ESTIMATING A BINOMIAL PARAMETER: IS ROBUST BAYES REAL BAYES?

by

Mei-Mei Zen and Anirban DasGupta
Purdue University

Technical Report #90-63C

Department of Statistics
Purdue University

November 1990
Revised December 1991

Estimating a Binomial Parameter:
Is Robust Bayes Real Bayes?*

Abstract

In robust Bayesian analysis, a prior is assumed to belong to a family instead of being specified exactly. The multiplicity of priors naturally leads to a collection of Bayes actions (estimates), and these often form a convex set (an interval in the case of a real parameter). It is clearly essential to be able to recommend one action from this set to the user. We address the following problem: if we systematically choose one action for each X thereby constructing a decision rule, is it going to be Bayes? Is it Bayes with respect to a prior in the original prior family? Even if it is not genuine Bayes, is it admissible?

This problem is addressed in the context of estimating an unknown Binomial parameter. Several prior families are considered. We look at the midpoint of the interval of Bayes estimates; this has a minimax interpretation, apart from its obvious simplistic appeal. We establish that unless the prior family includes unreasonable priors, use of this estimate guarantees good behavior and indeed it is usually admissible or even genuine Bayes.

1. Introduction

1.1 The Basic Problem. In the robust Bayesian viewpoint, the parameter θ of a random variable X is assumed to have a prior distribution belonging to an appropriate family of priors, say, Γ. Considerable work has now been done in this area on the aspect of sensitivity of posterior expectations to the choice of the prior. See Berger (1990) for a general exposition and references to other works. The typical conclusion in all of these works is that if the body of the prior(s) and that of the likelihood match, then posterior means and such other measures are moderately insensitive to the tail of the prior and otherwise they are not. In a broad sense, therefore, the central phenomenon in these problems is now quite well understood.

A related problem of interest to frequentists and Bayesians alike is the following: if, from the set of all Bayes actions corresponding to the priors in the family Γ, a particular

* Research supported by NSF Grant DMS 89-230-71

action is chosen in a systematic way for each value of the sample X, what would be the behavior of the resulting decision rule? A clearly natural question for a Bayesian here is if this procedure actually corresponds to valid prior opinion; to put it in the language of decision theory, one likes to know if there is a prior $g(\theta)$ (hopefully itself belonging to the family Γ) such that the espoused rule is Bayes with respect to $g(\theta)$. Note we require that g be a bonafide prior; that is, g should not depend on the data X. We thus not only want the procedure to be conditionally Bayes for each fixed X, we further want the same prior to work for every X. In such a case, we are satisfied that our 'robust Bayes' procedure actually corresponds to prior opinion. In the frequentist analysis, the similar question of interest is the risk behavior of the procedure. In particular, one wants to know if it is admissible. Notice that the two questions raised here are related because of the well known relations between Bayes and admissible procedures, although they are not (usually) completely equivalent. The questions raised here are of clear practical importance, since in the final analysis we cannot recommend a range of actions to an untrained user and specific actions must be advised.

1.2. The Selection Criterion. The literature on robust Bayes analysis does not as yet contain substantial work on how exactly a specific action should be chosen. It may even be the case that no single method works well all the time. Indeed, this is very likely. In DasGupta and Studden (1989), several methods for choosing a specific action are proposed and the resulting procedures are derived and evaluated for estimating a multivariate normal mean. In this article, we will use posterior regret as the selection criterion and use the corresponding minimax action. Specifically, given fixed X, a fixed action δ and a fixed prior π in Γ, let $r(\pi, \delta, X)$ denote the posterior expected loss of δ under π and let $r(\pi, X) = \inf r(\pi, \delta, X)$; the posterior regret is defined to be $r^*(\pi, \delta, X) = r(\pi, \delta, X) - r(\pi, X)$. The minimax action $\delta_0 = \delta_0(X)$ then minimizes $\sup_\pi r^*(\pi, \delta, X)$. If the basic loss in the decision problem is ordinary squared error and if the range of actions for π in Γ is an interval, then this minimax action simply corresponds to the midpoint of the interval of Bayes actions. Note that the use of the midpoint may seem to be rather naive at first glance. But we will in fact give evidence that in the problem we consider, this simple action often turns out to be itself Bayes and also admissible. More comments about similarities with the normal problem are made in section 3.

1.3. The Specific Problem. The problems we describe above are relevant in arbitrary decision problems in any dimension. Since we believe that this is the first work on this issue, we confine ourselves to the one dimensional problem of estimating an unknown Binomial success probability. A technical reason for considering this problem is that usable characterizations of admissible and Bayes procedures are available here; see the works of Johnson (1971), Kozek (1982), Skibinsky and Rukhin (1989) and Brown (1981).

Formally, then, one has $X \sim$ Bin (n, θ), $0 \leq \theta \leq 1$. θ has a prior π; five different families of priors will be considered:

$$\begin{aligned}
\Gamma_1 &= \text{ set of all symmetric Beta } (\alpha, \alpha) \text{ priors with } \alpha_1 \leq \alpha \leq \alpha_2; \\
\Gamma_2 &= \text{ set of all priors symmetric and unimodal about } \frac{1}{2}; \\
\Gamma_3 &= \left\{ \pi : \ \pi(\theta) = (1-\varepsilon)\pi_0(\theta) + \varepsilon q(\theta), \right. \\
&\qquad \left. \text{where } 0 < \varepsilon < 1, \pi_0(\theta) \text{ are fixed, and } \pi_0, \ q \in \Gamma_2 \right\}; \\
\Gamma_4 &= \text{ set of all priors symmetric about } \frac{1}{2}; \\
\text{and } \Gamma_5 &= \text{ set of all Beta } (\alpha, \beta) \text{ priors with } (\alpha, \beta) \text{ belonging to a} \\
&\quad\ \text{suitable bounded triangle in the two dimensional plane.}
\end{aligned}$$

1.4. Discussion on the Families of Priors. First note that Γ_i for $i = 1, 2, 3, 4$ only contain symmetric prior distributions. If asymmetric priors are considered plausible, Γ_5 is a possible choice.

Γ_1 is a convenient choice as far as calculations are concerned; the bounds on α can be justified by proving that these correspond to imposing bounds on any one or a fixed number of quantiles of the prior. Note that unlike in the situation of estimating a normal mean, use of conjugate priors is not bad here because of the true richness of conjugate priors in this case.

Use of Γ_2 is an attempt to include nonconjugate priors while preserving reasonable shape features; Γ_3 is like Γ_2 but is less conservative since π_0 is kept fixed. Many authors have used Γ_3 or families like Γ_3 in various robustness works; Huber (1973), Berger and

Sivaganesan (1989), DasGupta and Bose (1988), DasGupta and Delampady (1990), DasGupta (1990) are some references here. Γ_4 is the largest of the first four families and is really quite conservative. Note that $\Gamma_1, \Gamma_3 \subseteq \Gamma_2 \subseteq \Gamma_4$ but neither of Γ_1 and Γ_3 contains the other. Regarding Γ_5, it will later be demonstrated that varying (α, β) in a bounded triangle implies that the mean and the variance of the prior are bounded in both directions. Since priors which allow the parameter to be a fixed constant with a very large probability are generally not thought of as highly plausible priors in problems where robustness is a concern, a lower bound on the variance is very natural. An upper bound on the variance is automatic since θ is bounded. More will be said of Γ_5 later.

1.5. Description of Results. It is proved that the midpoint equals the Bayes estimate corresponding to a Beta prior (depending on n) for Γ_1 and Γ_2; thus it is Bayes as well as admissible. Indeed, if in Γ_1 one takes $\alpha_1 = 1$ and $\alpha_2 = \infty$ (in which case $\alpha = \alpha_2$ is not allowed), the midpoint for Γ_1 and Γ_2 coincide. Notice that with $\alpha_1 = 1$ and $\alpha_2 = \infty$, Γ_1 becomes the collection of all symmetric unimodal Beta priors. For the family Γ_4, the midpoint is shown to be a piecewise constant rule equaling $\frac{1}{4}$ for $X < [\frac{n}{2}]$, $\frac{1}{2}$ for $X = [\frac{n}{2}]$ and $\frac{3}{4}$ for $X > [\frac{n}{2}]$. This is neither Bayes nor admissible for $n > 2$. The midpoint is explicitly evaluated for $n \leq 4$ for the family Γ_3 and shown to be Bayes and admissible for any ε and any π_0. For $n > 4$, explicit evaluation of the estimate becomes difficult, but we have numerical evidence using characterizations of admissible estimates that for any ε and π_0, the midpoint is admissible for all n. A similar phenomenon holds for the family Γ_5. Here also we prove analytically the Bayes nature of the estimate for $n \leq 4$. Interestingly, even though in most of the cases where the midpoint is Bayes, it is Bayes with respect to (infinitely) many priors many of which are 'nice' plausible priors belonging to the family in consideration, in quite a few cases it is Bayes with respect to the unique distribution which is the upper or the lower principal representation of a moment problem, i.e., the unique discrete distribution which maximizes or minimizes a particular moment for given values of lower order moments. In these cases, the prior is finitely supported and does not belong to the original family under consideration.

Throughout the entire article, the midpoint estimate will be referred to as $\delta_M(X)$.

2. Main Results

2.1 The Family Γ_1. As stated in the introduction, if $\theta \sim Be(\alpha, \alpha)$, then bounding α between α_1 and α_2 is equivalent to bounding a fixed quantile between two fixed numbers. Since intersection of a finite number of closed intervals is also a closed interval, bounding α in an interval is also equivalent to a consistent set of bounds on a finite number of quantiles. Since subjective elicitation of bounds on quantiles is considered a relatively easier part of eliciting prior knowledge, bounding α in an interval is reasonable.

Lemma 2.1.1. Let $\theta \sim Be(\alpha, \alpha)$, $\alpha > 0$. Let $Z = \theta(1-\theta)$. Then the density of Z given α has monotone likelihood ratio in Z.

Proof: Straightforward.

Lemma 2.1.2. Let $\theta \sim Be(\alpha, \alpha)$, $\alpha > 0$. Then $P_\alpha(\theta \le \tau)$ is increasing in α if $\tau \ge \frac{1}{2}$ and decreasing in α if $\tau \le \frac{1}{2}$.

Proof: Let $\tau \ge \frac{1}{2}$. Then

$$\begin{aligned} P_\alpha(\theta \le \tau) &= \frac{1}{2} + P_\alpha(\frac{1}{2} \le \theta \le \tau) \\ &= \frac{1}{2} + \frac{1}{2} P_\alpha(1-\tau \le \theta \le \tau) \\ &= \frac{1}{2} + \frac{1}{2} P_\alpha(\theta(1-\theta) \ge \tau(1-\tau)), \end{aligned}$$

from which the assertion follows by using Lemma 2.1.1.

The proof for the case $\tau \le \frac{1}{2}$ is similar.

Lemma 2.1.3. Let $\theta \sim Be(\alpha, \alpha)$. Then bounding a fixed percentile of θ in an interval is equivalent to bounding α in an interval.

Proof: Follows from Lemma 2.1.2.

Theorem 2.1.4. Let $X \sim$ Bin (n, θ) and $\theta \sim Be(\alpha, \alpha)$, where $\alpha_1 \leq \alpha \leq \alpha_2$. Then

$$\delta_M(X) = \frac{X + \frac{n(\alpha_1+\alpha_2)+4\alpha_1\alpha_2}{2(n+\alpha_1+\alpha_2)}}{n + \frac{n\alpha_1+n\alpha_2+4\alpha_1\alpha_2}{n+\alpha_1+\alpha_2}}.$$

Hence, δ_M is Bayes with respect to the $Be(\alpha^*, \alpha^*)$ prior where $\alpha^* = \frac{n(\alpha_1+\alpha_2)+4\alpha_1\alpha_2}{2(n+\alpha_1+\alpha_2)}$ and is admissible for all n.

Proof: Use the fact that for any α, the Bayes estimate with respect to the $Be(\alpha, \alpha)$ prior is $\frac{X+\alpha}{n+2\alpha}$.

Remark: Notice $\alpha_1 \leq \alpha^* \leq \alpha_2$; hence δ_M is Bayes with respect to a prior in Γ_1 itself. Also $\alpha^* \approx \frac{(\alpha_1+\alpha_2)}{2}$ for large n.

2.2 The Family Γ_2.

Theorem 2.2.1. For the family Γ_2,

$$\delta_M(X) = \frac{X + \frac{n}{2} + 2}{2n + 4}.$$

Hence, δ_M is Bayes with respect to the $Be(\frac{n}{2} + 2, \frac{n}{2} + 2)$ prior and is admissible for all n.

To prove this theorem, we need the following two lemmas.

Lemma 2.2.2. Let $X \sim g(x)$, $-\infty < x < \infty$. Suppose for $x > 0$, $T(x) = \frac{g(-x)}{g(x)}$ is nondecreasing and $T(x) \geq 1$. Then $E(X|\,|x| \leq z)$ is nonincreasing in z.

Proof: By easy algebra,

$$E(X|\,|x| \leq z) = \frac{-\int_0^z x \cdot \frac{T(x)-1}{T(x)+1}(T(x)+1)g(x)dx}{\int_0^z (T(x)+1)g(x)dx} \tag{2.1}$$

Define a new family of densities $L(x|z)$ now as

$$L(x|z) = \frac{(T(x)+1)g(x)I\ 0 < x \leq z}{\int_0^z (T(x)+1)g(x)dx} \tag{2.2}$$

$L(x|z)$ has monotone likelihood ratio in X. The result now follows from (2.1) on noting that

$$E(X|\,|x| \le z) = -\int \phi(x) L(x|z)dx, \tag{2.3}$$

where

$$\phi(x) = x \cdot \frac{T(x)-1}{T(x)+1} \tag{2.4}$$

is nondecreasing.

Lemma 2.2.3. For $0 \le z \le \frac{1}{2}$, $0 \le X \le n$, define

$$f(z) = \frac{\int_{\frac{1}{2}-z}^{\frac{1}{2}+z} \theta^{X+1}(1-\theta)^{n-X} d\theta}{\int_{\frac{1}{2}-z}^{\frac{1}{2}+z} \theta^{X}(1-\theta)^{n-X} d\theta} \qquad \text{if } 0 < z \le \frac{1}{2},$$

$$= \frac{1}{2} \qquad \text{if } z = 0. \tag{2.5}$$

Then f is monotone decreasing for $X \le \frac{n}{2}$ and monotone increasing for $X \ge \frac{n}{2}$.

The case $X \le \frac{n}{2}$ follows from Lemma 2.2.2 on noting that $f(z) = \frac{1}{2} + E(\theta - \frac{1}{2}|\,|\theta - \frac{1}{2}| \le z)$, when $\theta \sim Be(X+1, n-X+1)$. The case $X \ge \frac{n}{2}$ follows by symmetry.

Proof of Theorem 2.2.1: Denote a general prior in Γ_2 by G and the corresponding Bayes estimate by δ_G. By Theorem 2.3.1 in Sivaganesan and Berger (1989),

$$\sup_G \delta_G(X) = \sup_{0 \le z \le \frac{1}{2}} f(z)$$

and

$$\inf_G \delta_G(X) = \inf_{0 \le z \le \frac{1}{2}} f(z),$$

where f is as in Lemma 2.2.3.

Hence, by Lemma 2.2.3, for $X < [\frac{n}{2}]$,

$$\sup_G \delta_G(X) = f(0) = \frac{1}{2}$$

and

$$\inf_G \delta_G(X) = f(\frac{1}{2}) = \frac{X+1}{n+2},$$

giving the required expression for δ_M. For $X > [\frac{n}{2}]$, the proof is similar and for $X = [\frac{n}{2}]$, the proof is trivial.

Remark: A comparison of the assertions of Theorems 2.1.4 and 2.2.1 shows that δ_M is the same for the family of all symmetric unimodal priors and all symmetric unimodal Beta priors (these correspond to Γ_1 with $\alpha_1 = 1$ and $\alpha_2 = \infty$). This result is not entirely obvious since Γ_1 is not weakly dense in Γ_2.

2.3. The Family Γ_3. Technically, this is the hardest family to handle. Indeed, even an explicit formula for δ_M involves formidable calculations and we have succeeded in doing so for $n \le 4$. For these values of n and any arbitrary symmetric unimodal π_0 and any $0 \le \varepsilon \le 1$, the midpoint δ_M is an admissible Bayes procedure. Once an explicit expression for δ_M is found, the Bayes nature and the admissibility of δ_M is proved by using the characterizations of Bayes or admissible procedures as in Kozek (1982) or Skibinsky and Rukhin (1989). Since this indeed involves nontrivial calculations, we give a complete proof for the case $n = 2$ and give the main steps of the proof for $n = 3$ and 4. The details are available from the authors. First we need some notation.

Let $\mu_j^0 = E_{\pi_0}(\theta^j)$, $\mu_j = E_q(\theta^j)$ and $\alpha = \frac{\varepsilon}{1-\varepsilon}$ (the case $\varepsilon = 0$ is uninteresting and the case $\varepsilon = 1$ corresponds to the family Γ_2).

If $\theta \sim q$ where q is symmetric and unimodal then $2\theta - 1$ admits the representation

$$\theta - \frac{1}{2} = u \cdot Z,$$

where $u \sim u[-1,1], \frac{1}{2} \ge Z \ge 0$, and u, Z are independent (see Khintchine (1938)). Let $W = 2Z$, so that $0 \le W \le 1$. Consequently,

$$\mu_j = \frac{1}{2^j} \sum_{i=0}^{[\frac{j}{2}]} \binom{j}{2i} \frac{1}{2i+1} c_i, \tag{2.6}$$

where $c_i = EW^{2i}$.

Let $c_1^0, c_2^0, \ldots$ correspond to π_0 and $c_1, c_2, \ldots$ correspond to a general q.

Theorem 2.3.1. For $n = 2$ and the family Γ_3,

$$\delta_M(0) = 1 - \delta_M(2) = \frac{\frac{3-6\mu_2^0+\alpha}{12\mu_2^0+4\alpha} + \frac{2-4\mu_2^0+\alpha}{8\mu_2^0+2\alpha}}{2},$$

$$\text{and } \delta_M(1) = \frac{1}{2}.$$

Furthermore, δ_M is admissible and Bayes.

Proof: Let $G = (1-\varepsilon)\pi_0 + \varepsilon q$ denote a general prior in Γ_3 and δ_G denote the corresponding Bayes rule. Then, direct calculation gives

$$\delta_G(X) = \frac{\sum_{j=0}^{n-X} \binom{n-X}{j}(-1)^j \mu_{j+X+1}^0 + \alpha \sum_{j=0}^{n-X} \binom{n-X}{j}(-1)^j \mu_{j+X+1}}{\sum_{j=0}^{n-X} \binom{n-X}{j}(-1)^j \mu_{j+X}^0 + \alpha \sum_{j=0}^{n-X} \binom{n-X}{j}(-1)^j \mu_{j+X}}. \tag{2.7}$$

which simplifies to

$$\begin{aligned} \delta_G(0) &= \frac{\frac{1}{4} - \frac{\mu_2^0}{2} + \alpha(\frac{1}{4} - \frac{\mu_2}{2})}{\mu_0^2 + \alpha\mu_2}, \\ \delta_G(1) &= \frac{1}{2}, \\ \delta_G(2) &= 1 - \delta_G(0) \end{aligned} \tag{2.8}$$

Using (2.6), it follows that μ_2 varies in the interval $\frac{1}{4} \le \mu_2 \le \frac{1}{3}$. Maximizing and minimizing $\delta_G(0)$ over this range, the expression for $\delta_M(0)$ follows. The facts that $\delta_M(1) = \frac{1}{2}$ and $\delta_M(2) = 1 - \delta_M(0)$ are immediate.

To prove that δ_M is Bayes and admissible, it is enough to prove that $0 < \delta_M(0) \le \delta_M(1) = \frac{1}{2} \le \delta_M(2) < 1$ (see Skibinsky and Rukhin (1989)). That $\delta_M(0) > 0$ follows from the fact that $3 - 6\mu_2^0 + \alpha$ and $2 - 4\mu_2^0 + \alpha$ are both positive. That $\delta_M(0) \le \frac{1}{2}$ follows from the fact that $\delta_G(0) \le \delta_G(1) = \frac{1}{2}$ for each G by monotonicity of Bayes rules. The remaining two inequalities are immediate since $\delta_M(0) = 1 - \delta_M(2)$. This proves the Theorem.

Remark: Since the plausible priors were thought to belong to Γ_3, the more natural question is whether δ_M is Bayes with respect to some prior in Γ_3 itself. The following example illustrates this.

Example 1. Suppose π_0 is the uniform distribution on $[0,1]$. Then, using (2.7), δ_M is Bayes with respect to $G = (1-\varepsilon)\pi_0 + \varepsilon q$ if and only if there exists a symmetric unimodal q with second moment μ_2 satisfying

$$\frac{\frac{1}{4} + \frac{\alpha+\frac{2}{3}}{2(\alpha+\frac{4}{3})}}{2} = \frac{\frac{1}{12} + \alpha(\frac{1}{4} - \frac{\mu_2}{2})}{\frac{1}{3} + \alpha\mu_2}$$

$$\Leftrightarrow \mu_2 = \frac{\frac{1}{12} + \frac{\alpha}{2} - \frac{\alpha+\frac{2}{3}}{6(\alpha+\frac{4}{3})}}{\alpha(\frac{5}{4} + \frac{\alpha+\frac{2}{3}}{2(\alpha+\frac{4}{3})})} \tag{2.9}$$

Recalling that for symmetric unimodal distributions $\frac{1}{4} \leq \mu_2 \leq \frac{1}{3}$, it follows that δ_M is Bayes with respect to a prior in Γ_3 if and only if the right hand side of (2.9) is between $\frac{1}{4}$ and $\frac{1}{3}$. This is indeed the case for each ε between 0 and 1. Also, since the value of μ_2 in (2.9) does not correspond to an upper or a lower principal representation, for any ε, δ_M is Bayes with respect to infinitely many priors in Γ_3.

Theorem 2.3.2. For $n = 3$ and the family Γ_3,

$$\begin{aligned} \delta_M(0) &= \frac{1}{20}\left(\frac{5-y+5\alpha}{1+\alpha+x} + \frac{5-x^2+4\alpha}{1+2\alpha+x}\right), \\ \delta_M(2) &= \frac{3}{20}\left(\frac{5-y+5\alpha}{3-x+3\alpha} + 2c^*\right), \\ \delta_M(1) &= 1 - \delta_M(2) \\ \text{and } \delta_M(3) &= 1 - \delta_M(0), \end{aligned}$$

where $x = c_1^0$, $y = c_2^0$ and c^* is as in (2.15). Furthermore, δ_M is admissible and Bayes.

Proof: Using the notation of Theorem 2.3.1,

$$\delta_G(0) = \frac{\frac{1}{16} - \frac{y}{80} + \alpha(\frac{1}{16} - \frac{c_2}{80})}{\frac{1}{8}(1+x) + \frac{\alpha}{8}(1+c_1)},$$

$$\delta_G(2) = \frac{\frac{1}{16} - \frac{y}{80} + \alpha(\frac{1}{16} - \frac{c_2}{80})}{\frac{1}{8} - \frac{x}{24} + \alpha(\frac{1}{8} - \frac{c_1}{24})}. \tag{2.10}$$

It is clear that $\delta_G(0)$ is decreasing in both c_1 and c_2. Since c_1, c_2 belong to the set

$$\{(c_1, c_2)\colon\ 0 \leq c_1^2 \leq c_2 \leq c_1 \leq 1\},$$

elementary calculations give

$$\sup_G \delta_G(0) = \frac{\frac{1}{16} - \frac{y}{80} + \frac{\alpha}{16}}{\frac{1}{8}(1+x) + \frac{\alpha}{8}}$$

$$\text{and } \inf_G \delta_G(0) = \frac{\frac{1}{16} - \frac{y}{80} + \frac{\alpha}{20}}{\frac{1}{8}(1+x) + \frac{\alpha}{4}}. \tag{2.11}$$

The expression for $\delta_M(0)$ follows from (2.11). Also, $\delta_G(2)$ is clearly decreasing in c_2 for fixed c_1, implying

$$\inf_G \delta_G(2) = \inf_{0 \le c \le 1} \frac{\frac{1}{16} - \frac{y}{80} + \alpha(\frac{1}{16} - \frac{c}{80})}{\frac{1}{8} - \frac{x}{24} + \alpha(\frac{1}{8} - \frac{c}{24})}. \tag{2.12}$$

The expression in c on the right side of (2.12) is increasing in c, from which it follows that

$$\inf_G \delta_G(2) = \frac{\frac{1}{16} - \frac{y}{80} + \frac{\alpha}{16}}{\frac{1}{8} - \frac{x}{24} + \frac{\alpha}{8}}. \tag{2.13}$$

On the other hand,

$$\sup_G \delta_G(2) = \sup_{0 \le c \le 1} \frac{\frac{1}{16} - \frac{y}{80} + \alpha(\frac{1}{16} - \frac{c^2}{80})}{\frac{1}{8} - \frac{x}{24} + \alpha(\frac{1}{8} - \frac{c}{24})}. \tag{2.14}$$

This supremum is attained at

$$c^* = \frac{3 - x + 3\alpha - \sqrt{(3 - x + 3\alpha)^2 - \alpha(5 - y + \alpha)}}{\alpha}, \tag{2.15}$$

which, on algebra, gives

$$\sup_G \delta_G(2) = \frac{3c^*}{5}. \tag{2.16}$$

The expression for $\delta_M(2)$ now follows from (2.13), (2.16).

To prove that δ_M is admissible, we verify that

$$(\delta_1 - \delta_0)(1 - \delta_1)(\delta_3 - \delta_2)\delta_2 \ge \delta_0(1 - \delta_3)(\delta_2 - \delta_1)^2, \tag{2.17}$$

(see Kozek (1982) or Skibinsky and Rukhin (1989)), where $\delta_i = \delta_M(i)$. On simple algebra, this reduces to

$$\frac{\delta_2(1 - \delta_2)}{3\delta_2 - 1} \ge \delta_2 \tag{2.18}$$

(2.18) can be verified by a lengthy argument which essentially shows that if (2.18) holds for $\alpha = 0$ and ∞, then it holds for all $0 < \alpha < \infty$ too. We omit these details.

Finally, since δ_M is admissible and $0 < \delta_M(i) < 1$ for all i, it is also Bayes (see Johnson (1971)). This proves the result.

Theorem 2.3.3. For $n = 4$ and the family Γ_3,

$$\delta_M(0) = \frac{1}{4}\left(\frac{15 + 10x - 9y + 16\alpha}{15 + 30x + 3y + 48\alpha} + \frac{15 + 10x - 9y + 15\alpha}{15 + 30x + 3y + 15\alpha}\right),$$

$$\delta_M(1) = \frac{1}{4}\left(\frac{15 - 10x + 3y + 15\alpha}{15 - 3y + 15\alpha} + \frac{3y - 5x + 5\sqrt{\Delta}}{15 - 3y + 15\alpha}\right),$$

$$\delta_M(2) = \frac{1}{2},\ \delta_M(3) = 1 - \delta_M(1) \ \text{ and } \ \delta_M(4) = 1 - \delta_M(0),$$

where

$$\Delta = (3 - x + 3\alpha)^2 - \alpha(5 - y + 5\alpha).$$

Furthermore, δ_M is admissible and Bayes.

Proof: The proof uses arguments similar to that of Theorem 3 and will be omitted.

Example 2. The expressions for $\delta_M(x)$ get much simplified for specific priors π_0. For example, if π_0 is $u[0,1]$, then

$$\delta_M(0) = \frac{32 + 20\alpha}{6(32 + 10\alpha)}$$

$$\delta_M(1) = \frac{32 + 50\alpha}{3(32 + 40\alpha)}$$

$$\delta_M(2) = \frac{1}{2}$$

Notice that for small α, this estimate is nearly equal to δ_{π_0}, an expected result. Also notice that $|\delta_M(X) - \delta_{\pi_o}(X)| \leq \frac{1}{6}$ uniformly in α and X.

For any given π_0, $0 < \varepsilon < 1$, and n, the admissibility status of δ_M can be verified by using the characterization of admissibility in Kozek (1982). Numerical evidence suggests that δ_M is admissible for all π_0, ε and n.

From Kozek (1982), δ_M is admissible if and only if the matrices M_1, M_2 are nonnegative definite where $M_1 = ((b_{i+j}))$, $0 \leq i,\ j \leq k$ and $M_2 = ((b_{i+j+1}))$, $0 \leq i,\ j \leq k-1$ for

$n = 2k-1$ and $M_1 = ((b_{i+j+1}),\ 0 \le i,j \le k$ and $M_2 = ((b_{i+j})),\ 0 \le i,j \le k$, for $n = 2k$; see (2.21) for the definition of b_i. The following table gives the value λ of the minimum of the minimum eigenvalues of M_1 and M_2 for various π_0, ε and n. In each case, π_0 is taken as a Beta (α, α) distribution so that π_0 is indexed by α.

Table: Values for λ

		n			
ε	α	5	10	15	20
	1	3.73×10^{-5}	7.9×10^{-10}	0	0
.01	4	6.84×10^{-5}	3.26×10^{-9}	0	0
	8	3.46×10^{-4}	2.87×10^{-9}	0	0
	1	3.8×10^{-5}	8.43×10^{-10}	0	0
.05	4	5.51×10^{-5}	4.81×10^{-9}	0	0
	8	3.45×10^{-5}	2.67×10^{-9}	0	0
	1	3.9×10^{-5}	9.09×10^{-10}	0	0
.1	4	4.72×10^{-5}	6.68×10^{-9}	0	0
	8	3.45×10^{-5}	2.64×10^{-9}	0	0

2.4. The Family Γ_4. This is the broadest family of symmetric priors. As will be seen, however,use of such a large prior family leads to bad consequences. The midpoint estimate δ_M here is a rather unattractive estimate and becomes inadmissible for every $n \ge 3$.

Theorem 2.4.1. For the family Γ_4,

$$\begin{aligned} \delta_M(X) &= \frac{1}{4} && \text{if } 0 \le X < [\frac{n}{2}] \\ &= \frac{1}{2} && \text{if } X = [\frac{n}{2}] \\ &= \frac{3}{4} && \text{if } [\frac{n}{2}] < X \le n. \end{aligned}$$

Furthermore δ_M is admissible if and only if $n \le 2$.

Proof: Since symmetric distributions are mixtures of two point symmetric distributions, by a familiar argument (see, for example, Casella and Berger (1987)), $\sup_{G \in \Gamma_4} \delta_G(X)$ and $\inf_{G \in \Gamma_4} \delta_G(X)$ can be calculated by evaluating $\sup_{G \in \Gamma_4'} \delta_G(X)$ and $\inf_{G \in \Gamma_4'} \delta_G(X)$, where

$$\Gamma_4' = \{G \in \Gamma_4\colon\ G\{\frac{1}{2} + z\} = G\{\frac{1}{2} - z\} = \frac{1}{2}, 0 \le z \le \frac{1}{2}\}.$$

Now, clearly,

$$\sup_{G\in\Gamma_4'} \delta_G(X) = \sup_{0\le z\le \frac{1}{2}} \frac{(\frac{1}{2}+z)^{X+1}(\frac{1}{2}-z)^{n-X} + (\frac{1}{2}-z)^{X+1}(\frac{1}{2}+z)^{n-X}}{(\frac{1}{2}+z)^{X}(\frac{1}{2}-z)^{n-X} + (\frac{1}{2}-z)^{X}(\frac{1}{2}+z)^{n-X}} \tag{2.19}$$

On making the transformation $y = \frac{\frac{1}{2}+z}{\frac{1}{2}-z}$, the supremum on the right in (2.19) equals

$$\sup_{1\le y<\infty} \frac{1}{y+1}\cdot\frac{y+y^{n-2X}}{1+y^{n-2X}};$$

calculus shows that for $X < [\frac{n}{2}]$, the supremum is attained at $y = 1$, giving $\sup_{G\in\Gamma_4} \delta_G(X) = \frac{1}{2}$. A similar argument shows that for $X < [\frac{n}{2}]$, $\inf_{G\in\Gamma_4} \delta_G(X) = 0$, giving $\delta_M(X) = \frac{1}{4}$ for $X < [\frac{n}{2}]$. The case $X = [\frac{n}{2}]$ is immediate and the case $X > [\frac{n}{2}]$ follows from symmetry.

That δ_M is admissible for $n = 1, 2$ is trivial. We will now prove it is inadmissible for $n \ge 3$. We will prove this for the case n is odd; the case when n is even is similar. Let then $n = 2k-1$, $k \ge 2$. It follows from Kozek (1982) that it suffices to prove that the matrix

$$M_1 = \begin{bmatrix} b_0 & b_1 & \cdots & b_k \\ b_1 & b_2 & \cdots & b_{k+1} \\ \vdots & & & \\ b_k & b_{k+1} & \cdots & b_{2k} \end{bmatrix} \tag{2.20}$$

is not nonnegative definite, where,

$$\begin{aligned} b_{n+1} &= \delta_0\delta_1\ldots\delta_n &&= (\frac{1}{4})^k(\frac{3}{4})^k \\ b_n &= (1-\delta_n)\delta_0\delta_1\ldots\delta_{n-1} &&= (\frac{1}{4})^{k+1}(\frac{3}{4})^{k+1} \\ &\vdots \\ b_1 &= (1-\delta_n)(1-\delta_{n-1})\ldots\delta_0 &&= (\frac{1}{4})^{k+1}(\frac{3}{4})^{k+1} \\ b_0 &= (1-\delta_n)(1-\delta_{n-1})\ldots(1-\delta_0) &&= (\frac{1}{4})^k(\frac{3}{4})^k. \end{aligned} \tag{2.21}$$

$$\begin{aligned} \therefore M_1 &= \frac{1}{4^{n+1}} \begin{bmatrix} 3^k & 3^{k-1} & \cdots & 3^1 & 3^0 \\ 3^{k-1} & 3^{k-2} & \cdots & 3^0 & 3^1 \\ \vdots & & & & \\ 3^0 & 3^1 & \cdots & 3^{k-1} & 3^k \end{bmatrix} \\ &= \frac{1}{4^{n+1}} G \text{ (say)}. \end{aligned}$$

We claim that the matrix G is not nonnegative definite. To prove this, one can check directly that $|G| = (-1)^{\frac{k(k+3)}{2}} 8^k \neq 0$ and every other leading principal minor of G is zero. If $|G|$ is negative, clearly M_1 is not nonnegative definite. If $|G|$ is positive and G is nonnegative definite, then each eigenvalue of G must be strictly positive which contradicts the fact that G has leading principal minors equal to zero. Hence G cannot be nonnegative definite, proving the theorem.

2.5. The Family Γ_5. Although symmetry may be natural in some problems, it may not be so in others. Asymmetry of the prior can be modeled in various ways. Since conjugate priors form a rich family of priors in the Binomial problem, it is natural to study the behavior of the midpoint δ_M by taking general Beta priors. The family of all beta priors is unreasonably large. We will therefore take Beta priors with some additional constraints which seem natural.

Basically, we will have a location constraint and a variability constraint. A location constraint can be given in terms of any of the mean, median or the mode (if one exists) or a combination of these. We will use a restriction of the form $\frac{1}{2} - b \leq \mu \leq \frac{1}{2} + b$ where μ is a measure of location of the prior. A restriction on only the mean or the median does not seem tenable as it cannot rule out priors of the form Beta (α, α) where α is arbitrarily close to zero. Since the priors Beta (α, α) weakly converge to the two point distribution with mass $\frac{1}{2}$ at 0 and 1 each as $\alpha \to 0$ and this latter prior is hardly ever a realistic prior, a bound on the mean alone cannot give us a satisfactory class of priors. The same objection applies to the median. The problem can be overcome by using only such Beta priors which have a unique mode in the interior of (0,1) and by assuming the mode is between $\frac{1}{2} \pm b$ for some fixed $0 < b < \frac{1}{2}$. As we will momentarily see, this constraint automatically implies that the mean and the median of the prior are also each between $\frac{1}{2} \pm b$. This really provides a very satisfactory resolution of the problem.

A constraint on the mode alone, on the other hand, leaves open the possibility that the prior can be Beta (α, α) for arbitrarily large α; since these converge weakly to a point mass at $\frac{1}{2}$ which again is rarely realistic, it seems necessary that we also require the prior to satisfy a variability constraint.

If $\theta \sim$ Beta (α, β), the contour of equal variance (i.e., the set of all (α, β) such that

$\frac{\alpha\beta}{(\alpha+\beta)^2(\alpha+\beta+1)}$ = constant) is symmetric about the line $\alpha = \beta$. The tangent to this contour at its unique point of intersection with the 45° line is a line of the form $\alpha + \beta = \gamma$. We take all Beta (α, β) priors of the form

$$\Gamma_5 = \{\pi\colon\ \pi(\theta) = \frac{\theta^{\alpha-1}(1-\theta)^{\beta-1}}{B(\alpha,\beta)},\ \frac{1}{2} - b \leq \text{ mode of } \theta \leq \frac{1}{2} + b,\ \alpha + \beta \leq \gamma\} \tag{2.22}$$

Note that bounding $\alpha + \beta$ automatically places a lower bound on the variance of θ; also recall that the mode of θ is $\frac{\alpha-1}{\alpha+\beta-2}$ if $\alpha > 1,\ \beta > 1$. Thus the set of (α, β) satisfying the constraints in (2.22) is a triangular region on the plane; Figures 1 and 2 describe the contours of equal variance and also the triangular regions in two special cases. Notice that using the triangular region instead of the exact set satisfying the modal and the variance bound enlarges the class of priors slightly while making the problem technically much simpler.

We first describe two propositions which relate to the analysis here.

Proposition 2.5.1. Let $\theta \sim$ Beta (α, β), $\alpha > 1,\ \beta > 1$. If $a \leq$ mode of $\theta \leq b$, where $0 < a < \frac{1}{2} < b$, then $a \leq E(\theta) \leq b$.

Proof: Trivial.

Proposition 2.5.2. Let $\theta \sim$ Beta (α, β), $\alpha > 1,\ \beta > 1$. If mode of $\theta = b \geq \frac{1}{2}$, then median of $\theta \leq b$; if mode of $\theta = a \leq \frac{1}{2}$, then median of $\theta \geq a$.

Cor. If $\theta \sim$ Beta (α, β), $\alpha > 1, \beta > 1$, and if for some $0 < b < \frac{1}{2}$, $\frac{1}{2} - b \leq$ mode of $\theta \leq \frac{1}{2} + b$, then $\frac{1}{2} - b \leq$ median of $\theta \leq \frac{1}{2} + b$.

Proof of Proposition 2.5.2: We will only prove the first part. The second part follows on transforming to $1 - \theta$.

Clearly, it is enough to show that $\int\limits_0^b \frac{\theta^{\alpha-1}(1-\theta)^{\beta-1}}{B(\alpha,\beta)} d\theta \geq \frac{1}{2}$. Towards this end, observe that

$$1 = \int\limits_0^{2b-1} \frac{\theta^{\alpha-1}(1-\theta)^{\beta-1}}{B(\alpha,\beta)} d\theta + \int\limits_{2b-1}^{b} \frac{\theta^{\alpha-1}(1-\theta)^{\beta-1}}{B(\alpha,\beta)} d\theta + \int\limits_b^1 \frac{\theta^{\alpha-1}(1-\theta)^{\beta-1}}{B(\alpha,\beta)} d\theta \tag{2.23}$$

Consequently, it will suffice to show that

$$\int_{2b-1}^{b} \frac{\theta^{\alpha-1}(1-\theta)^{\beta-1}}{B(\alpha,\beta)}d\theta \geq \int_{b}^{1} \frac{\theta^{\alpha-1}(1-\theta)^{\beta-1}}{B(\alpha,\beta)}d\theta$$

$$\Leftrightarrow \int_{0}^{1-b} [(b-u)^b(1-b+u)^{1-b}]^{\frac{\alpha-1}{b}} d\theta \geq \int_{0}^{1-b} [(b+u)^b(1-b-u)^{1-b}]^{\frac{\alpha-1}{b}} d\theta, \quad (2.24)$$

on substituting $\theta = b - u$ in the first integral and $\theta = b + u$ on the second integral and using the fact that the mode of θ equals b. However, (2.24) follows since calculus gives that $\frac{(b+u)^b(1-b-u)^{1-b}}{(b-u)^b(1-b+u)^{1-b}} \leq 1$ for $0 < u < 1 - b$.

As in the case of Γ_3, we can analytically prove that δ_M is admissible and Bayes for all $0 < b < \frac{1}{2}$ and all $\gamma \geq 2$ if $n \leq 4$. δ_M is easy enough to describe for this family. Let the three vertices of the triangular region in which (α, β) are assumed to lie be (1,1), (α_1^*, β_1^*) and (α_2^*, β_2^*) (see Figures 1 and 2). Then

$$\begin{aligned} \delta_M(X) &= \frac{\frac{X+1}{n+2} + \frac{x+\alpha_2^*}{n+\gamma}}{2} && \text{if } 0 \leq X \leq n_1 \\ &= \frac{X + \frac{\gamma}{2}}{n+\gamma} && \text{if } n_1 < X \leq n_2 \\ &= \frac{\frac{X+1}{n+2} + \frac{x+\alpha_2^*}{n+\gamma}}{2} && \text{if } n_2 < X \leq n, \end{aligned}$$

where $n_1 = [n(\frac{1}{2} - b) - 2b]$ and $n_2 = [n(\frac{1}{2} + b) + 2b]$. Note $\alpha_1^* = \gamma(\frac{1}{2} - b) + 2b$ and $\alpha_2^* = \gamma(\frac{1}{2} + b) - 2b$ and $\beta_i^* = \gamma - \alpha_i^*$. Interestingly, δ_M always equals the average of Bayes estimates with respect to two Beta priors whose parameters coincide with two vertices of the triangle and also δ_M is symmetric even though Γ_5 contains asymmetric priors. This is because for any X, the supremum as well as the infimum of the Bayes estimates with respect to the family Γ_5 are attained at one of the three vertices of the triangle. The following theorem will be stated without proof since the techniques are the same as before.

Theorem 2.5.3. For each $0 < b < \frac{1}{2}$, $\gamma \geq 2$ and $n \leq 4$, δ_M is admissible and Bayes. For each $n \leq 4$, there exists b_n such that $b_n \leq b < \frac{1}{2}$ implies that δ_M is Bayes with respect to the Beta $(\frac{\gamma}{2}, \frac{\gamma}{2})$ prior. Furthermore $b_1 = \frac{1}{6}$, $b_2 = \frac{1}{4}$, $b_3 = \frac{3}{10}$ and $b_4 = \frac{1}{3}$.

Discussion: For $n > 4$, an analytical proof of the admissibility of δ_M was not possible. However, a numerical implementation of Kozek's (1982) characterization again indicates that δ_M is admissible for any b, γ and n, although we are less sure that it is Bayes.

3. Risk Behavior

Since a number of prior classes was considered, it seems necessary that a comparison of their risk functions be made. This is to help understand in which cases the use of δ_M is reasonable and in which cases it is not. Figures 6, 7 and 8 give the estimates δ_M themselves, while figures 3, 4 and 5 give the plots of the risk functions of δ_M for various prior classes; the unbiased estimate is used as a standard. The plots are obtained by calculating each estimate and then exactly calculating its risk function as a polynomial in θ. Since the plots are self explanatory, no effort will be made to illustrate them verbally. The very clear moral is that use of δ_M is very reasonable and safe when either Γ_3 or Γ_5 is used (with respect to Γ_3, π_0 was taken as $u[0,1]$ in the plots), or even Γ_1 unless n is very small. This is saying that as long as Γ does not contain unreasonable priors, the use of δ_M should be sensible in this problem. Note how Γ_2 and Γ_4 are both rather conservative classes of priors. Γ_2 includes a prior degenerate at .5 and Γ_4 contains in addition the prior which assigns mass only at 0 and 1. By the same token, even Γ_5 is conservative without a lower bound on the variance since degenerate priors will be included. Not surprisingly, for each of $n = 4$, 10 and 20, the risk function of δ_6 is unattractive.

The results indicate that in general, the minimax method for choosing among Bayes estimates leads to good risk properties. This was also evidenced in the normal problem, as in DasGupta and Studden (1989). However, we do not expect that in other problems, such estimates will be actually Bayes corresponding to a fixed prior. We believe the Binomial case may be an exception in this sense. The lack of a verifiable characterization is also going to make resolution to this question difficult in most other problems.

4. Conclusion

Analytical derivation of systematically chosen Bayes actions from a set of actions is a hard mathematical problem. Geometrically intuitive estimates such as the midpoint are much easier to evaluate. Of course, whether such apparently naive estimates have any optimality property is the important issue.

Our results here indicate that they in fact do, although excessive conservatism by using very large classes of priors will not help here. The problems we address here are equally and perhaps more important in higher dimensions. It is usually the case that the collection of Bayes actions for many natural families of priors form a convex set in the appropriate dimension (see DasGupta and Studden (1988)). Epicenters of these convex sets are easy enough to identify and their use can be more justified if they turn out to be genuine Bayes or admissible.

References

1. Berger, J. (1990). Robust Bayesian analysis: sensitivity to the prior, *Jour. Stat. Planning and Inf.*, **25**, 303–328.

2. Brown, L. D. (1981). A complete class theorem for statistical problems with finite sample spaces, *Ann. Stat.*, **9**, 960–994.

3. Casella, G. and Berger, R. (1987). Reconciling Bayesian and frequentist evidence in the one-sided testing problem (with discussion). *Jour. Stat. Assoc.*, **82**, 106–111.

4. DasGupta, A. and Bose, A. (1988). Γ-minimax and restricted risk Bayes estimation of multiple Poisson means, *Stat. and Decisions*, **6**, 311–341.

5. DasGupta, A. and Studden, W. J. (1988). Robust Bayesian analysis in normal linear models with many parameters, Technical Report, Department of Statistics, Purdue University.

6. DasGupta, A. and Studden, W. J. (1989). Frequentist behavior of robust Bayes estimates of normal means, *Stat. and Decisions*, **7**, 333–361.

7. DasGupta, A. and Delampady, M. (1990) Γ-minimax estimation of vector parameters when the parameter space is a bounded convex set, Technical Report, Department of Statistics, Purdue University.

8. DasGupta, A. (1990). Pitman efficiencies of nonparametric tests in contamination families, Technical Report, Department of Statistics, Purdue University.

9. Huber, P. J. (1973). Robust Statistics, John Wiley, New York.

10. Johnson, B. M. (1971). On admissible estimators for certain fixed sample Binomial problems, *Ann. Math. Statist.* **42**, 1529–1587.

11. Karlin, S. (1968). Total Positivity, Vol. 1, Stanford University Press, Stanford, California.

12. Khintchine, A. K. (1938). On unimodal distributions, Izv. Nauchino. Issled. Inst. Mat. Mech. Tomsk. Gos. Univ. **2**, 1–7.

13. Sivaganesan, S. and Berger, J. (1989). Ranges of posterior measures for priors with unimodal contaminations, *Ann. Statist.*, **10**, 868–889.

14. Skibinsky, M. and Rukhin, A. L. (1989). Admissible estimators of Binomial probability and the inverse Bayes rule map, *Ann. Inst. Stat. Math.*, **41**, 699–716.

FIGURE 1

Region of prior parameters in family Γ_5

$\gamma=24 \quad b=0.1$

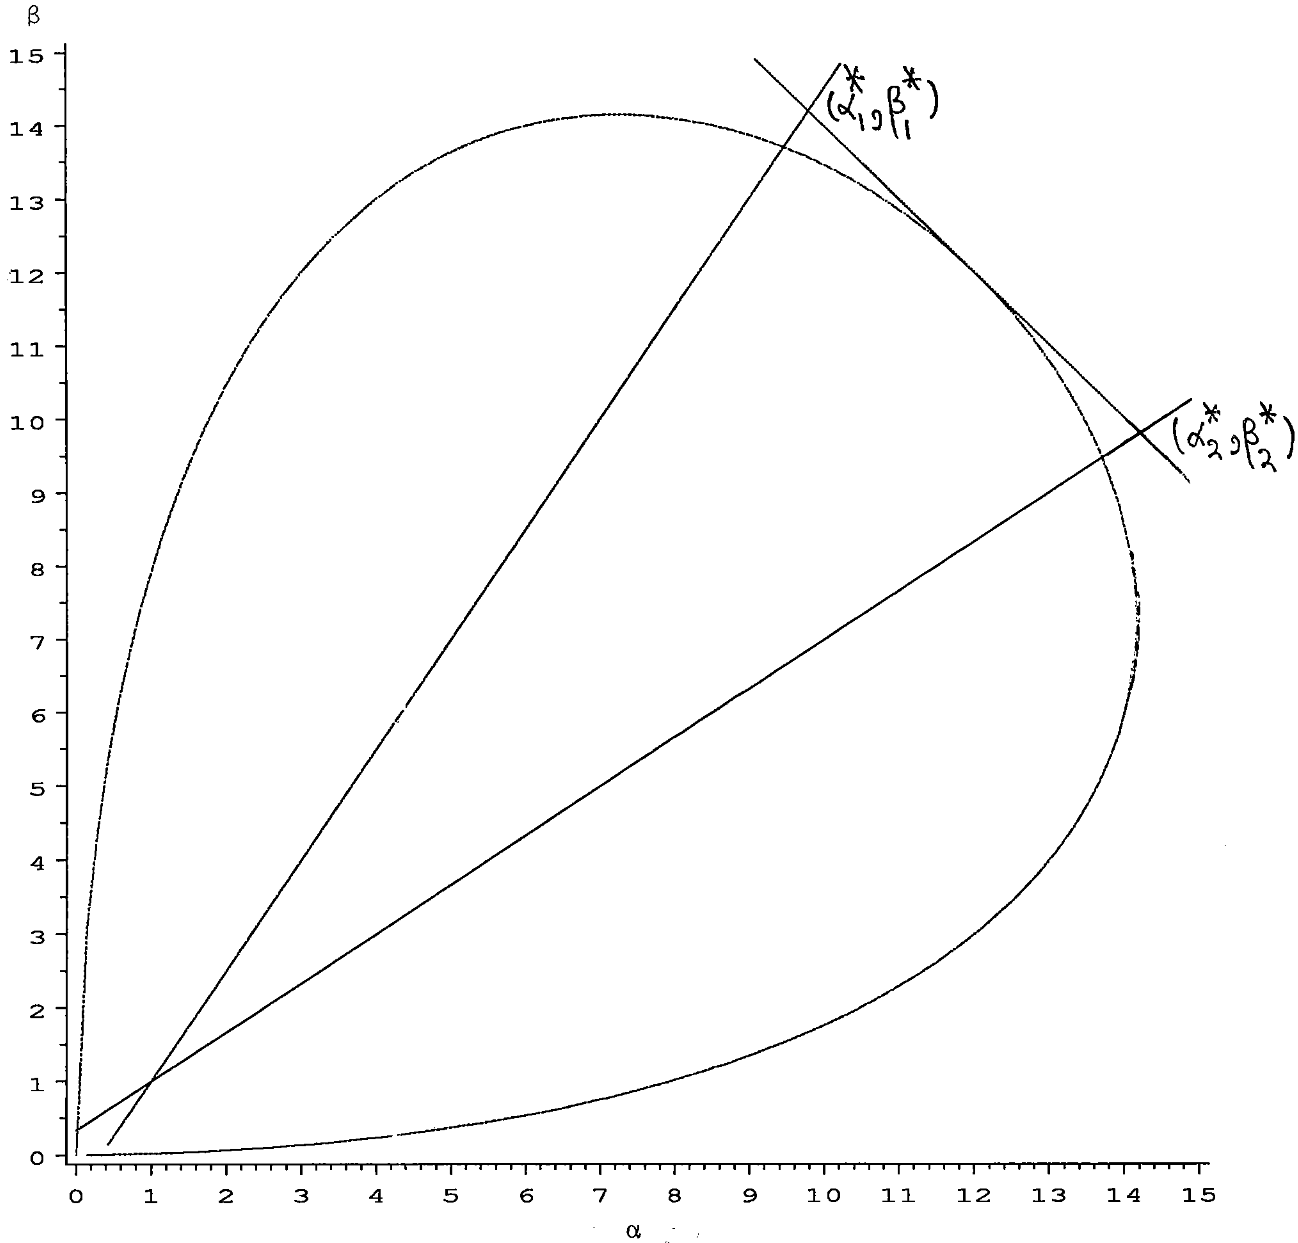

FIGURE 2

Region of prior parameters in family Γ_5

γ=11.5 b=0.1

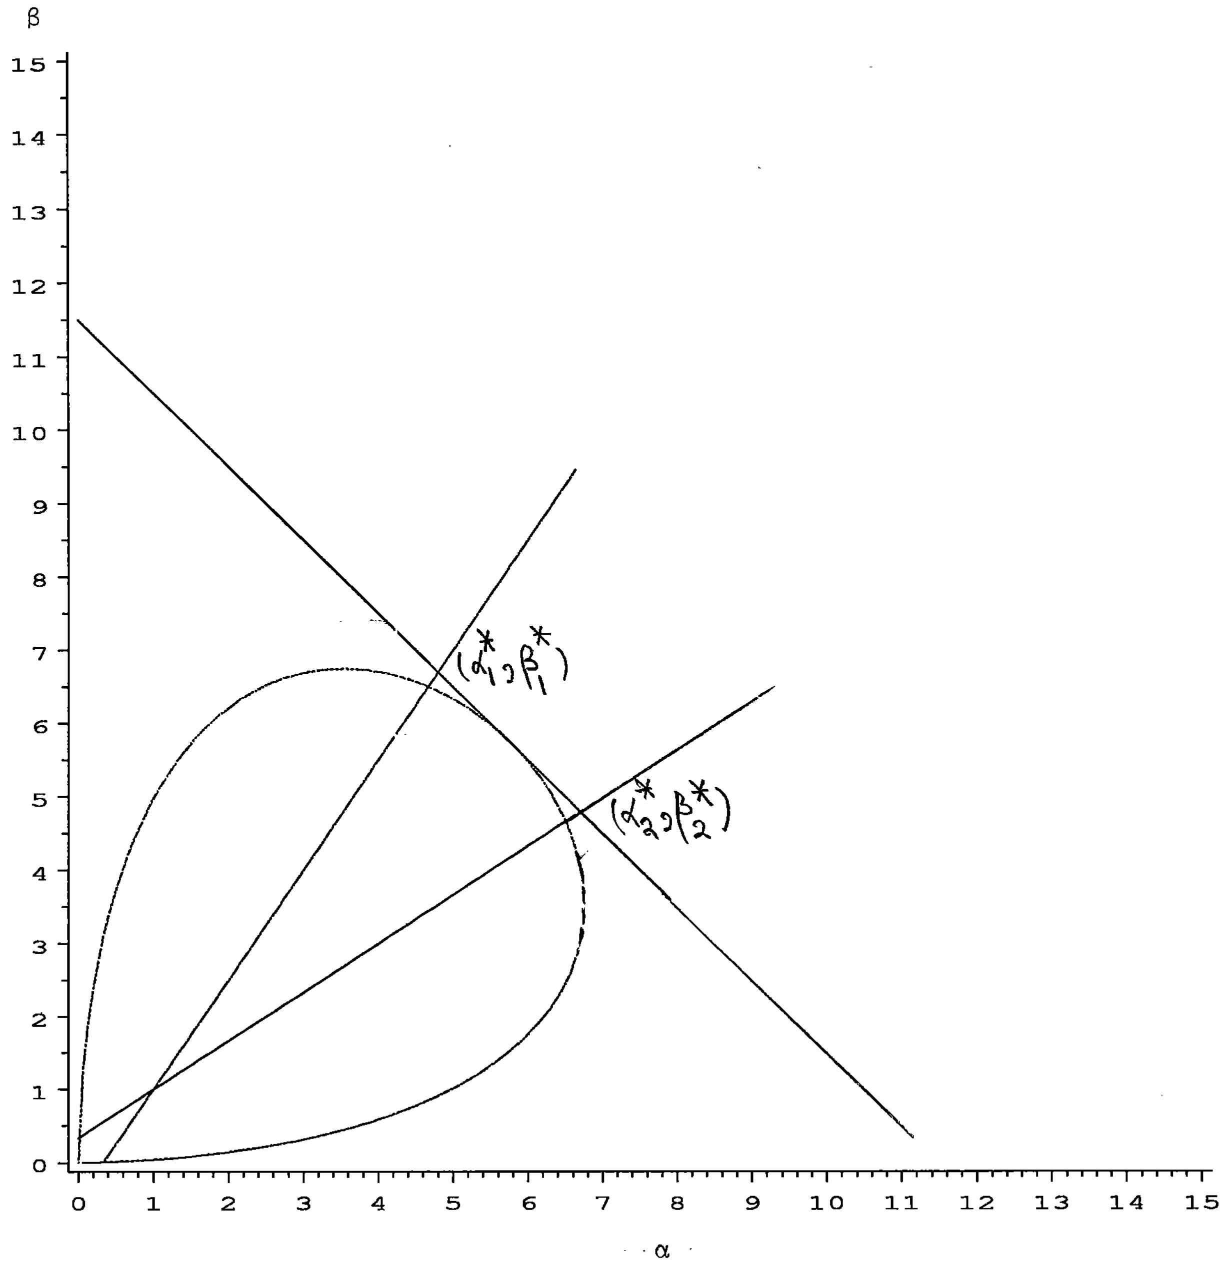

FIGURE 3

The Risk Functions of Estimates for 7 Classes of Priors
#1: Be(a,a), 1<a<6 #2: Be(a,b), γ=4 b=0.05
#3: symmetric #4: symmetric unimodal
#5: e-contaminated e=0.05 #6: Be(a,b), γ=infinite b = 0.05
#7: unbiased

n=4

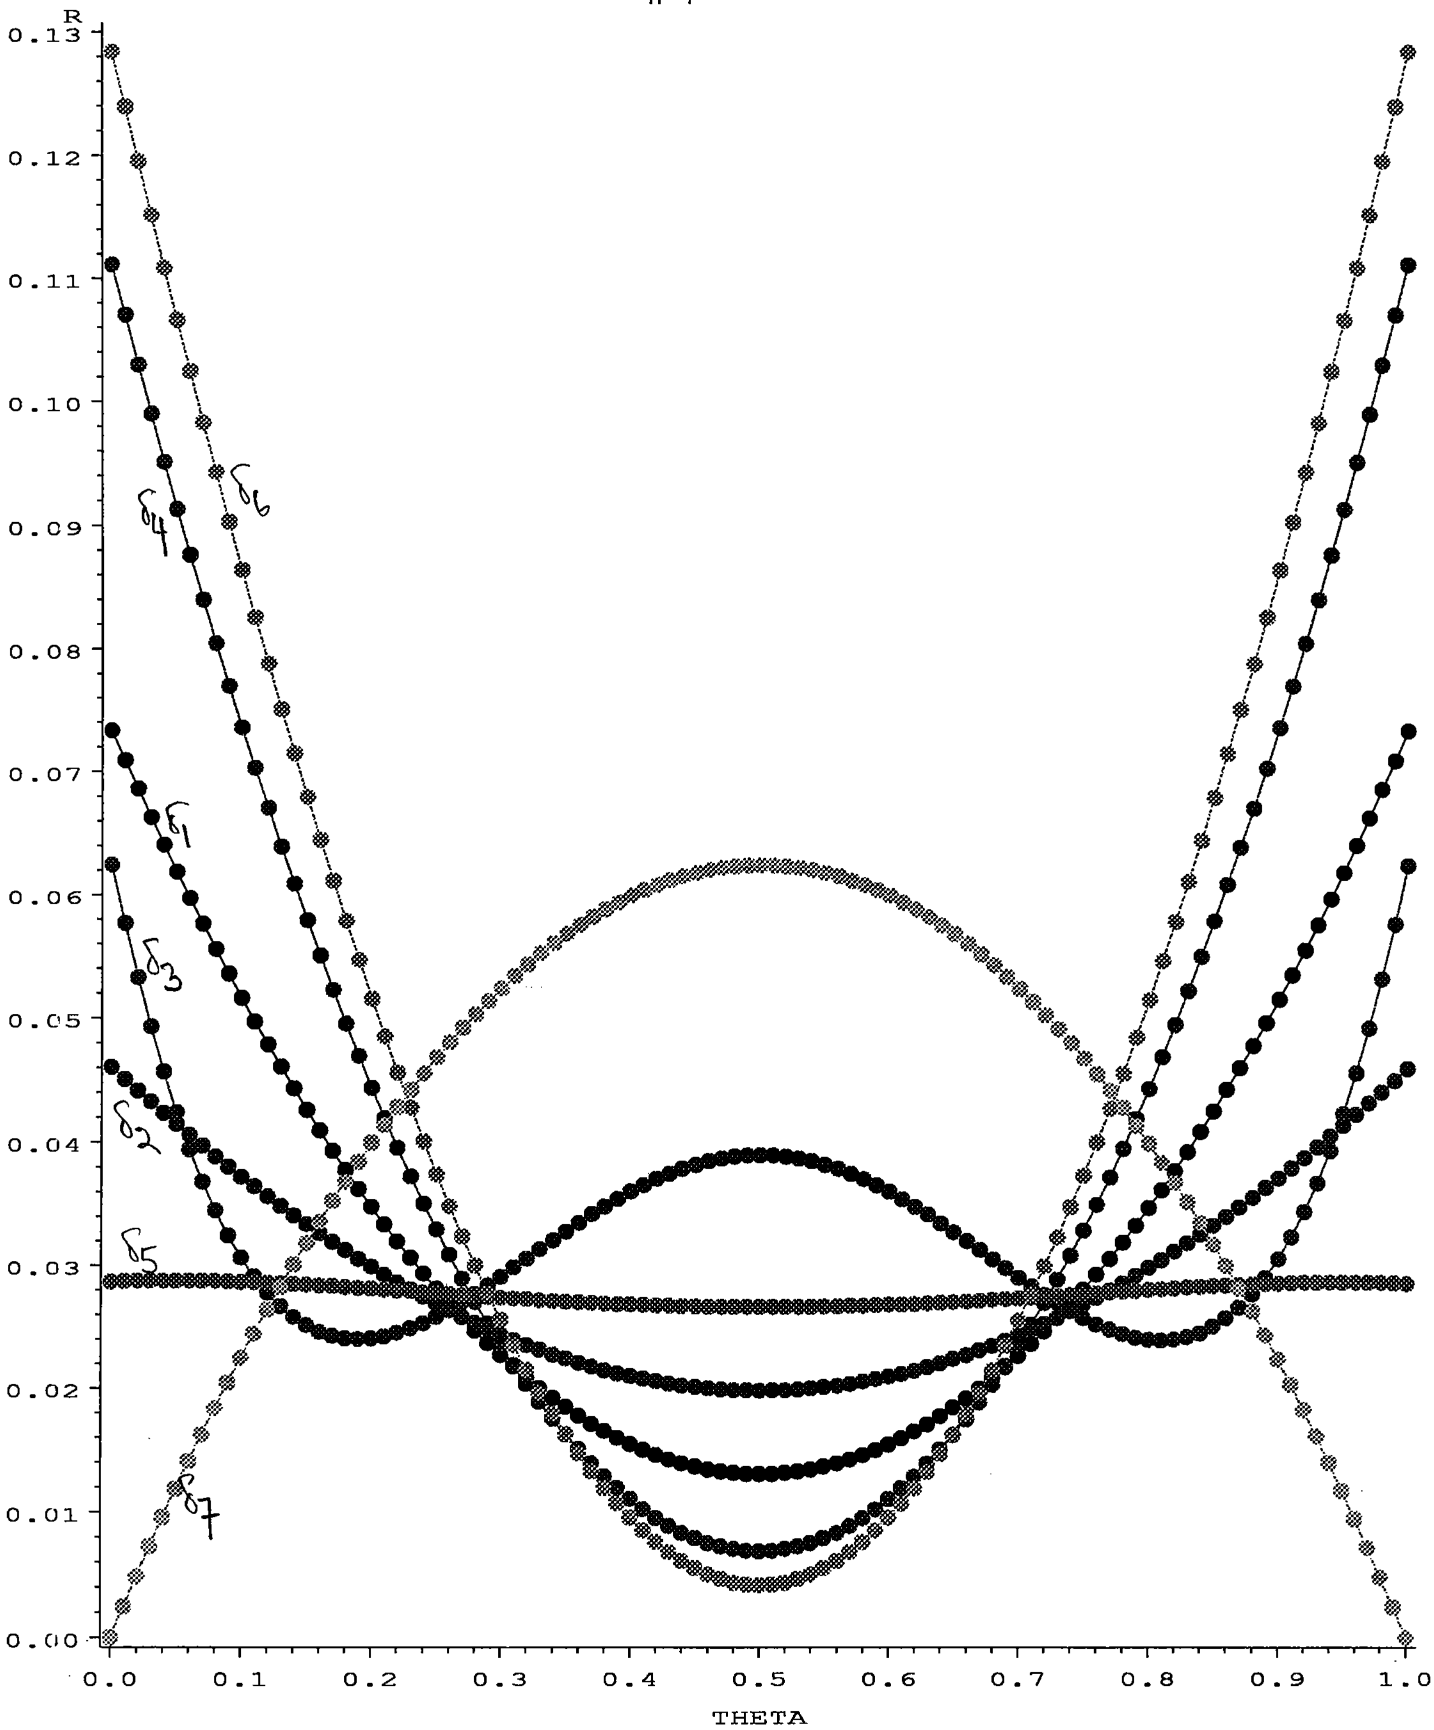

FIGURE 4

The Risk Functions of Estimates for 7 Classes of Priors

#1: Be(a,a), 1<a<6 #2: Be(a,b), γ=4 b=0.05

#3: symmetric #4: symmetric unimodal

#5: e-contaminated e=0.05 #6: Be(a,b), γ=infinite b = 0.05

#7: unbiased

n=10

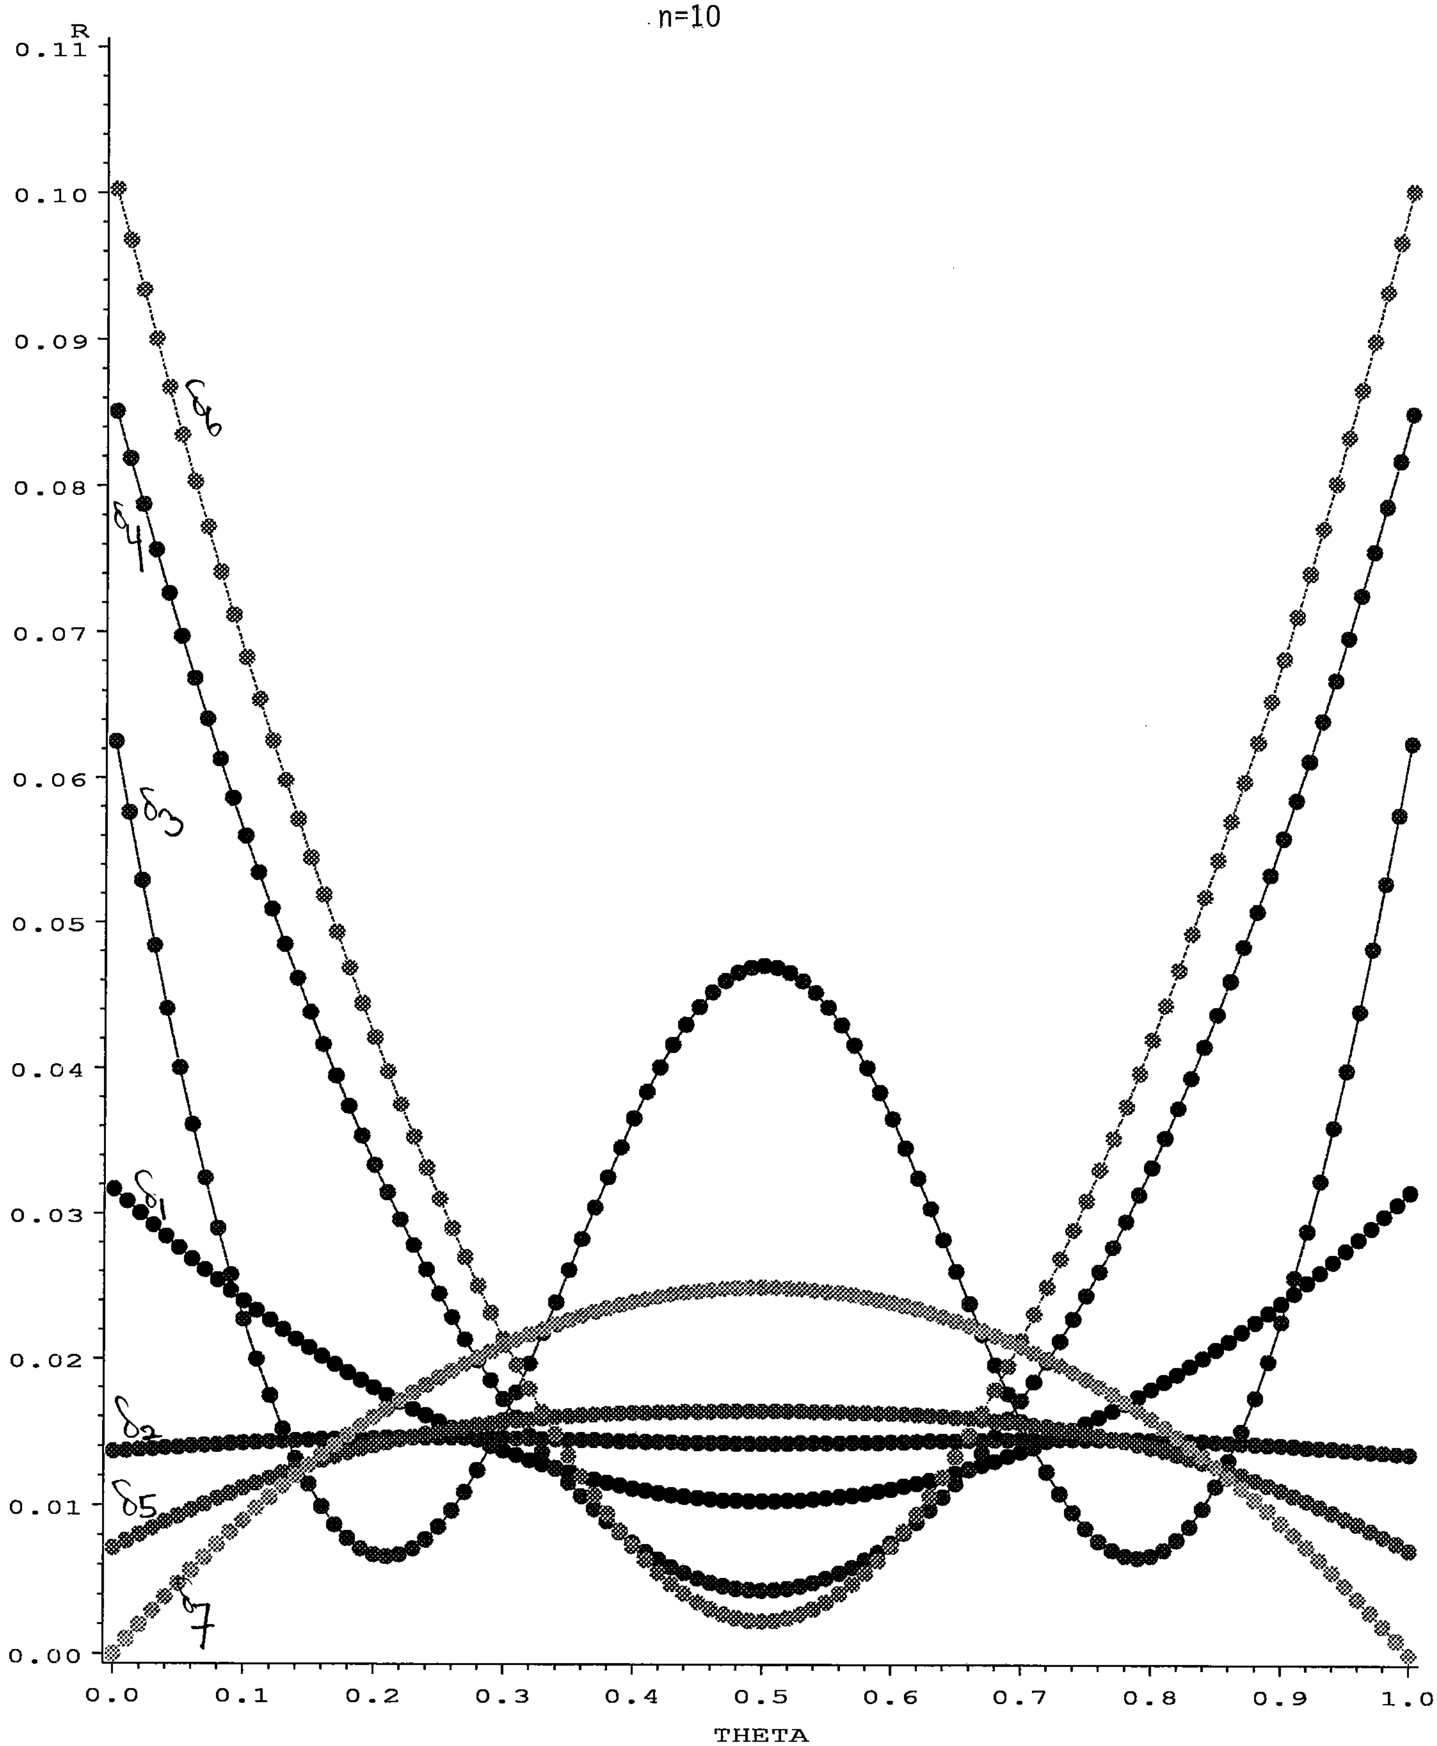

FIGURE 5

The Risk Functions of Estimates for 7 Classes of Priors
#1: Be(a,a), 1<a<6 #2: Be(a,b),γ=4 b=0.05
#3: symmetric #4: symmetric unimodal
#5: e-contaminated e=0.05 #6: Be(a,b), γ=infinite b = 0.05
#7: unbiased

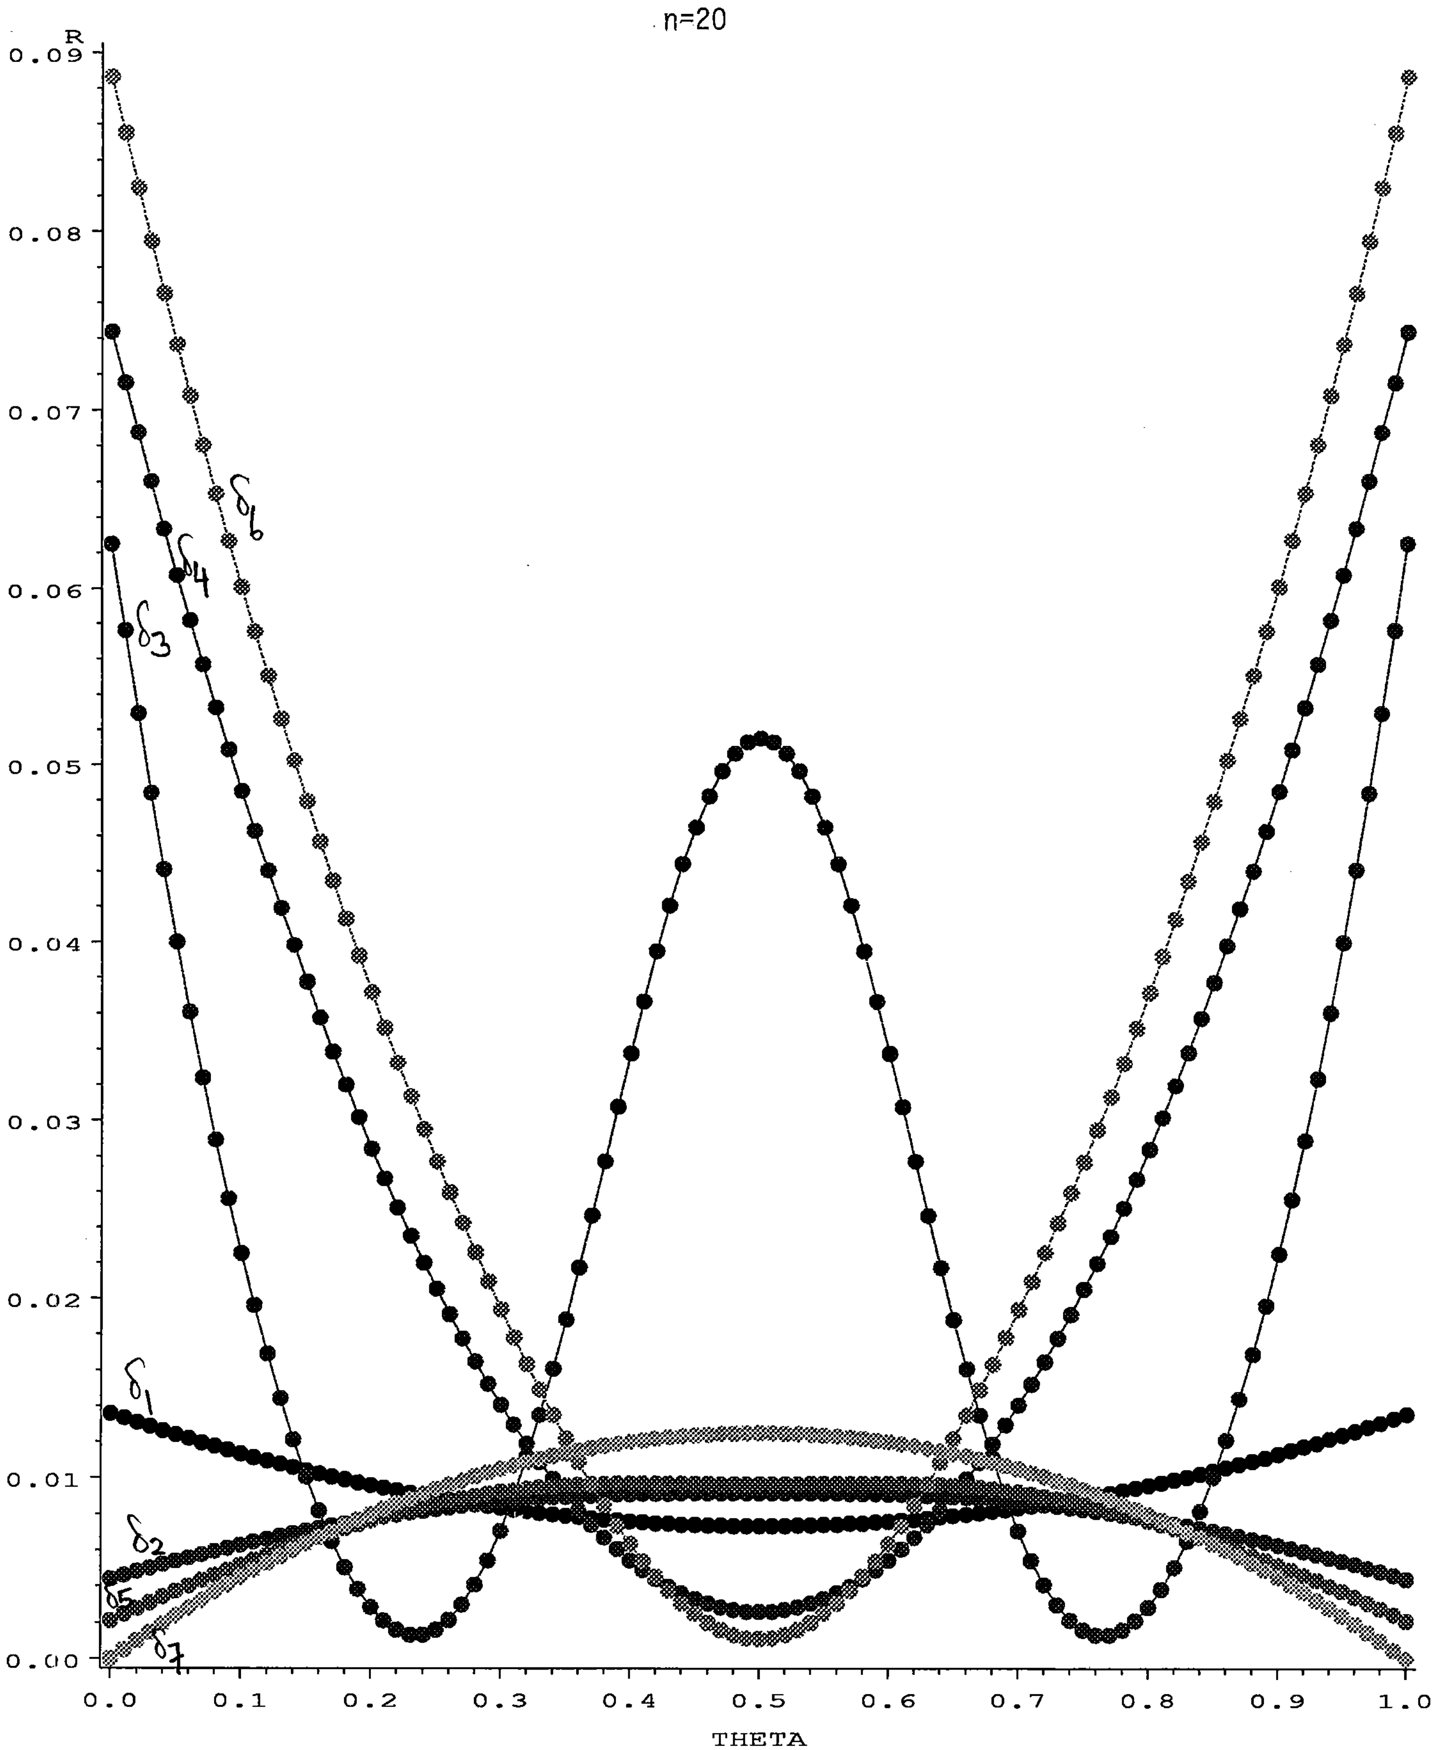

FIGURE 6

The Midpoint Estimates for 7 Classes of Priors
#1: Be(a,a), 1<a<6 #2: Be(a,b), γ=4 b=0.05
#3: symmetric #4: symmetric unimodal
#5: e-contaminated e=0.05 #6: Be(a,b), γ=infinite b = 0.05
#7: unbiased

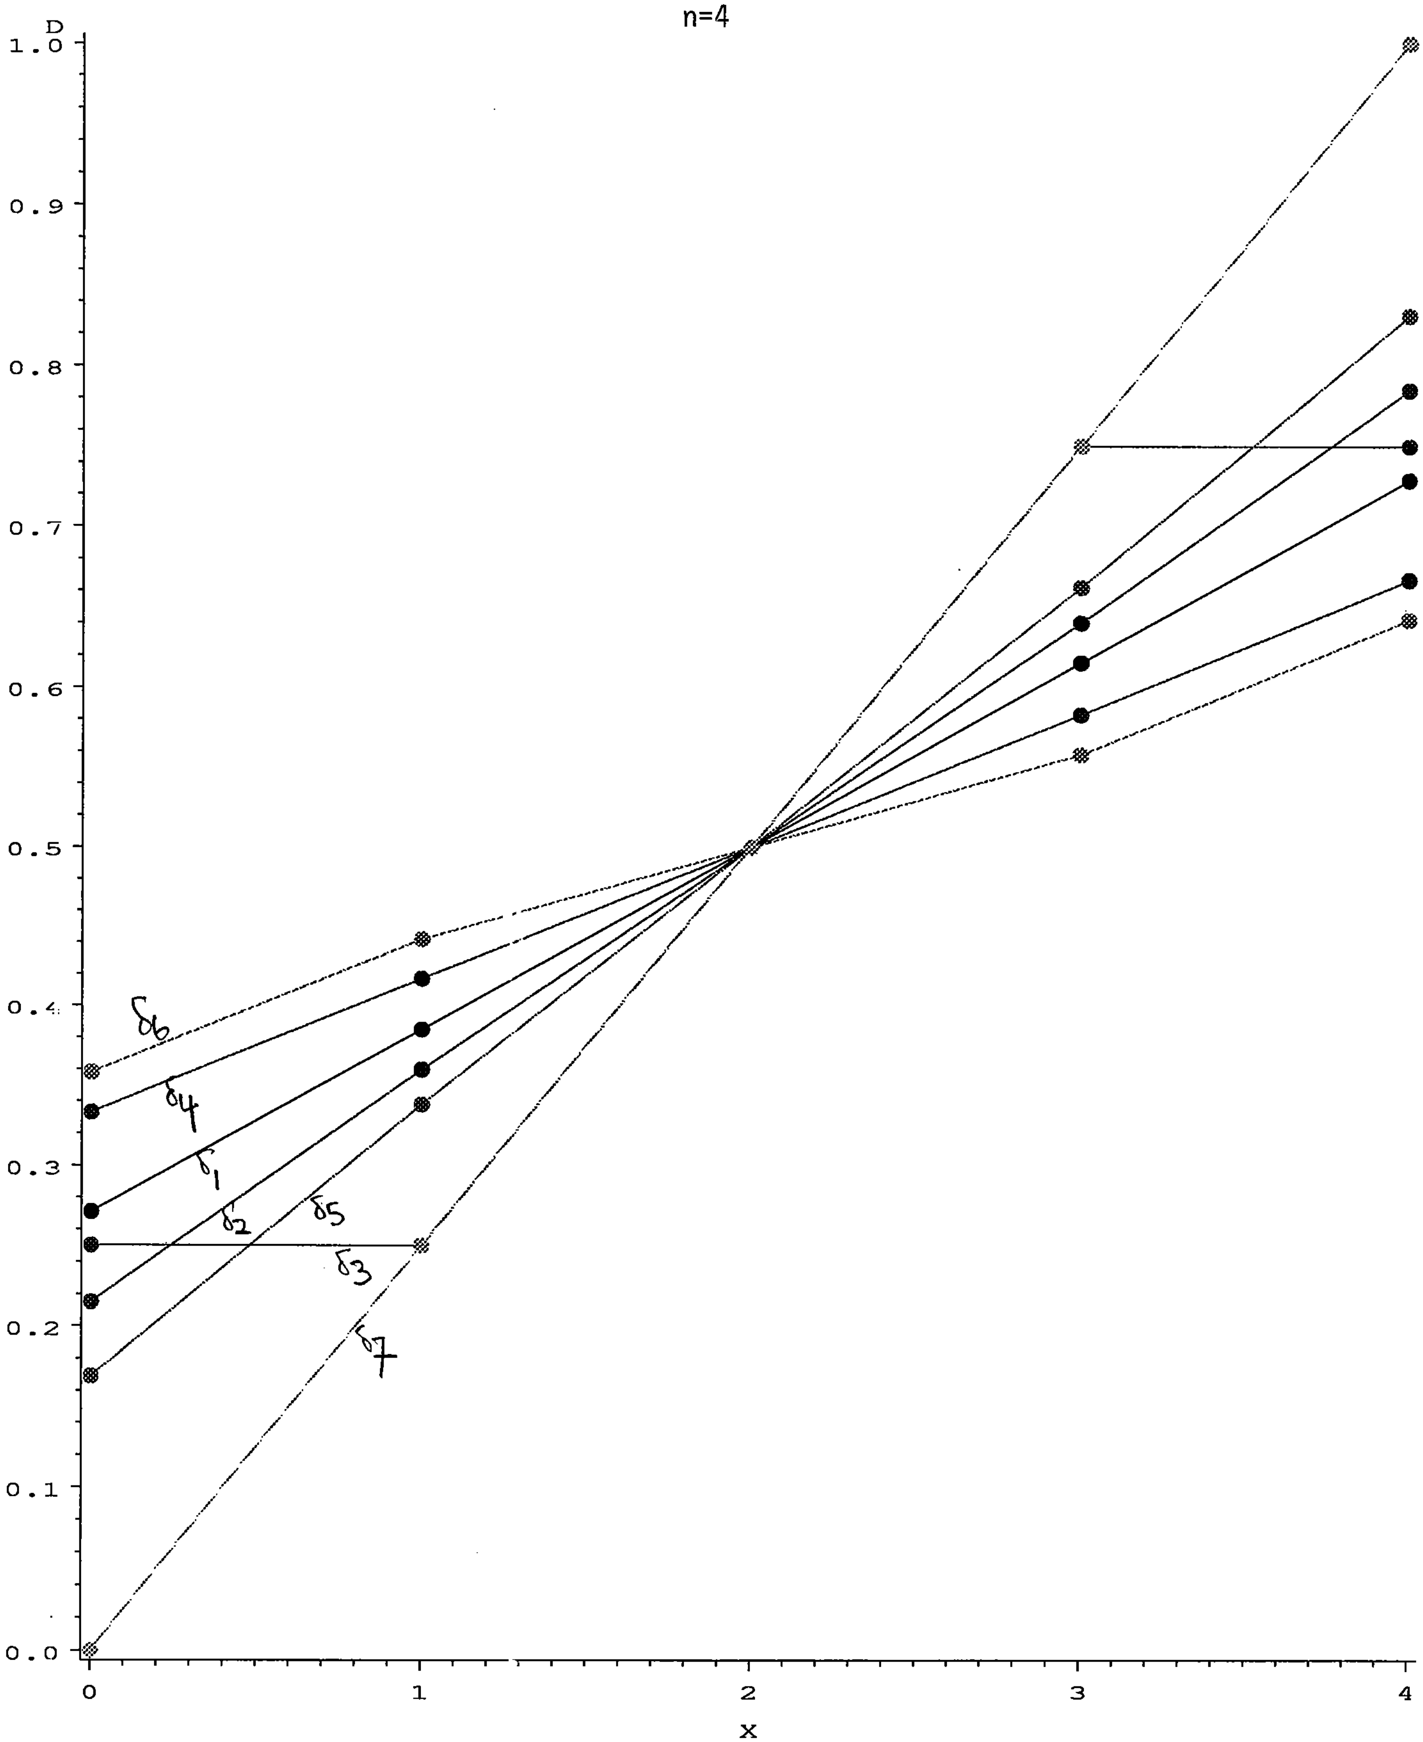

FIGURE 7

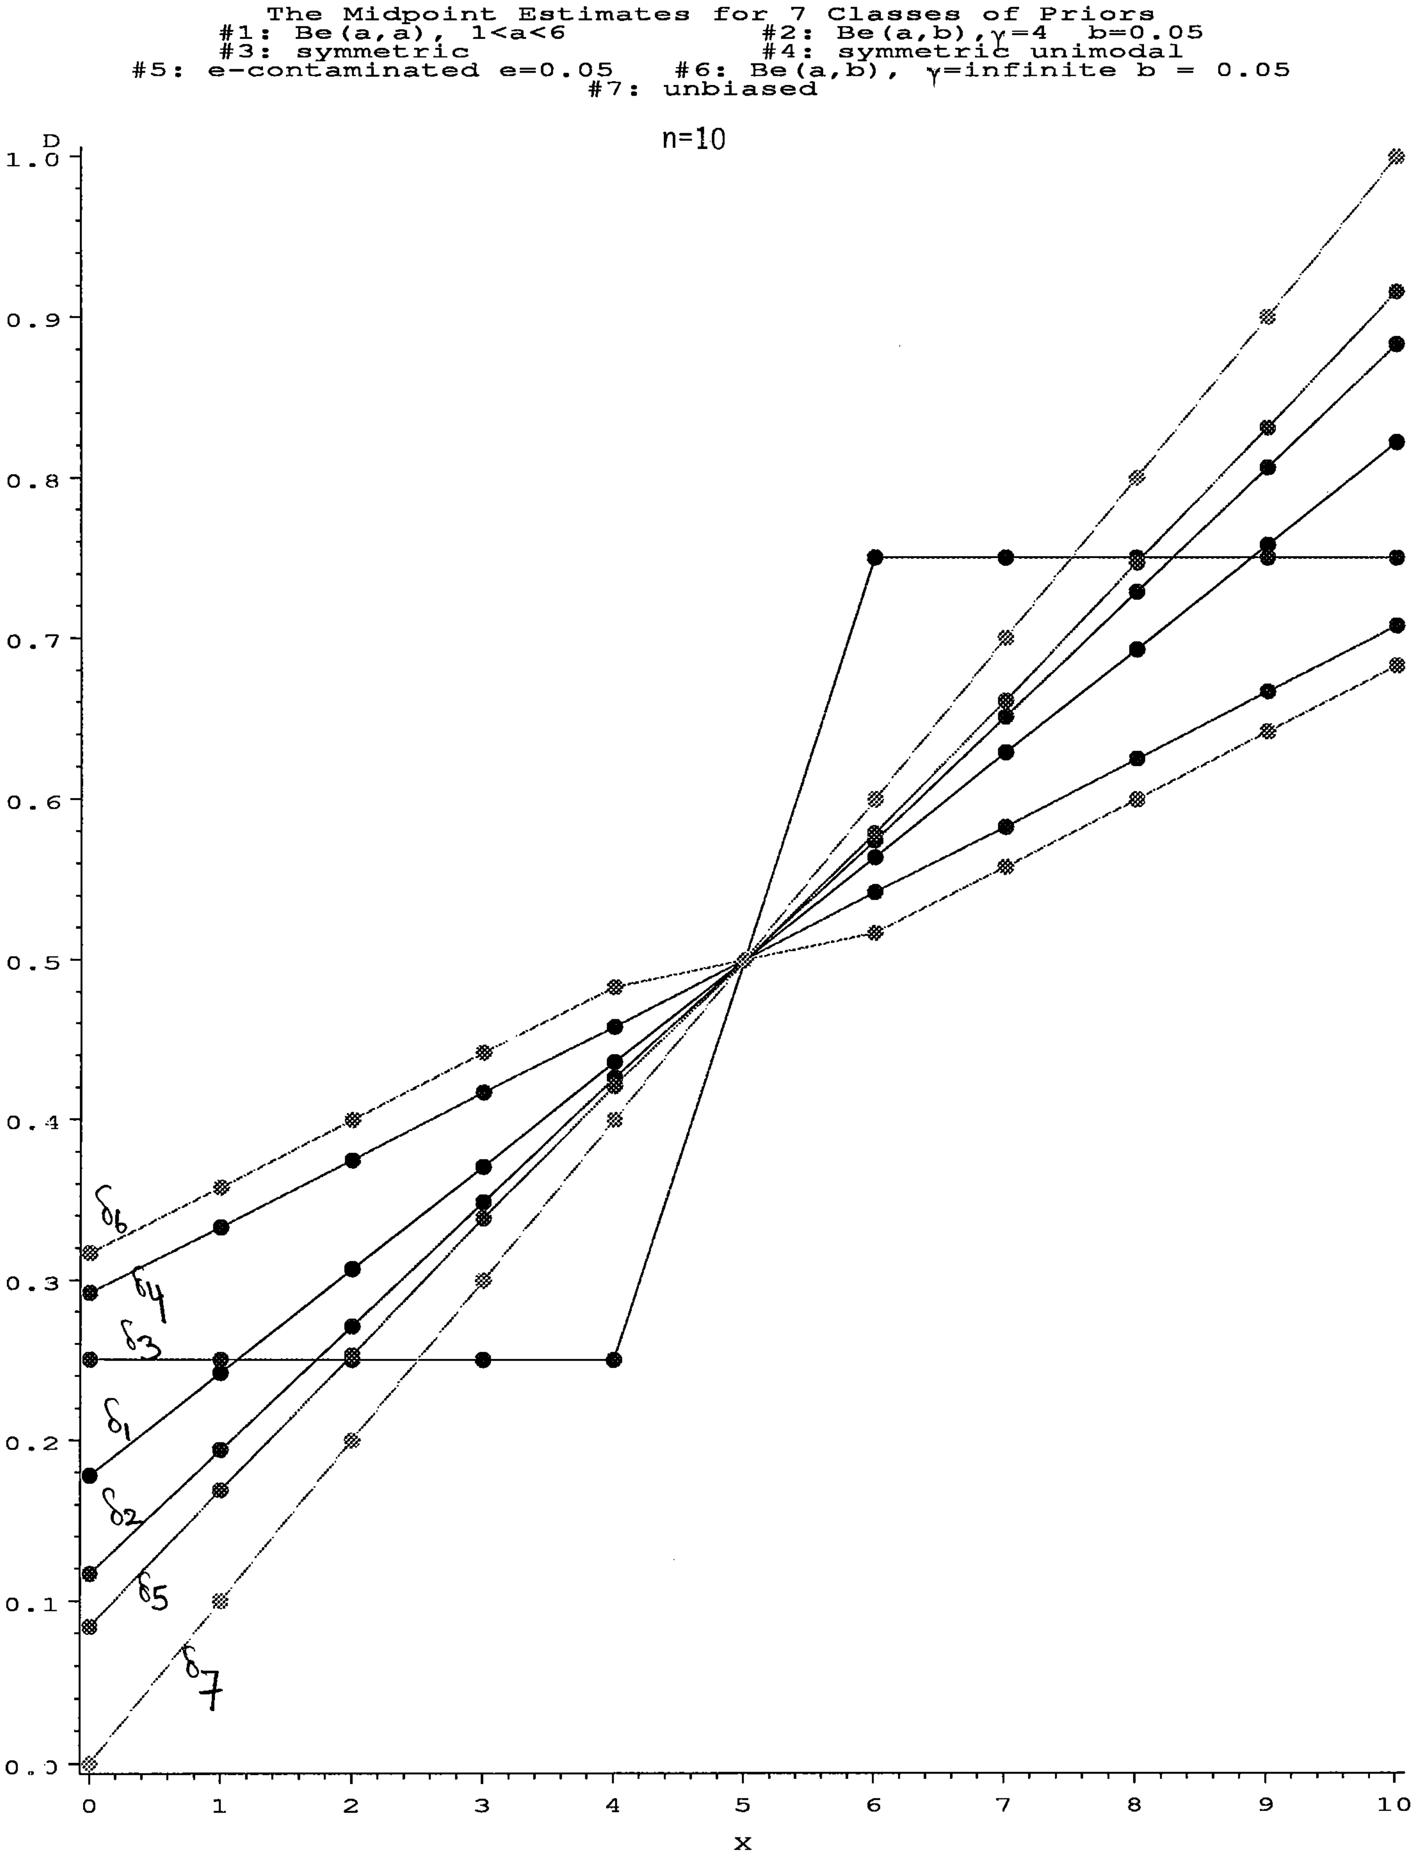
The Midpoint Estimates for 7 Classes of Priors
#1: Be(a,a), 1<a<6
#2: Be(a,b), γ=4 b=0.05
#3: symmetric
#4: symmetric unimodal
#5: e-contaminated e=0.05
#6: Be(a,b), γ=infinite b = 0.05
#7: unbiased
n=10
D
1.0
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0.0
0
1
2
3
4
5
6
7
8
9
10
x
δ6
δ4
δ3
δ1
δ2
δ5
δ7

FIGURE 8

The Midpoint Estimates for 7 Classes of Priors

#1: Be(a,a), 1<a<6 #2: Be(a,b), γ=4 b=0.05
#3: symmetric #4: symmetric unimodal
#5: e-contaminated e=0.05 #6: Be(a,b), γ=infinite b = 0.05
#7: unbiased