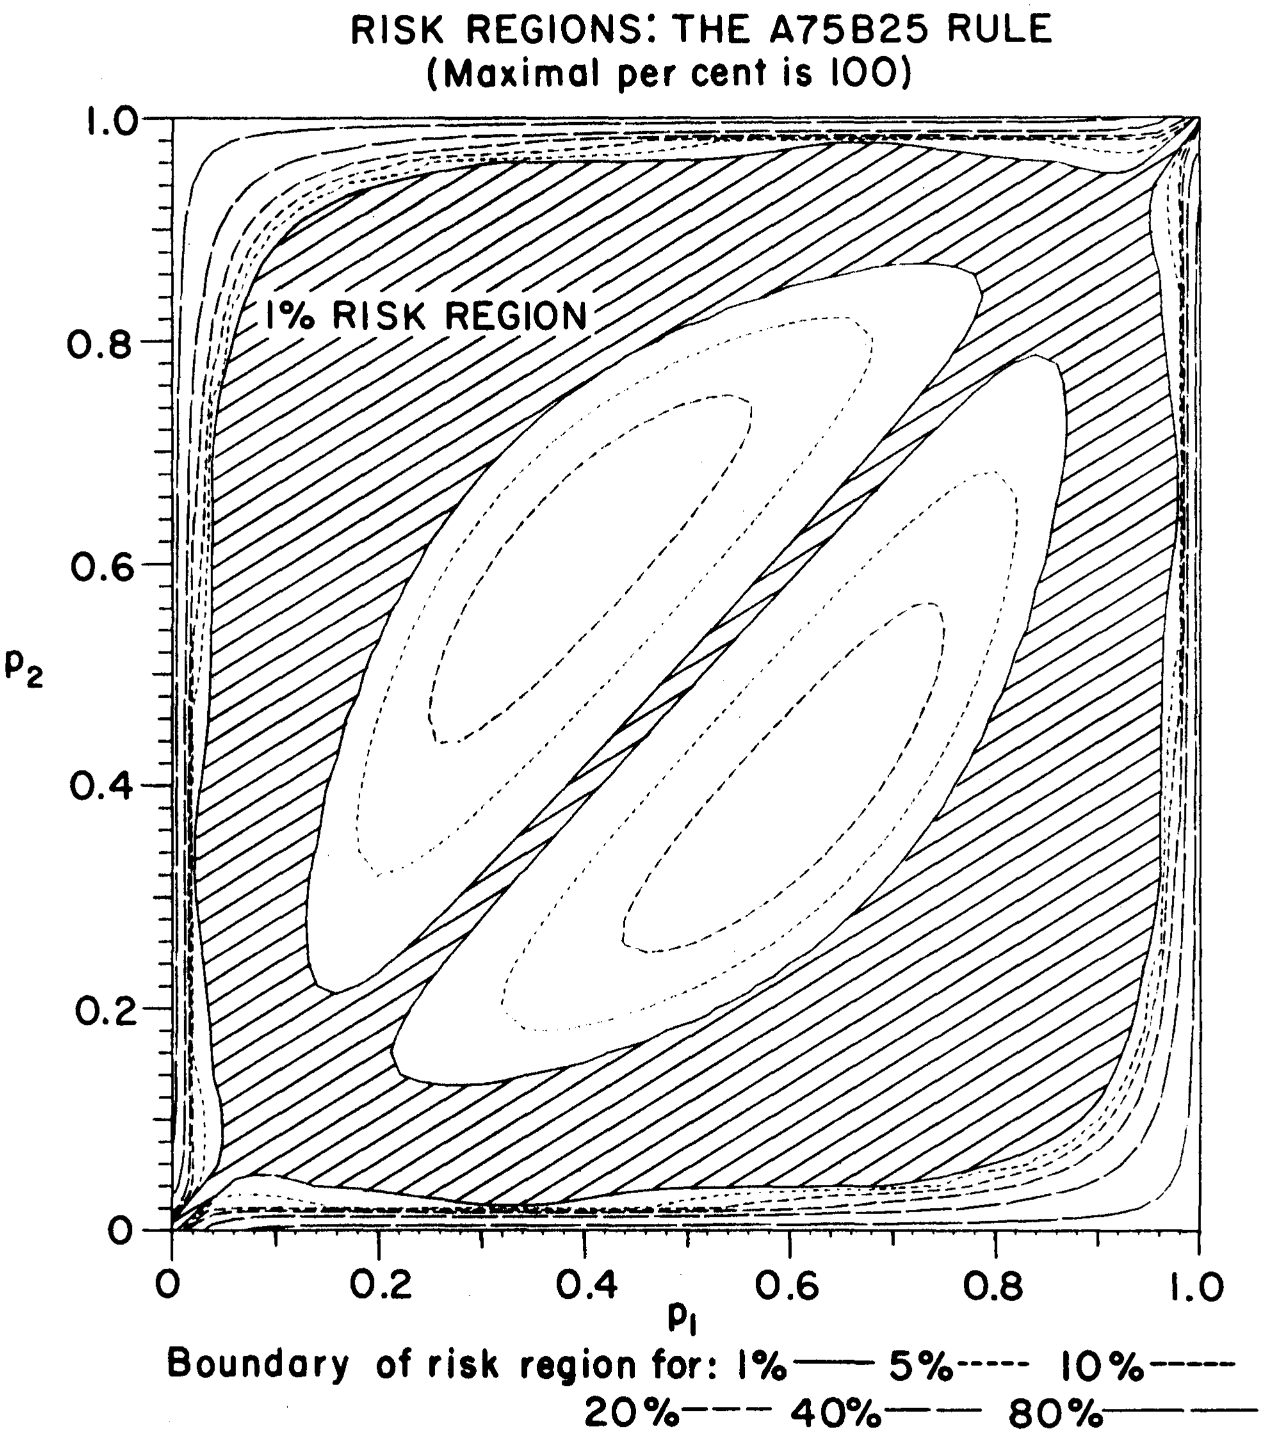
RISK REGIONS: THE A75B25 RULE
(Maximal per cent is 100)
1% RISK REGION
1.0
0.8
0.6
0.4
0.2
0
0
0.2
0.4
0.6
0.8
1.0
p_2
p_1
Boundary of risk region for: 1% 5% 10%
20% 40% 80%

Figure II g

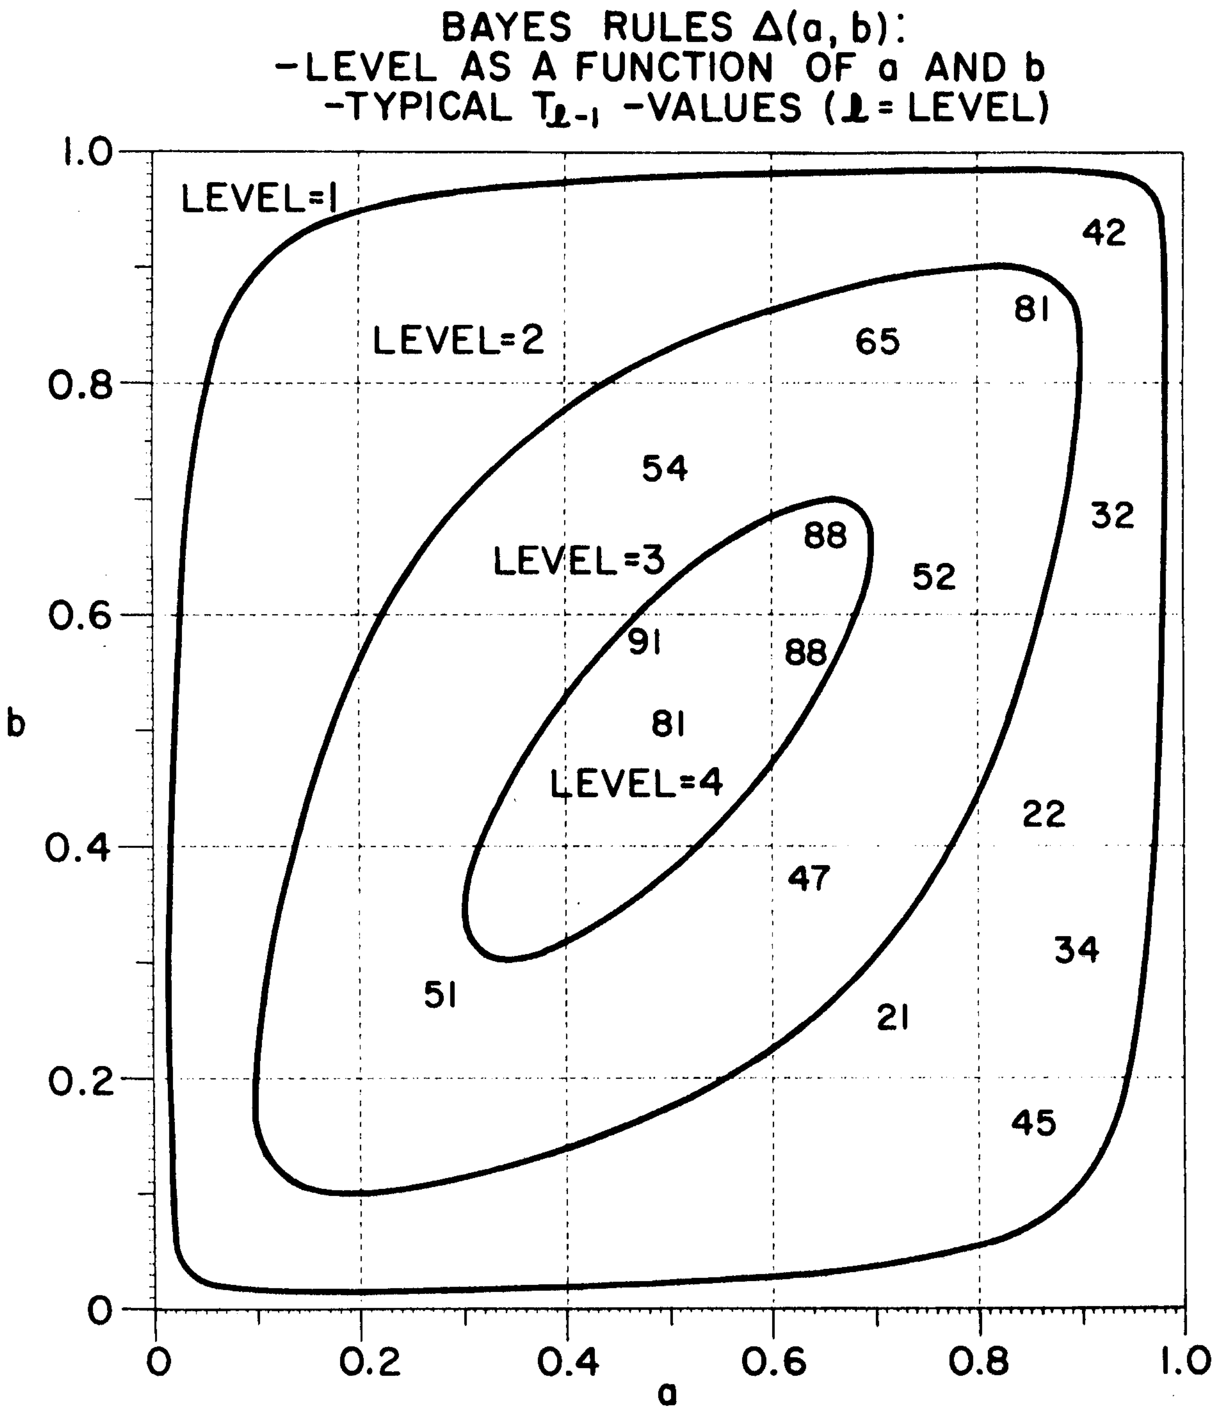
BAYES RULES Δ(a, b):
-LEVEL AS A FUNCTION OF a AND b
-TYPICAL $T_{\ell-1}$ -VALUES (ℓ = LEVEL)
LEVEL=1
LEVEL=2
LEVEL=3
LEVEL=4
42
81
65
54
32
88
52
91
88
81
22
47
34
51
21
45
b
a
1.0
0.8
0.6
0.4
0.2
0
0
0.2
0.4
0.6
0.8
1.0

Figure III

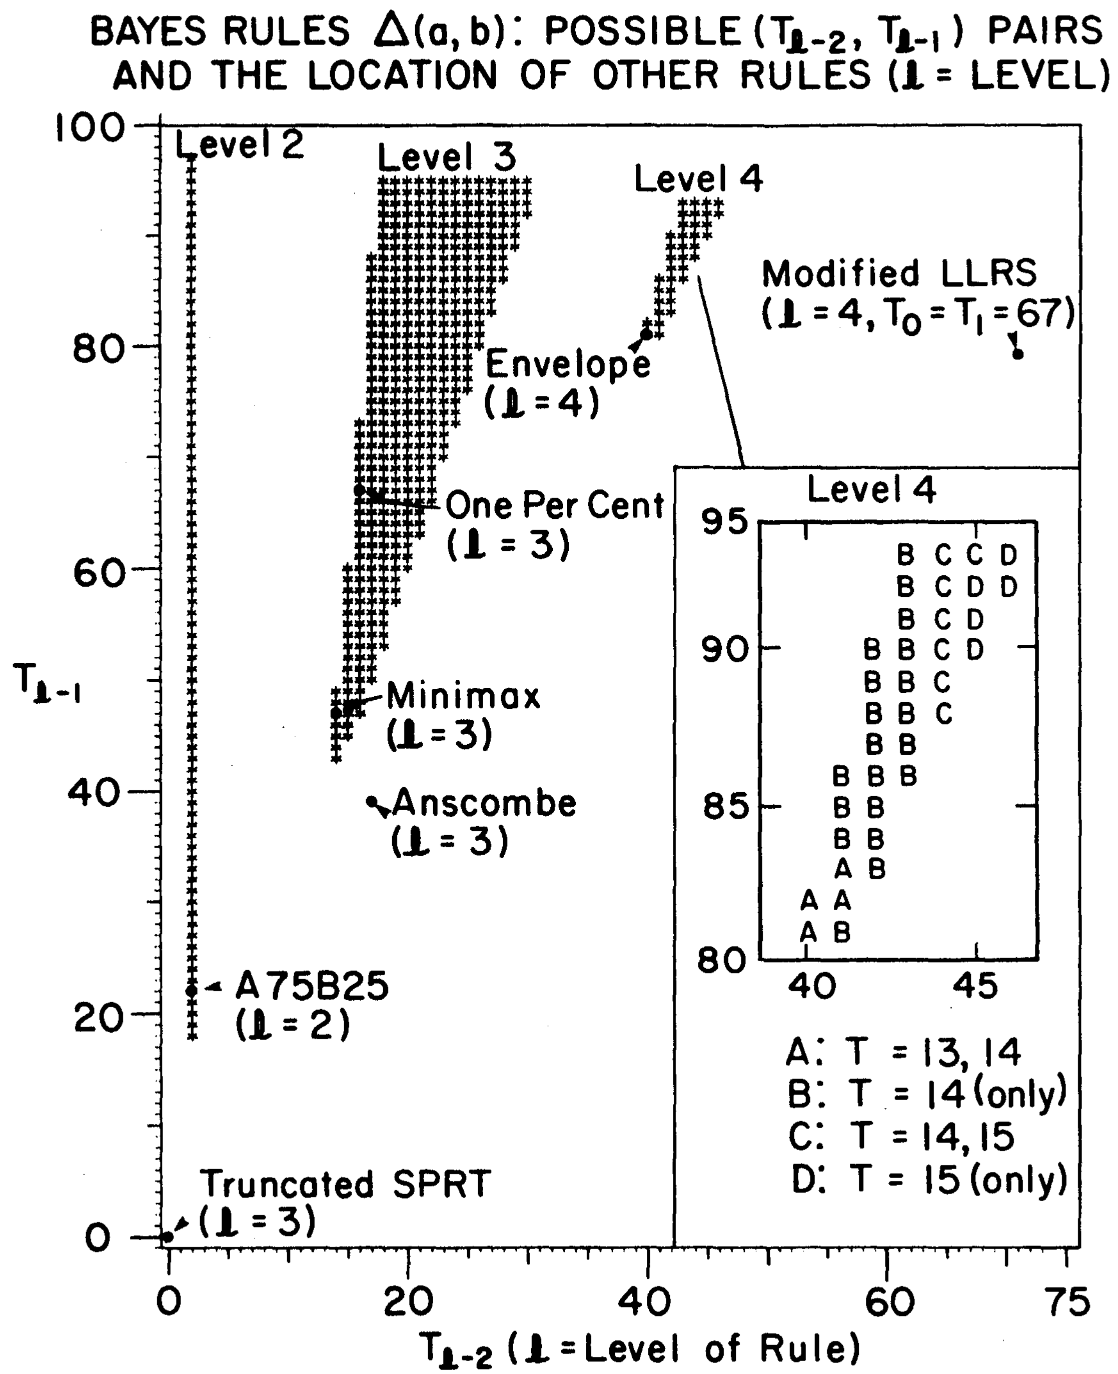
BAYES RULES Δ(a, b): POSSIBLE ($T_{\ell-2}$, $T_{\ell-1}$) PAIRS
AND THE LOCATION OF OTHER RULES (ℓ = LEVEL)
Level 2
Level 3
Level 4
Modified LLRS
(ℓ = 4, T_0 = T_1 = 67)
Envelope
(ℓ = 4)
One Per Cent
(ℓ = 3)
Minimax
(ℓ = 3)
Anscombe
(ℓ = 3)
A75B25
(ℓ = 2)
Truncated SPRT
(ℓ = 3)
$T_{\ell-1}$
$T_{\ell-2}$ (ℓ = Level of Rule)
Level 4
A: T = 13, 14
B: T = 14 (only)
C: T = 14, 15
D: T = 15 (only)

Figure IV

BAYES RULES Δ(a, b): MINIMAL BAYES RISKS
(Maximum of Minimal Bayes Risks is about 3.7)

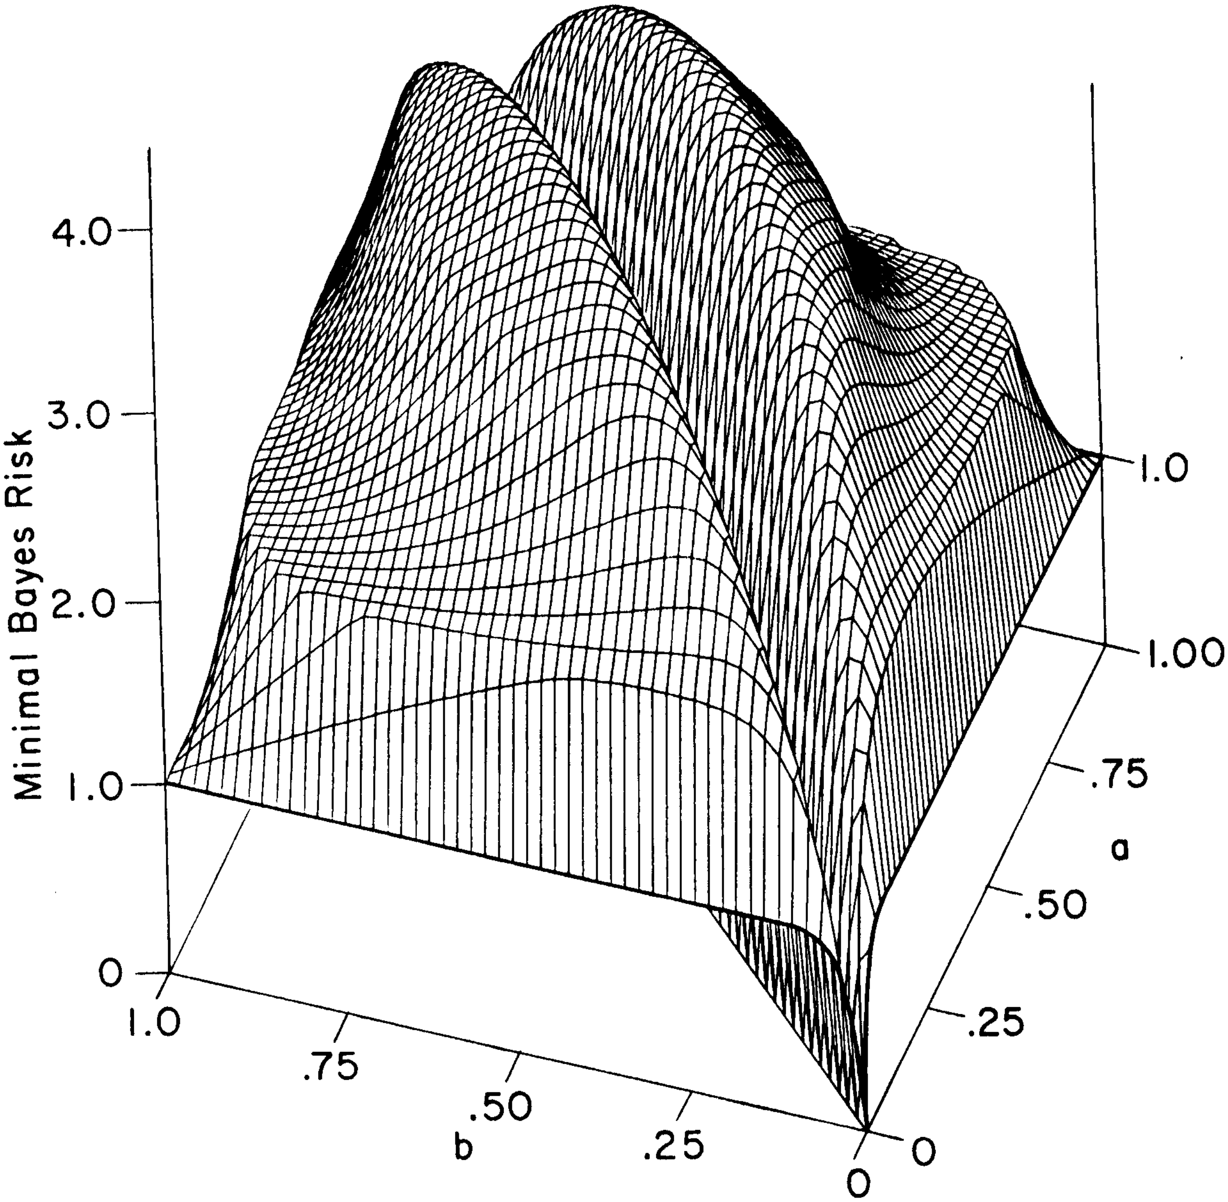

Figure V

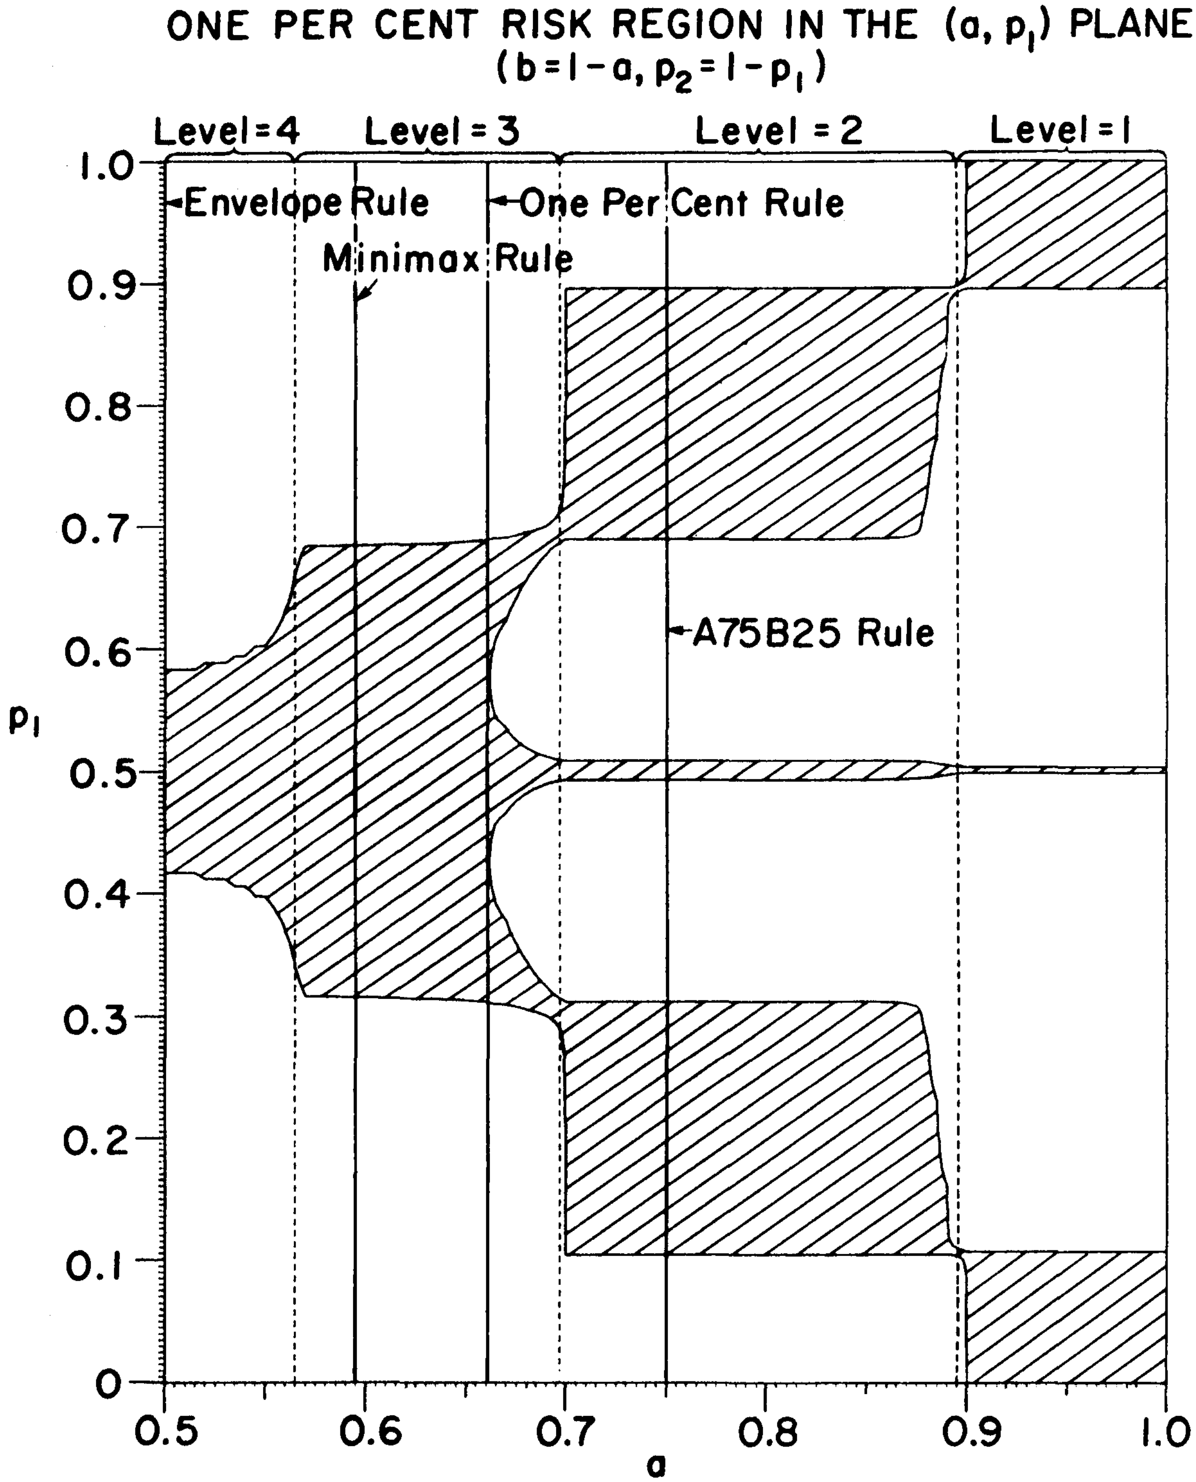
ONE PER CENT RISK REGION IN THE (a, p_1) PLANE
(b = 1 − a, p_2 = 1 − p_1)
Level = 4
Level = 3
Level = 2
Level = 1
Envelope Rule
One Per Cent Rule
Minimax Rule
A75B25 Rule
p_1
a
1.0
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
0.5
0.6
0.7
0.8
0.9
1.0

Figure VI