It is worth mentioning that the investigation in this section was inspired by the following experimental finding: better quantizers than those described as "optimal" for a Laplacian source by Farvardin and Modestino (column 6 of the table on p. 495 of [26])[5] are easily obtainable using either a randomized search or the technique to be described in Section 2.5.

2.4 Design of Locally-Optimal Alphabet-Constrained Quantizers

When no entropy-constraint is imposed, the design problem is conceptually very simple: locate a relative minimum of a nonlinear scalar function, $D^{(\nu)}$, of a vector $\vec{d}$ (the decision-level vector). In the MSE case the problem simplifies considerably, and efficient methods of solution are described by Lloyd [12] and Max [13]. In the general case, any of a number of well-known procedures for nonlinear optimization can be used [39][40][41]. In the present study, the simple *ad hoc* procedure presented in Appendix A was found to be effective in all cases. Locally-optimal decision and reconstruction levels found by that procedure for the case $K = 8$ are shown against the distortion exponent ν in Figure 2.5.

It is interesting to note that the effect of increasing the distortion exponent is to cause a spreading of decision and reconstruction levels. This is consistent with intuition — as the exponent increases, the distortion measure becomes dominated by the larger instantaneous deviations. For the sources of interest the largest instantaneous deviations occur when the absolute value of the source sample is large. By spreading out the decision and reconstruction levels for large values of ν, the optimization procedure is effectively trading average resolution in the most probable region (near the origin) for a reduction in size and frequency of the largest errors. As will be seen shortly, this is in marked contrast to the case of entropy-constrained quantization, where the locally-optimal decision levels change little with ν for a given entropy.

It should be emphasized that even in the comparatively simple case of alphabet-constrained quantization, little is known in general about the global optimality of

[5] The same numbers appear again in a later paper by the same authors (see Table II of [11]).

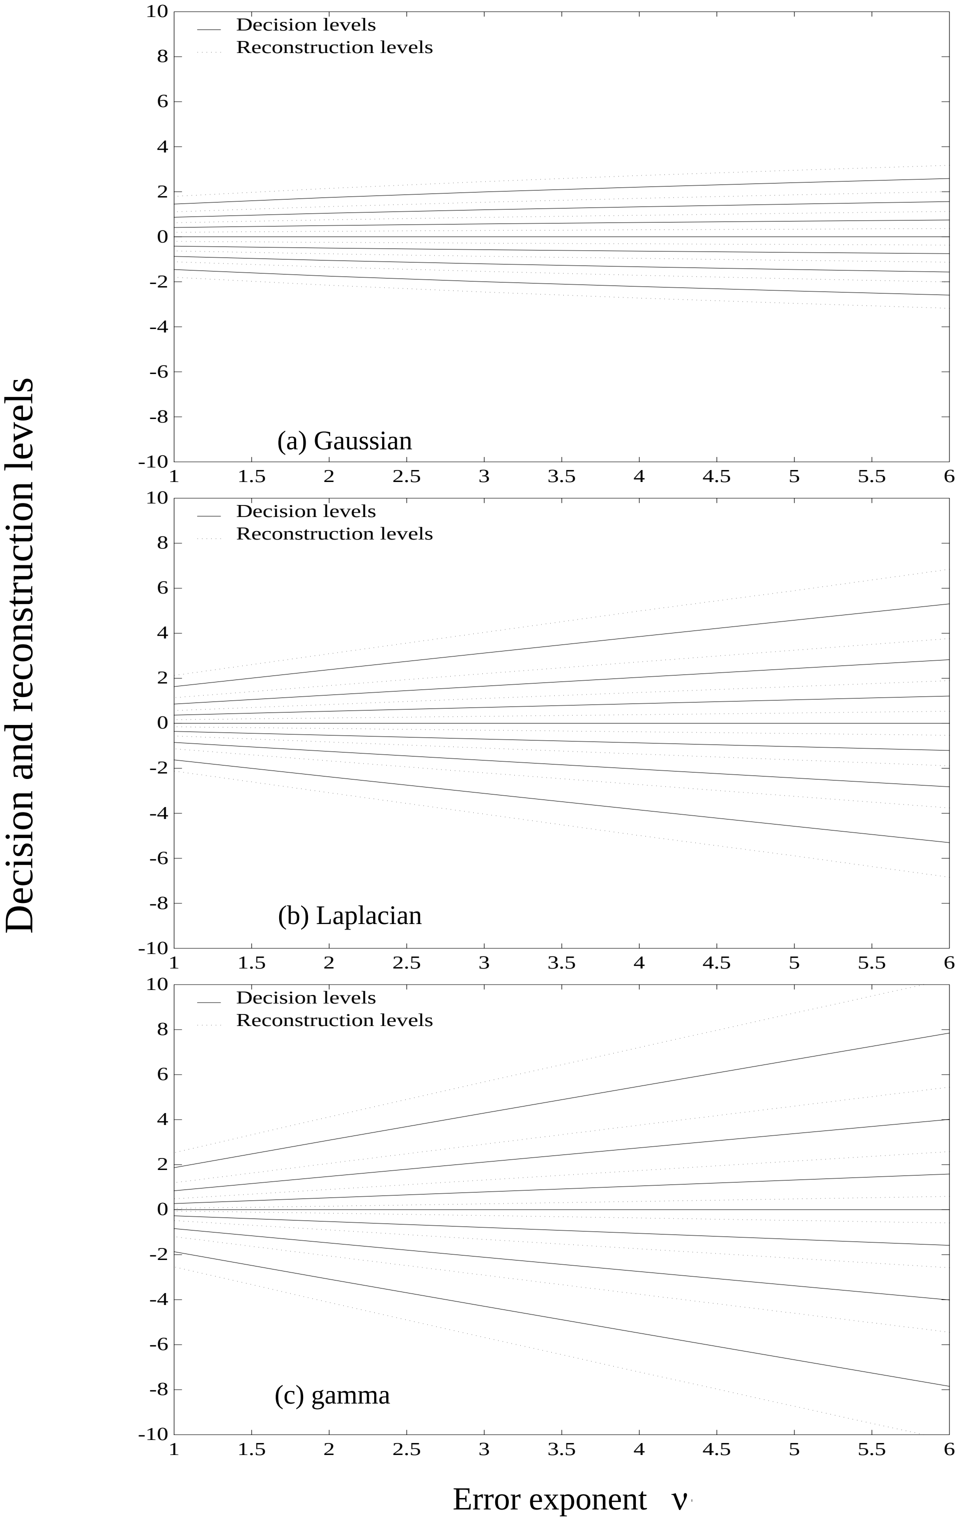

Figure 2.5: Decision and reconstruction levels of locally-optimal alphabet-constrained (8-level) quantizers for (a) Gaussian, (b) Laplacian, and (c) gamma distributed unit-variance sources, versus error exponent ν. The continuous lines are obtained by connecting corresponding decision levels computed for integer values of ν.

a particular quantizer for a given source. However, a few special cases have been successfully handled [27][28][42]. In particular, Kabal [42] shows that in the case of a gamma source and mean-square error, the symmetric locally-optimal alphabet-constrained quantizers obtained in the present study are *not* globally optimal.

It was observed at the beginning of the present section that the problem of obtaining a locally-optimal alphabet-constrained quantizer is conceptually very simple. This is not to say that there are no computational difficulties. The remainder of this section outlines some of the major ones encountered in the present study.

In general, three quantities must be computed repetitively in terms of the d_k's in the course of a gradient-based optimization. They are the r_k's, the distortion $D^{(\nu)}$, and the $\partial D^{(\nu)}/\partial d_k$'s.

Computation of the r_k's is considered first. Recall from Section 2.1, Equation (2.5) that for a given pair of adjacent decision levels d_{k-1} and d_k, r_k must satisfy

$$\frac{\partial}{\partial r_k}\int_{d_{k-1}}^{d_k} |r_k - X|^\nu f_x(X)dX = 0.$$

When ν is even, the above can be rewritten

$$\nu\int_{d_{k-1}}^{d_k} (r_k - X)^{\nu-1} f_x(X)dX = 0. \tag{2.11}$$

In this case, the optimal value of r_k for given a given pair d_{k-1} and d_k can be found by expanding the left-hand side of (2.11) as a $(\nu-1)$-degree polynomial in r_k, finding all roots of the polynomial that lie between d_{k-1} and d_k, and selecting the one that minimizes $D_k^{(\nu)}$. When $\nu = 2$, this approach becomes particularly simple, and yields the well-known result that the optimal r_k is the conditional expected value of x, given that $d_{k-1} < x \leq d_k$.

If ν is large or odd, then the above approach becomes difficult, and the optimal r_k can be found more easily by direct numerical minimization of $D_k^{(\nu)}$.

To evaluate $D^{(\nu)}$ for a given set of decision and reconstruction levels, it is useful first to rewrite the expression (2.4) for $D_k^{(\nu)}$ as

$$D_k^{(\nu)} = \begin{cases} \int_{d_{k-1}}^{d_k} (r_k - X)^\nu f_x(X)dX & \text{for } \nu \text{ even;} \\ \int_{d_{k-1}}^{r_k} (r_k - X)^\nu f_x(X)dX + \int_{r_k}^{d_k} (X - r_k)^\nu f_x(X)dX & \text{for } \nu \text{ odd,} \end{cases} \tag{2.12}$$

eliminating the absolute value operation. By expanding the binomials in (2.12), $D_k^{(\nu)}$ can be expressed as a sum of moment terms of the form

$$G(r_k) \int_a^b X^i f_x(X)dX,$$

where $i = 1, \ldots, \nu$. For all three of the pdf's considered here, nearly analytic expressions (in particular, expressions which involve the error function $\operatorname{erf} x$, but are otherwise analytic) for the moment terms have been obtained with the aid of the *MACSYMA* symbolic algebraic manipulation program. Subroutines written in the C programming language to evaluate the moment terms for Gaussian, Laplacian, and gamma sources are provided in Appendix B.

As mentioned in Section 2.1, computing the partial derivative of $D_k^{(\nu)}$ with respect to d_k and d_{k-1} is complicated by the dependence of the integrand in (2.4) on r_k, which in turn depends on d_k and d_{k-1}. Since r_k is not known analytically, the only recourse is to approximate the partial derivative numerically. In the present study, a central difference approximation has been found to yield adequate results.

While the foregoing discussion has been presented in the context of alphabet-constrained optimization, the observations made about the computational difficulties apply equally well to entropy-constrained optimization, which is the topic of the next section.

2.5 Design of Locally-Optimal Entropy-Constrained Quantizers

A method of finding a locally optimal vector $\vec{d}^*$ of decision levels for an entropy-constrained quantizer is now described. Suppose a vector $\vec{d}_0$ that satisfies the constraint and that is known to be in the vicinity of $\vec{d}^*$ has been found (the problem of finding such a vector will be discussed shortly). This vector can serve as the starting point of a simple search procedure, each iteration of which involves the following steps. First, it is determined which decision level, when modified by a small amount, results in the greatest reduction in $D^{(\nu)}$ for a given reduction in H. That is, d_i is found where

$$i = \arg\max_k |\lambda_k|, \qquad \lambda_k \stackrel{\text{def}}{=} \frac{\partial D^{(\nu)}/\partial d_k}{\partial H/\partial d_k}, \quad k = 1, \ldots, K-1. \tag{2.13}$$

The current vector is adjusted by changing d_i by a small amount to reduce $D^{(\nu)}$. In general, the resulting vector no longer satisfies the entropy constraint; this can be remedied by changing another decision level, d_j, to return the vector to the constraining surface. Here, j is chosen to satisfy

$$j = \arg\max_k |\lambda_k - \lambda_i|. \tag{2.14}$$

If all of the λ_k's happen to have the same sign, which is usually the case in the situations of present interest,[6] then j is given more simply as

$$j = \arg\min_k |\lambda_k|. \tag{2.15}$$

The particular choice of d_i and d_j in this procedure is easy to understand if it is kept in mind that the goal is to reduce $D^{(\nu)}$ while leaving H unchanged. The first modification, adjustment of d_i, results in a relatively large reduction in $D^{(\nu)}$ with a small change in the entropy. This small but undesirable change in entropy can be regarded as a debt incurred in reducing $D^{(\nu)}$; the best way to repay the debt is by adjusting that decision level d_j which, when modified so that the entropy constraint is once again satisfied, results in the least reversal of the reduction of $D^{(\nu)}$ that was obtained by adjusting d_i.

[6] Note that if not all of the λ_k's have the same sign, then $D^{(\nu)}$ can be reduced while stepping away from the constraining surface, and then reduced again while returning to the constraining surface. It is unlikely that this situation would arise except perhaps during the first few iterations in the case of a particularly bad starting vector $\vec{d}_0$.

It is interesting to note that each iteration of the above procedure will continue to yield a reduction in $D^{(\nu)}$ while maintaining H until all of the λ_k's are equal — i.e., improvement will continue so long as entropy can be borrowed and repayed with unequal associated changes in $D^{(\nu)}$. Thus, the natural termination condition for this procedure is simply the Lagrange multiplier condition

$$\lambda_k = \lambda, \qquad k = 1, \ldots, K-1,$$

for some constant λ. Indeed, the above optimization procedure serves as an intuitively satisfying — albeit non-rigorous — derivation of the Lagrange multiplier condition as a necessary condition for local optimality in problems of this type (i.e., problems that require minimization or maximization of smooth functions subject to smooth equality constraints).

The foregoing description allows considerable freedom in selecting the step size used in modifying d_i. In contrast, the step size associated with d_j is completely determined by the requirement that modifying d_j must return the decision-level vector to the constraining surface. For the source distributions of present interest, determination of the step size for d_j is easily accomplished by direct numerical search (via Newton's method, for example). Although this may seem like a computationally expensive operation, in practice it has been found that most of the computational burden actually lies in the calculation of the λ_k's.[7] Applying Newton's method to adjust d_j so that H lies within a few thousandths of a bit of H_0 at each iteration does not contribute appreciably to the overall computation time. It is worthwhile noting that if calculation of the step size *had* proven to be excessively burdensome computationally, then the problem might have been mitigated by temporarily relaxing the constraint to require that H lie in some relatively large neighborhood around H_0. After the decision-level vector had come in the vicinity of the sought locally-optimal vector $\vec{d}^*$, the constraint might once again be tightened. In this way, the optimization

[7] To be more specific, the "bottleneck" appears to be in the computation of the optimal r_k's, which is an intermediate step in the computation of the λ_k's.

process would be allowed to proceed swiftly for most of its duration, slowing only at the final stages where strict compliance with the constraint is essential.

While the above procedure has been employed successfully in a large number of test cases, in certain situations (specifically, those in which both the numerator and denominator of λ_k as defined in (2.13) are extremely small, and those in which the step size used in modifying d_i is relatively large), adjusting d_j in the above manner occasionally fails to return the current decision-level vector to the constraining surface. An alternative, more robust technique takes advantage of the symmetry of the sources of interest. For these sources, the entropy associated with a particular decision-level vector can be adjusted by minutely scaling the vector by a factor close to unity, to effect a small increase or decrease in the entropy. Since scaling preserves the shape of the characteristic staircase function of the quantizer, it approximately preserves the gain in quantizer performance obtained by modifying d_i.[8] As in the case of modifying d_j, the appropriate scaling factor can be found numerically. The computational cost of this method of returning to the constraining surface is roughly K times that of searching in d_j alone; nevertheless, the cost remains negligible in comparison with other operations involved in the optimization. Essentially, scaling provides a reliable and computationally reasonable means of returning the current vector to the constraining surface in a performance-preserving way.

It has been observed that considerable freedom is allowed in selecting the step size for d_i. This is not to say that the choice can or should be made arbitrarily. As is normally the case with numerical optimizations, judicious adaptation of the step size as the search progresses is essential for timely convergence. In the present study, the following strategy has been found to work well. Initially, the step size for d_i is taken to be a predetermined fraction of the minimum spacing between decision levels

[8] Besides the assumption of a symmetric source distribution, the additional assumption being made here is that for a locally-optimal quantizer, the decision levels, regarded as functions of entropy, are continuous. While this may seem obviously true for the sources and distortion measures of interest, the ultimate justification of this assumption lies in the success of the algorithm in finding locally-optimal quantizers, and in the apparent continuity of the decision levels of those quantizers.

in the starting vector $\vec{d}_0$. As the search progresses, the step size is reduced whenever the variation among the λ_k's falls below a certain tolerance. The tolerance is reduced along with the step size, so that the λ_k's are guided towards the common value λ. The procedure terminates when all of the λ_k's are sufficiently close to one another, or (in this case equivalently) when the step size for adjusting d_i has become sufficiently small.

The optimization procedure described above is now summarized. Let $\vec{d}$ denote the current decision-level vector; let δ denote the step size used in modifying d_i; let ϵ denote the tolerance threshold for the variation of λ_k's; and let $\vec{d}_0$, δ_0, and ϵ_0 denote initial values for $\vec{d}$, δ, and ϵ, respectively. Let F_c be the minimum probability required of a quantization region for its bounding decision levels to be considered in the optimization.[9] The algorithm can be summarized as follows:

Step 1: Initialize $\vec{d} = \vec{d}_0$; $\delta = \delta_0$; $\epsilon = \epsilon_0$.

Step 2: For the current $\vec{d}$, let $\mathcal{K} = \{k : \max(F_k, F_{k+1}) > F_c\}$. Compute λ_k for all $k \in \mathcal{K}$.

Step 3: Compute

$$\mu = \frac{\max_{l,m \in \mathcal{K}} \{|\lambda_l - \lambda_m|\}}{\sum_{k \in \mathcal{K}} |\lambda_k|}.$$

If μ is sufficiently small, then halt. Otherwise, if $\mu < \epsilon$, then halve both ϵ and δ and continue. Otherwise, simply continue.

Step 4: Find $i = \arg\max_{k \in \mathcal{K}} |\lambda_k|$. Modify $\vec{d}$ to reduce (improve) $D^{(\nu)}$ by either increasing or decreasing (as necessary) d_i by the amount δ.

Step 5: Scale $\vec{d}$ so that $H = H_0$ and return to Step 2.

Suitable choices for δ_0, ϵ_0, and F_c in the present application have been found to

[9] The restriction that the search involve only those decision levels that have nonnegligible probability is imposed to prevent numerical underflow during computation.

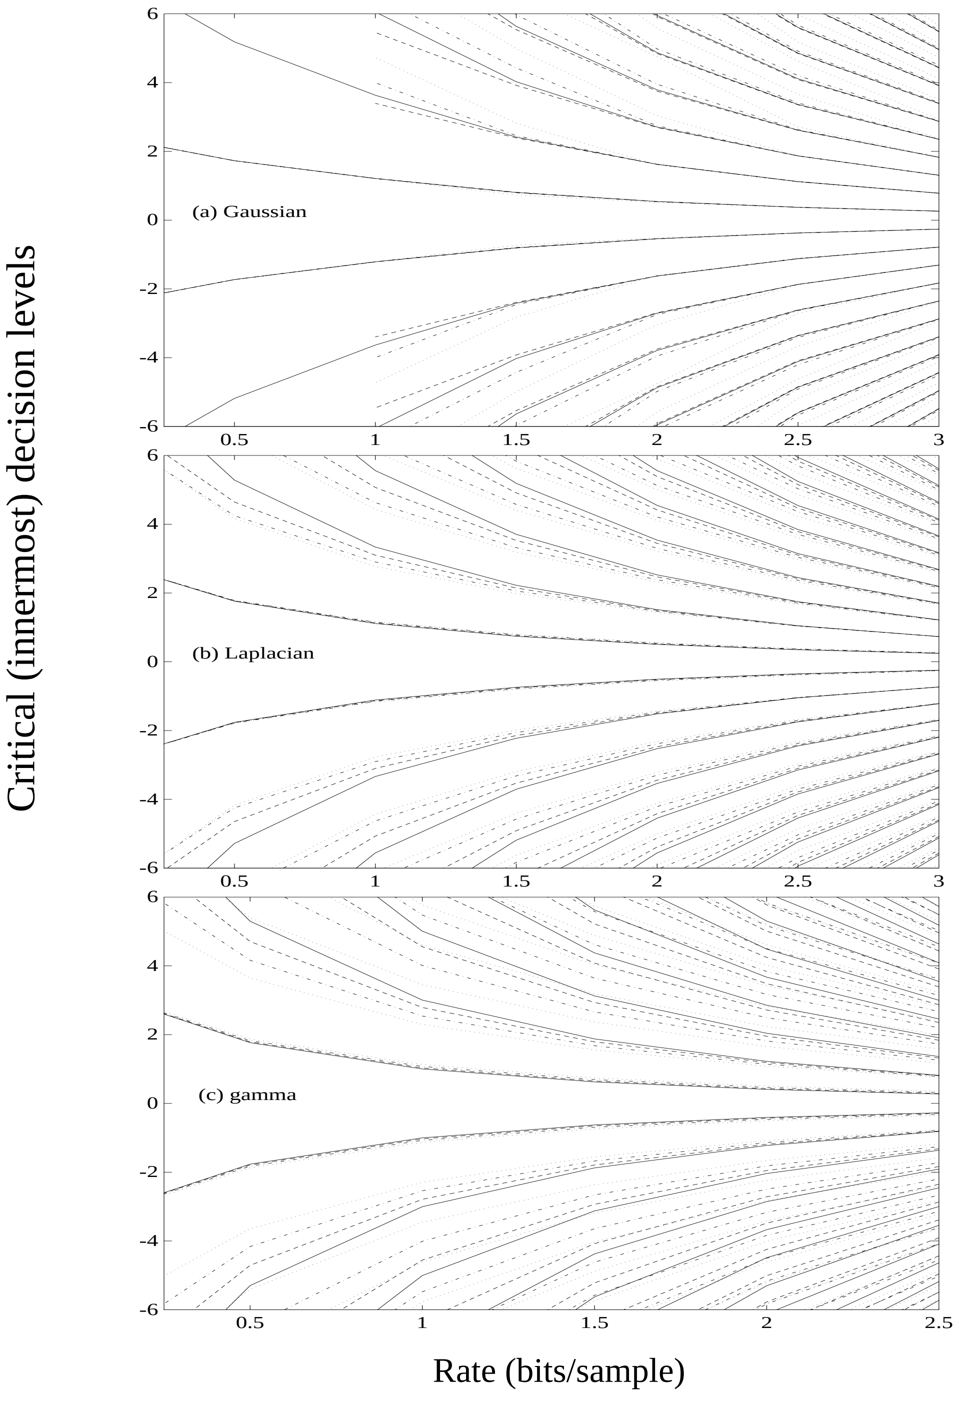

Figure 2.6: Critical (innermost) decision levels of locally-optimal entropy-constrained quantizers for (a) Gaussian, (b) Laplacian, and (c) gamma distributed sources, plotted against rate, for three values of the error exponent ν: 1 (dotted), 2 (dot-dashed), and 6 (dashed). For reference, the decision levels of the uniform midtread quantizer are also shown (solid curves).

be

$$\delta_0 = 0.005 \min_k \{d_k - d_{k-1}\};$$

$$\epsilon_0 = 0.2,$$

$$F_c = \begin{cases} 10^{-6} & \text{if } \nu > 4 \text{ and the source is gamma distributed;} \\ 10^{-8} & \text{if } \nu > 4 \text{ and the source is not gamma distributed;} \\ 10^{-10} & \text{otherwise.} \end{cases}$$

The suitability of these values was determined empirically.

The remaining issue is the selection of the starting point $\vec{d}_0$. Ideally, the starting point would be selected so that the locally-optimal quantizer resulting from the search would also be globally-optimal — i.e., the starting point would be in the vicinity of a global optimum. A case can be made for using the results of the random-search technique described in the following section as starting points. However, the approach taken here is to use a uniform midtread quantizer of the appropriate entropy as the starting point in each case. The reason for this choice is threefold: first, uniform quantization is known to be optimal in the asymptotic case of low distortion [21]; second, such quantizers are known to perform well in the MSE sense at all rates for the sources of present interest [22]; and third, using such a quantizer as a starting point provides a means of testing the degree to which the uniform midtread quantizer is locally-optimal (by comparing it with the result of the search). Figure 2.6 shows the decision levels in the most probable region versus rate, for uniform midtread and locally-optimal quantizers resulting from the search described in this section. Notice that as the rate increases, the decision levels for the locally-optimal quantizers tend to merge into those for the uniform quantizers. Further note that the locally-optimal quantizers that are the *least* uniform are those in the low-rate region for a gamma source, with mean-absolute distortion measure. The distortion-rate performance for these locally-optimal quantizers will be presented in Section 2.8. It will be seen there that although the locally-optimal quantizers are occasionally substantially nonuniform, the performance of uniform midtread quantization is always nearly that of locally-optimal quantization.

2.6 Random Search for Optimal Entropy-Constrained Quantizers

In the previous section a technique was described for searching a small region of the constraining surface in decision-level space for a relative minimum of $D^{(\nu)}$. A quantizer resulting from that technique satisfies necessary conditions for optimality. In contrast, the present section describes a technique wherein a very large portion of the constraining surface is explored in an effort to gain insight into the nature of *global* optima. The approach is to generate a very large number of candidate decision-level vectors at random, and to record those with the best performance.

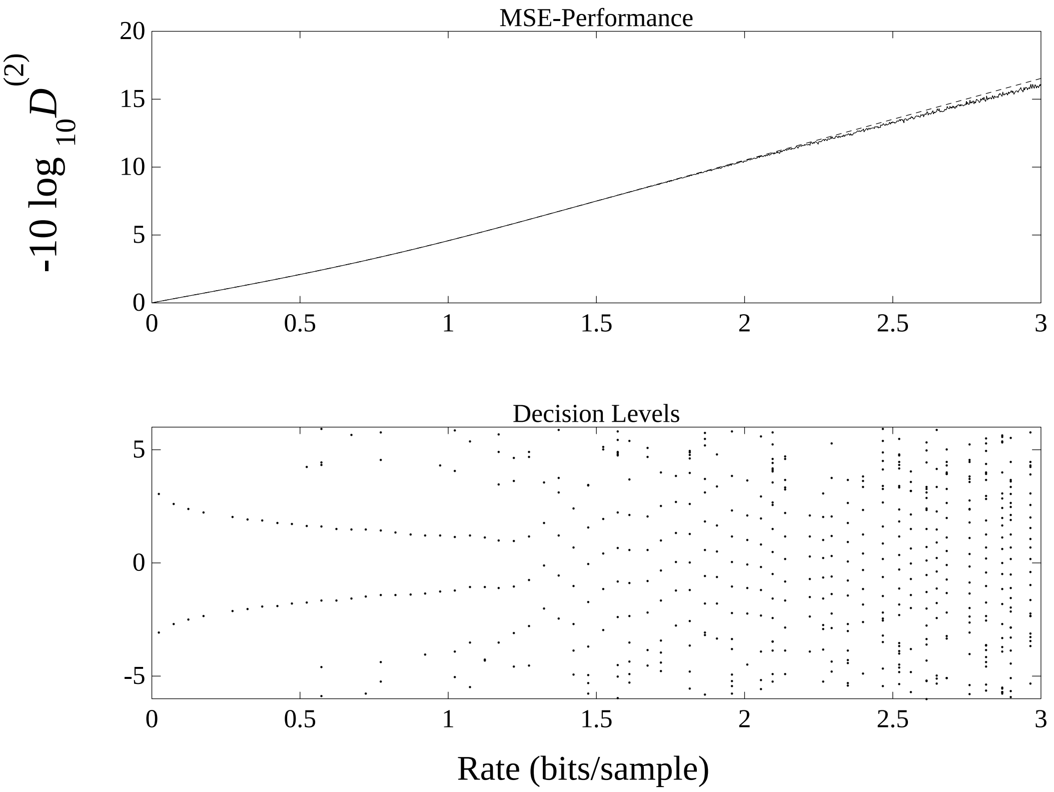

Figure 2.7: Results of random design for a Gaussian source in the case of mean-square error ($\nu = 2$). Performance (rate versus MSE) is shown in (a), while a scatter plot of the decision levels is shown in (b). For reference, the performance of uniform midtread quantizers is shown in (a) as well (dotted line). The number of candidate quantizers considered is 2×10^6.

The success of this approach depends in large part on striking a reasonable compromise in choosing the distribution of the candidate vectors. The candidate vectors

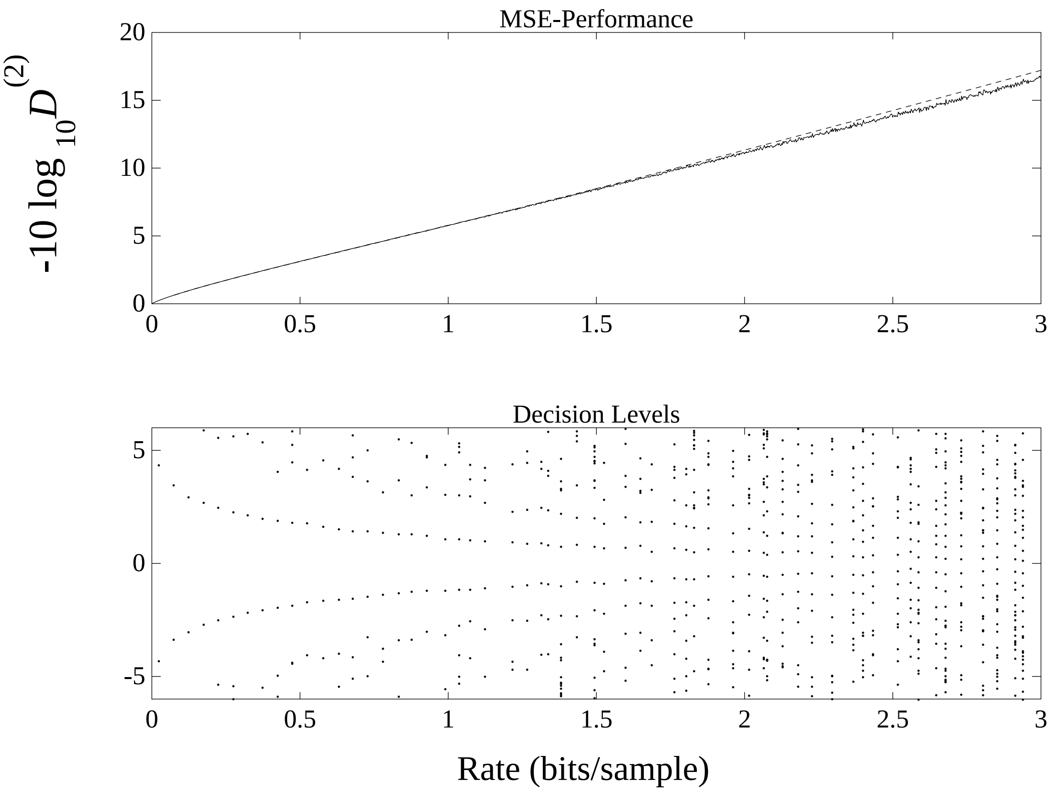

Figure 2.8: Results of random design for a Laplacian source in the case of mean-square error ($\nu = 2$). Performance (rate versus MSE) is shown in (a), while a scatter plot of the decision levels is shown in (b). For reference, the performance of uniform midtread quantizers is shown in (a) as well (dotted line). The number of candidate quantizers considered is 2×10^6.

must be generated with a distribution that covers all regions of decision-level space likely to contain optima. The distribution should not be so concentrated in one region that it biases the outcome of the search, nor should it be so diffuse that a region containing an optimum is only sparsely covered by the search. In the present study, the following approach is used. First, it is assumed that decision levels falling outside of the range $[-W, W]$ for the zero-mean, unit-variance sources of interest have a negligible effect on both the distortion and the entropy, where $W = 6$ in the Gaussian case, $W = 12$ in the Laplacian case, and $W = 18$ in the gamma case. These numbers have been determined empirically. The rate axis between zero and 5 bits/sample is divided into M bins, and a $K = 4$ quantizer is designed by generating three random numbers (uniformly distributed between $-W$ and W), sorting them, and assigning

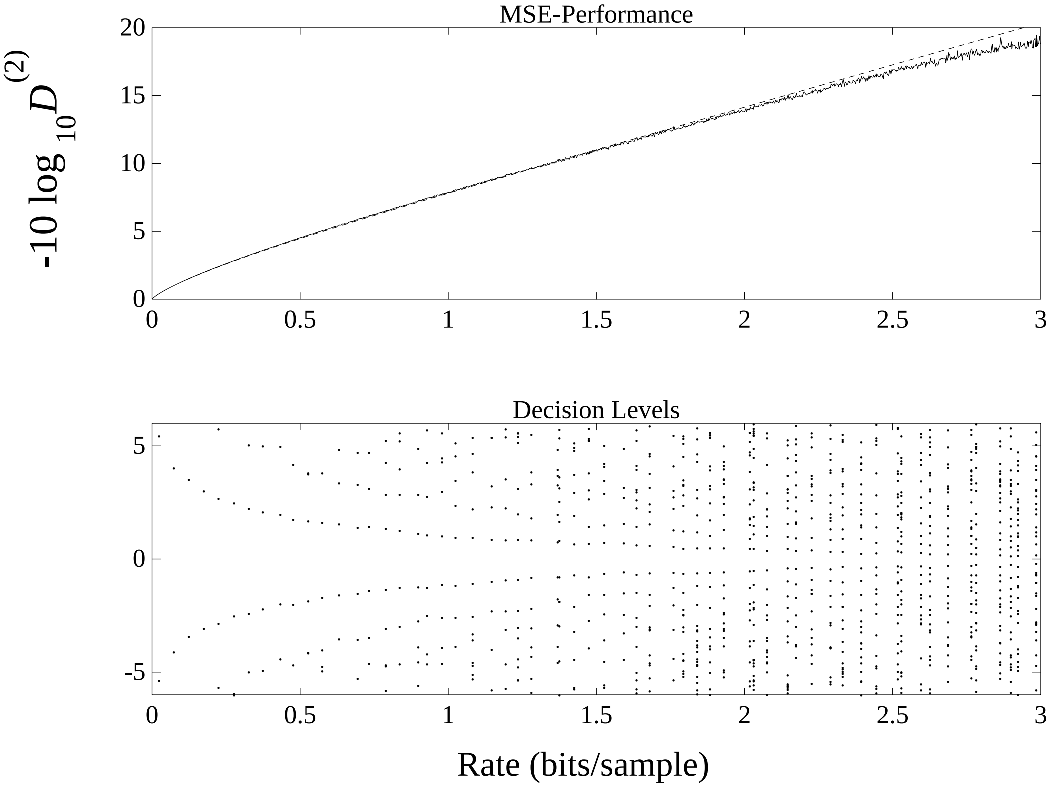

Figure 2.9: Results of random design for a gamma source in the case of mean-square error ($\nu = 2$). Performance (rate versus MSE) is shown in (a), while a scatter plot of the decision levels is shown in (b). For reference, the performance of uniform midtread quantizers is shown in (a) as well (dotted line). The number of candidate quantizers considered is 2×10^6.

them to the decision levels d_1, d_2, and d_3. Both the rate and distortion are computed, and the appropriate rate-bin is determined. If the ratio of the distortion (in dB) to the rate is lower than the lowest previously encountered ratio for that bin, then the performance and/or the decision levels are recorded. The rate is recorded for each bin as well. Next, the number of levels K is doubled and the process is repeated, yielding a quantizer with higher entropy. This process is carried out repeatedly, each time with double the number of levels as the previous time, until the measured entropy of the quantizer exceeds a predetermined threshold (in the present study, 5.0 bits/sample).

Selection of the number of entropy-bins M involves a compromise that is similar to the one made in selecting the distribution of the candidate vectors. If the bins are very small, then the entropy-resolution will be high, but the recorded value will

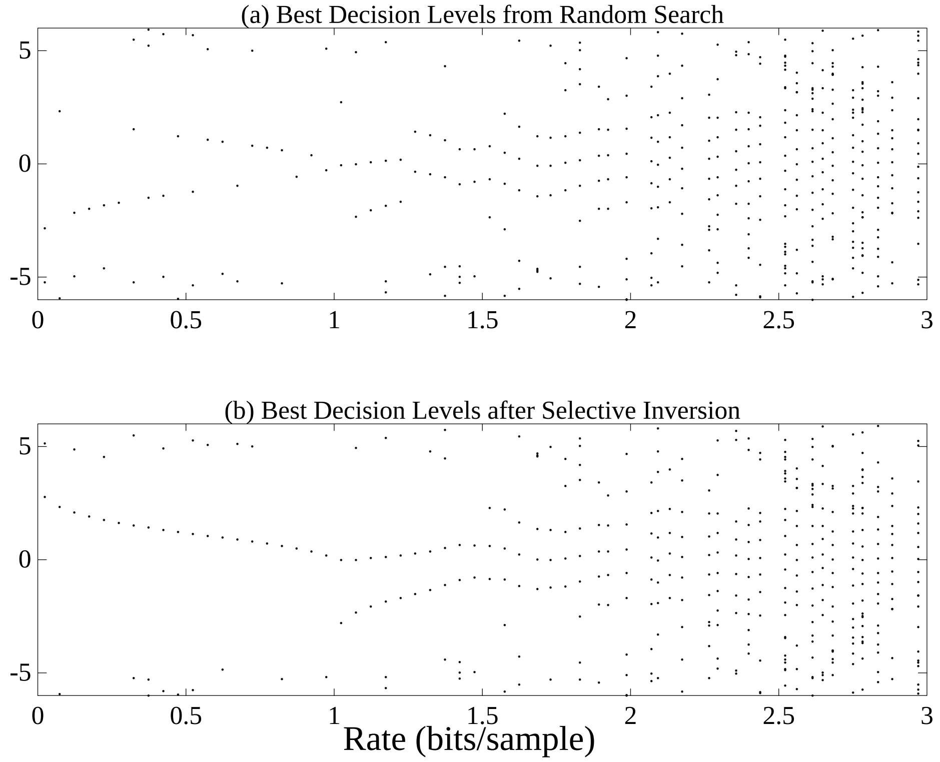

Figure 2.10: Scatter plot of best decision levels found by random search for a Gaussian source and mean-absolute distortion measure ($\nu = 1$). The raw results of the search are shown in (a), while in (b) some of the decision-level vectors have been inverted to obtain more continuous trajectories. In obtaining (b), the choice of overall polarity was arbitrary; an equally performing set of quantizers can be obtained by inverting the entire graph. The important characteristic to note here is the obvious asymmetry in the low-rate region. The number of candidate quantizers considered is 10^6.

have been drawn from a smaller population in each bin, reducing the quality of each recorded point. Alternatively, if the bins are large, then the entropy-resolution will be low, although the recorded value for each bin will have been selected from a large population and will therefore be of higher quality. In the present study, two different values of M are used — one in estimating optimal performance ($M = 1000$), and another in characterizing the corresponding optimal distribution of decision levels ($M = 100$). The results are shown in Figures 2.7– 2.9 for the MSE distortion measure, where in each case a total of 2×10^6 random candidate vectors are generated and considered. In each figure, (a) shows an estimate of optimal performance while (b) shows the corresponding scatter plot of decision levels. The uneven horizontal spacing in the scatter plots is a result of the best decision-level vectors occasionally being found near the edges of entropy-bins, while the falling-off of performance at higher bit rates reflects the inability of the random-search technique to adequately explore the decision-level space as the effective dimensionality of the problem (the number of important decision levels) increases. Note that in both the high-rate and low-probability regions of the graphs, the points exhibit reduced structure. This is expected, since the high-rate region is not thoroughly explored by the search and the low-probability regions are relatively unimportant in determining the performance of the quantizers.

The data in figures 2.7- 2.9 suggest that in the low-rate region, globally-optimal quantizers in the MSE case for all three sources resemble uniform-midtread quantizers. For Laplacian and gamma sources, the optimal quantizers continue to resemble uniform midtread quantizers at higher rates, but in the Gaussian case, the pattern abruptly changes near the rate of 1.3 bits/sample (see Figure 2.7(b)). This phenomenon persists when the number of candidate decision-level vectors is increased from 2×10^6 to 10^7, and when the distortion exponent is set to be greater than 2. Further study is required before meaningful conclusions can be drawn about this phenomenon.

The mean-square error criterion is used in the large-scale randomized searches because it allows fast computation, which is essential if several million quantizers are to be evaluated. For non-MSE cases, the random-search technique has been applied using fewer candidate decision-level vectors. In all cases but one, a quantizer resembling the uniform midtread quantizer appears as the best at low rates. That one exception is the case of Gaussian mean-absolute error (see Figure 2.10). This case is remarkable because the globally-optimal quantizer at low rates appears to be asymmetric and unique (allowing for sign change). Figure 2.10 (a) is the raw scatter plot; Figure 2.10 (b) was obtained by selectively inverting the decision-level vectors of Figure 2.10 (b) as needed to obtain continuous trajectories. While asymmetric optimal quantizers have been previously reported for symmetric sources in the alphabet-constrained case

[37][42] and the joint alphabet- and entropy-constrained case [26], none has previously been reported in the purely entropy-constrained case. The result strikes the author as mildly counterintuitive.

An interesting feature of the decision-level scatter plots is the occasional doubling-up of decision levels — that is, two or more decision levels lying very close to one another. This is to be expected, since for any given quantizer, a new decision level can be inserted infinitesimally close to an existing decision level while changing neither the entropy nor the distortion (since the decision region bounded by the two levels has infinitesimal probability). Hence the appearance of multiple decision levels comes as no surprise; in fact, it testifies to the efficacy of the random-search technique in exploring the vast search space.

It is worthwhile noting that the foregoing random search strategy would *not* be as appropriate a means of characterizing the best-case performance of alphabet-constrained quantizers, for the following reason. In the alphabet-constrained problem, the characteristic staircase function of the quantizer must have both the right shape and the right scaling for the quantizer to perform well. Consequently, the near-optimal region of the search space is likely to be relatively small and therefore difficult to explore by the random technique described here. This is in contrast to the entropy-constrained problem where, for a given shape, the scaling determines the entropy. In this case, the random quantizers in any particular neighborhood of entropy automatically have the proper scaling so that the search space is explored efficiently.

In concluding the discussion of the random-search technique, it should be emphasized that the results of the search can provide evidence to support a hypothesis about global optima; they can never prove the the hypothesis. Still, the evidence can be extremely compelling, particularly when the number of candidate decision-levels vectors is very large.

2.7 Calculation of Distortion-Rate Functions

In the previous chapter it was noted that for a given source and distortion measure, there exists a function, called the rate-distortion function $R(D^*)$, which is a lower bound on the rate required of any encoder to represent a source with distortion not exceeding D^*. As a theoretical upper bound on performance, the rate-distortion function provides a natural reference by which to gauge the performance of a given encoding system. For example, the performance of a given system can be specified in terms of the rate it requires in excess of $R(D^*)$ to represent the source with distortion D^*. Alternatively, the performance can be specified as the amount of distortion incurred in excess of D^* when the rate of encoding is $R(D^*)$. That is, it can be specified as the amount of distortion in excess of the distortion-rate function (the inverse function of the rate-distortion function). The latter approach is assumed for the remainder of the chapter.

Although distortion-rate function tables for the sources of interest have been published in a few cases,[10] they do not appear to have been published for a range of values of the difference distortion exponent ν. Consequently, the distortion-rate function values in these cases had to be computed in the present study. This section describes the approach.

First, it should be noted that in general, direct calculation of the distortion-rate function of a continuous-amplitude source is extremely difficult. In some cases, the Calculus of Variations can be successfully applied [6][43]. In general, however, the distortion-rate function is much more easily bounded than computed [5, pp. 639-640].

In contrast, in the case of a discrete-amplitude source, an algorithm devised independently by Arimoto [44] and by Blahut [45][46] always converges to the distortion-rate function. In the present study, the approach to calculating the distortion-rate

[10] Analytical expressions for the distortion-rate function are available in the Gaussian MSE case [7, p. 477] and in the Laplacian mean-absolute error case [6, p. 95], while numerically obtained values are available in the Laplacian and gamma MSE cases [5][25]. The author is not aware of published values in the other cases, though they undoubtedly exist somewhere in the literature.

functions for a continuous source is to approximate the source by a discrete one, and then to apply the Arimoto-Blahut algorithm to the approximation. Although no attempt is made here to prove the mathematical validity of the technique, the following plausibility argument is offered. As the size of the discrete approximating alphabet gets very large, the probability of any particular letter approaches zero, and the probability that a letter will fall in any particular interval approaches the corresponding value for the continuous source. Thus, as the approximation becomes fine, it becomes increasingly difficult to determine statistically whether a particular sample originated from the continuous source or from its discrete approximation. Since the measured performance of an encoder is actually a sample statistic, as the approximation improves, the performance of an encoder applied to the discrete source must approach the performance of the same encoder applied to the continuous source. This must be true for any means of encoding, including one with performance that approaches the distortion-rate bound. Therefore, the distortion-rate functions for the discrete and continuous sources must approach one another as the approximation improves.

A remaining issue is the rate of convergence of the distortion-rate function of the discrete source to that of the continuous source as the approximation improves. Here, the issue is addressed experimentally. Two parameters are set — the size of the source and reproducing alphabets, and the range of values that the letters cover. Both are increased until no appreciable effect is observed on the distortion-rate function over the range of rates of interest. It has been determined experimentally that for the particular zero-mean, unit-variance sources of interest, in the most difficult case considered (that of a gamma source at a rate of 4 bits/sample and with $\nu = 6$), no appreciable change results when the number of letters in the discrete approximating source is allowed to increase above roughly 500, or when the range of values is allowed to exceed roughly ± 25. The results of the technique described above agree with those previously published for the Gaussian, Laplacian, and gamma MSE case [5], and the Laplacian mean-absolute error case [6].

In the next section, the distortion-rate functions computed in the above manner are used as a reference by which to gauge the performance of alphabet-constrained and entropy-constrained quantization for Gaussian, Laplacian, and gamma sources.

2.8 Performance of Locally-Optimal and Uniform Quantization

This section presents distortion-rate performance curves for uniform and locally-optimal entropy-constrained quantization.

Figure 2.11 shows the distortion $D^{(\nu)}$ as a function of rate for uniform and locally-optimal entropy-constrained quantization, in decibels (dB) above the minimum possible distortion $D^{(\nu)*}$ given by the distortion-rate function. Several observations can be made. The most striking feature of the graphs is the near local optimality of uniform quantization at all rates, particularly in the Gaussian case. The largest difference between the performance of uniform and locally-optimal quantization is in the gamma mean-absolute error case; even there, the maximum difference is merely about 0.3 dB.

Another observation is that the performance of entropy-constrained quantization appears to approach the distortion-rate bound as the distortion exponent increases, for all rates. This had been previously established only in the asymptotic case of low distortion [21].

It was observed in the previous section that based on the results of a randomized search, the globally optimal quantizers appear to be asymmetric at low rates in the Gaussian mean-absolute error case. The same conclusion can be reached in a different way. In the Gaussian mean-absolute error case, contrary to all other cases investigated, the best uniform quantizer at rates slightly greater than 1 bit/sample can be shown (by varying ψ) to be the midriser quantizer, with arbitrarily large Δ as the rate approaches unity from above. However, as was observed in Section 2.2, the midriser quantizer can never have a rate of less than 1 bit/sample. At rates slightly below 1 bit/sample, there are two possibilities: either the optimal uniform quantizer becomes the midtread quantizer, which is the only other symmetric uniform quantizer, or it

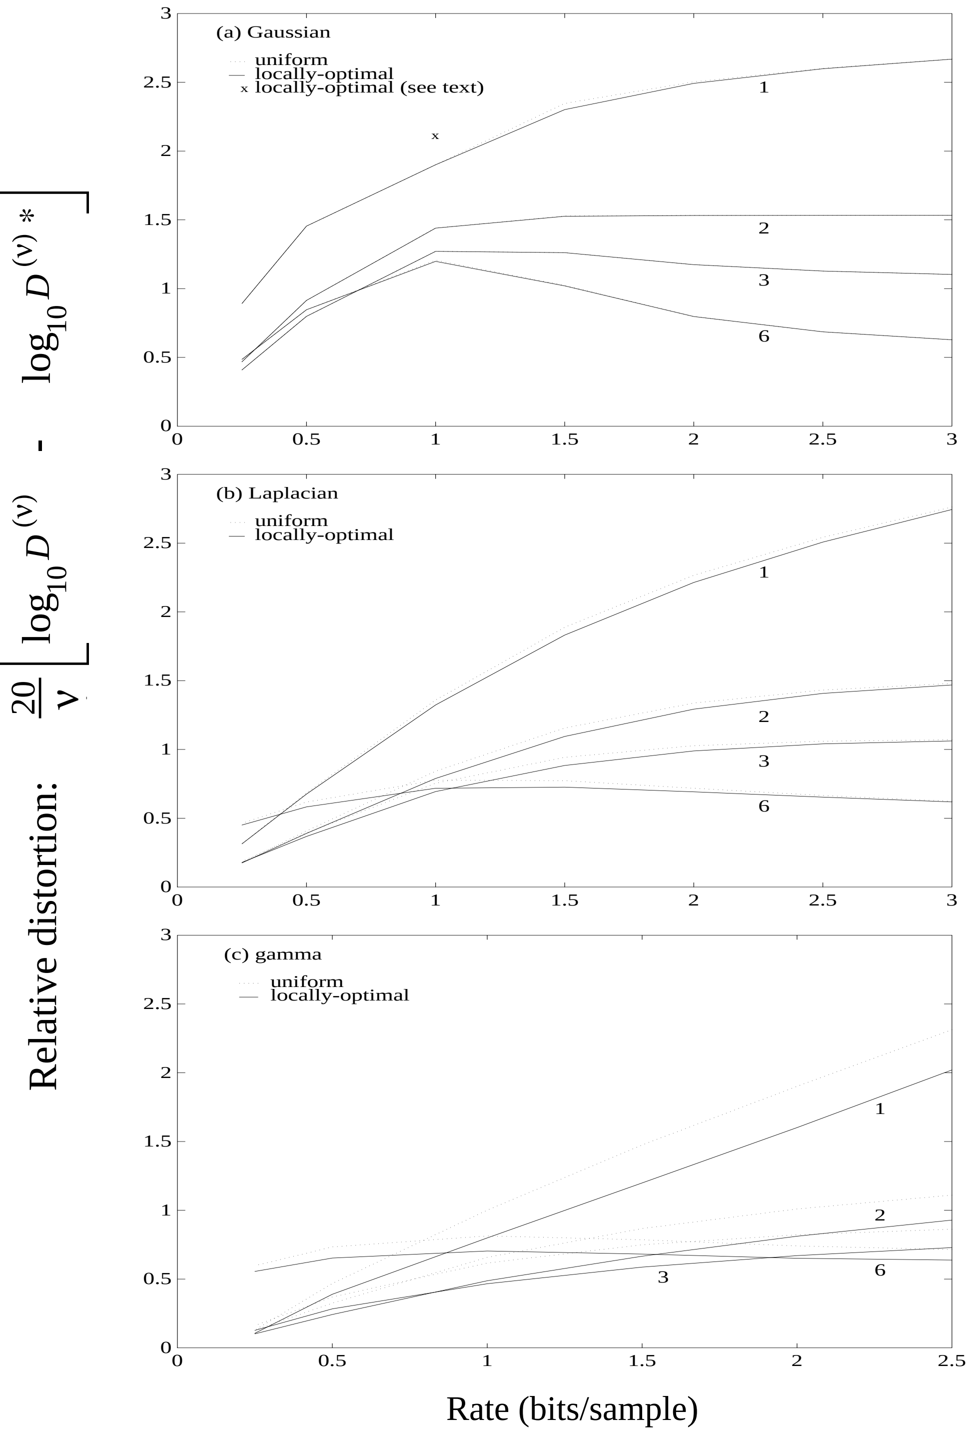

Figure 2.11: Relative distortion (expressed in dB above the distortion-rate bound) versus rate of locally-optimal (solid) and best uniform (dotted) quantizers as a function of entropy, for four different values of the error exponent ν: 1, 2, 3, and 6 (as indicated next to each solid curve).

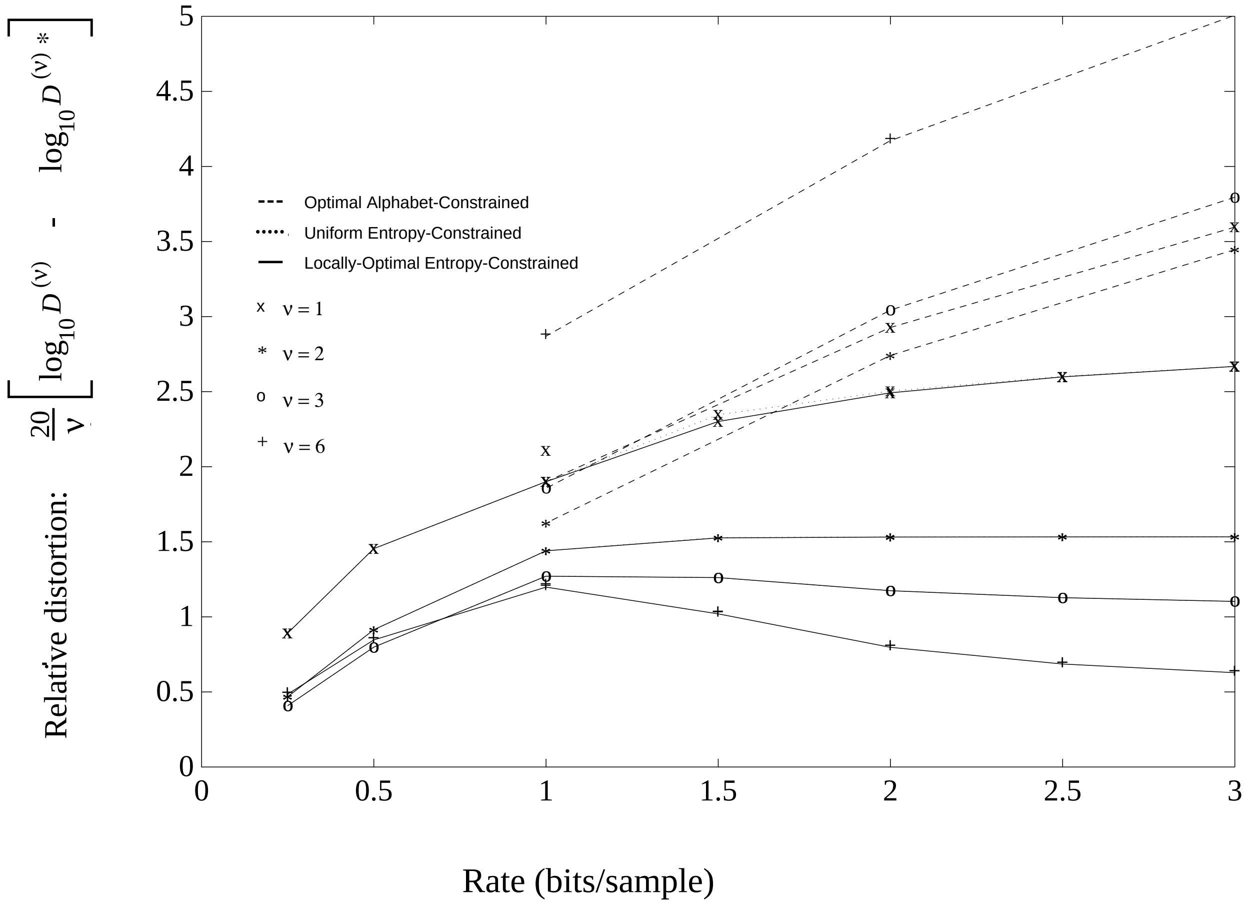

Figure 2.12: Performance (rate versus dB distortion in excess of distortion-rate bound) of locally-optimal alphabet-constrained (Lloyd-Max) quantization relative to uniform and locally-optimal entropy-constrained quantization, for a Gaussian source.

becomes asymmetric. The performance of the uniform midtread quantizer, marked by an 'x' in Figure 2.11(a), is significantly worse at 1 bit/sample than that of the midriser at the same rate (solid line). If the uniform midtread quantizer were optimal just below 1 bit/sample, then a discontinuity would have to exist in the performance curve for the optimal entropy-constrained uniform quantizer, which is clearly not possible. Therefore, the optimal uniform quantizer at rates below 1 bit/sample for the Gaussian mean-absolute error case must be asymmetric. The argument generalizes to nonuniform quantization by noting that the uniform midriser quantizer is (trivially) locally-optimal at 1 bit/sample and that there exists a locally-optimal quantizer at the same rate that is almost identical to the uniform midtread quantizer.

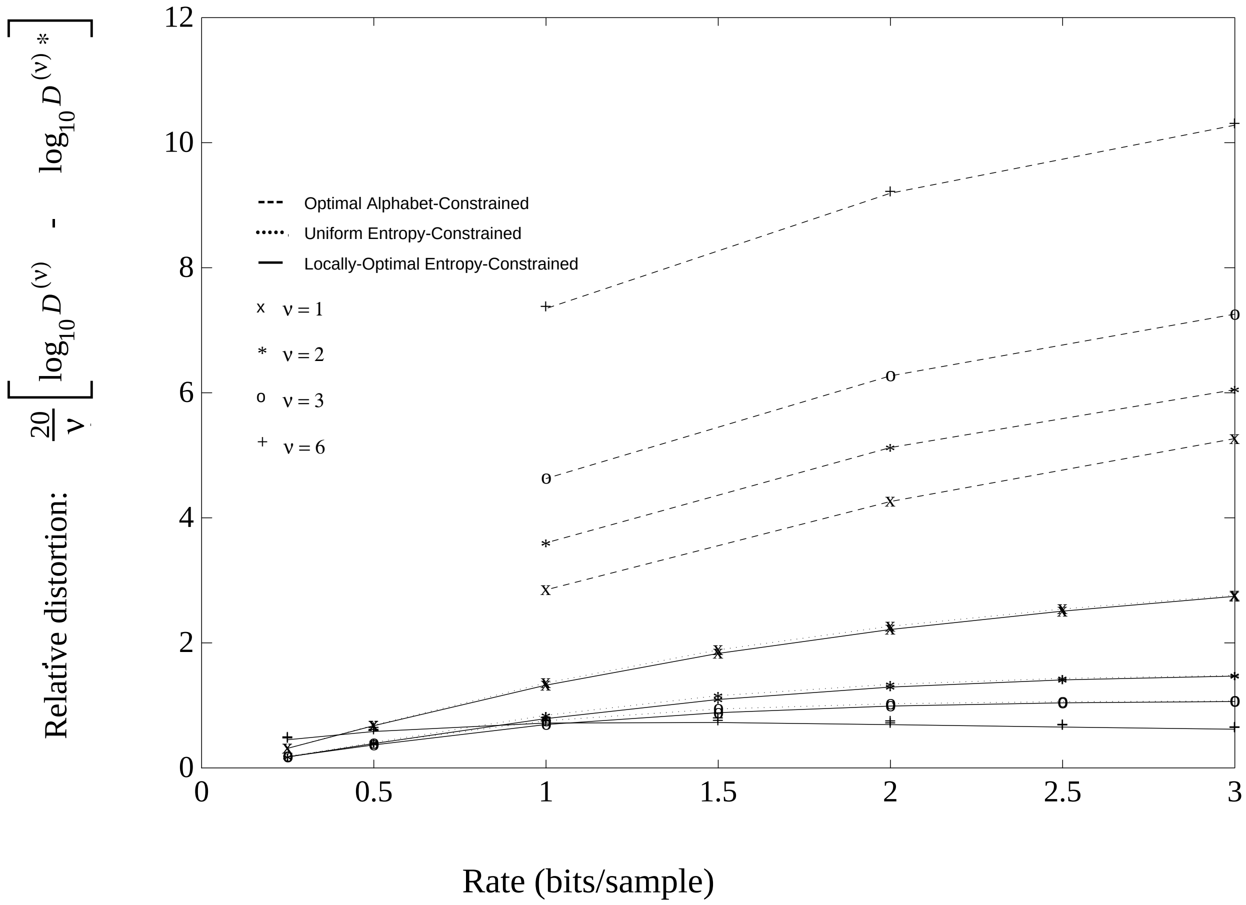

Figure 2.13: Performance (rate versus dB distortion in excess of distortion-rate bound) of locally-optimal alphabet-constrained (Lloyd-Max) quantization relative to uniform and locally-optimal entropy-constrained quantization, for a Laplacian source. Note that the distortion scale in this graph differs from that in Figure 2.12.

Figures 2.12—2.14 show the same performance curves, along with curves for locally-optimal alphabet-constrained quantization. Note that entropy-constrained quantization performs substantially better than alphabet-constrained quantization in most cases, particularly for large values of the error exponent ν.

It should be noted that the locally-optimal alphabet-constrained quantizers in the gamma MSE case shown in Figure 2.14 are *not* globally-optimal (this has been shown by Kabal [42]). However, the advantage of entropy-constrained quantization remains substantial even when the globally-optimal alphabet-constrained quantizers are considered in that special case.

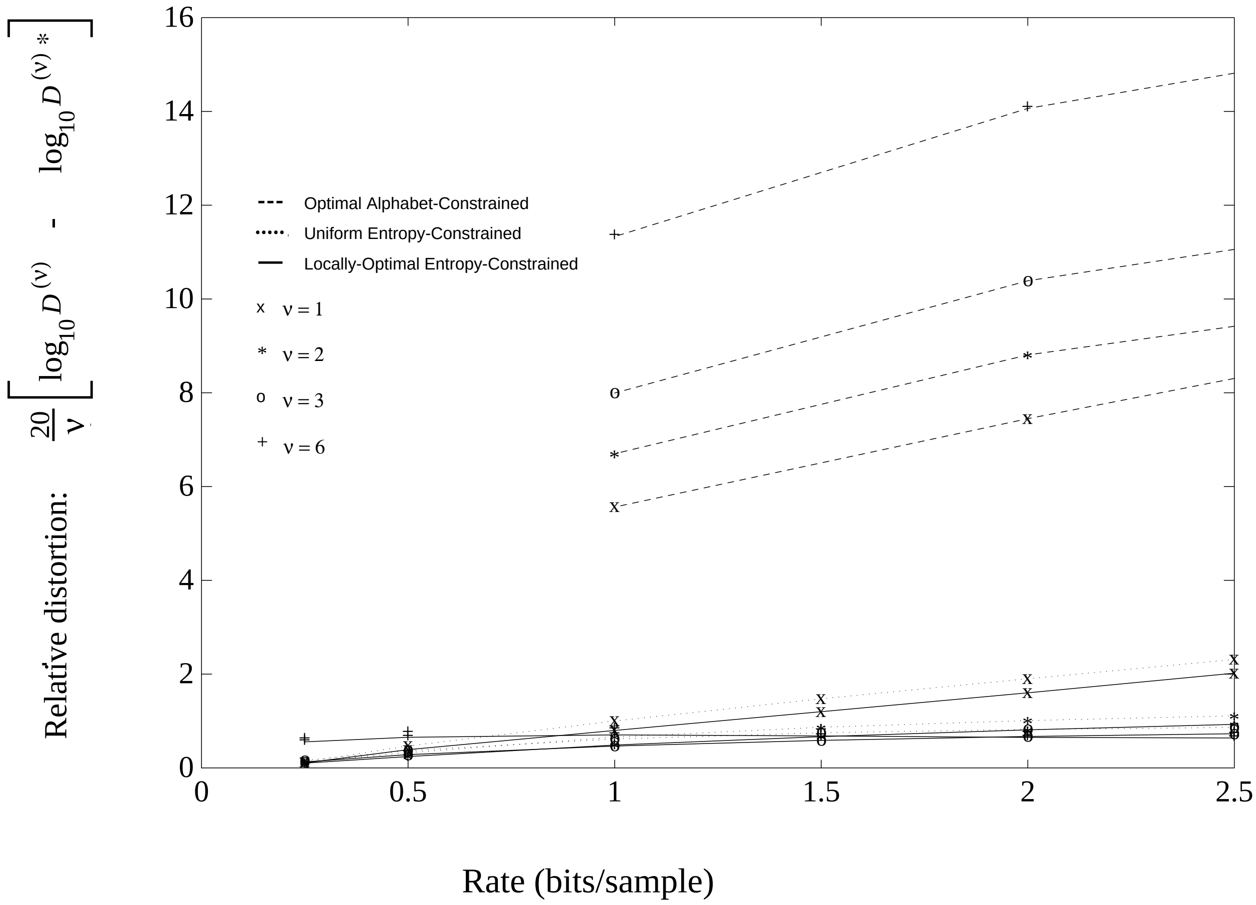

Figure 2.14: Performance (rate versus dB distortion in excess of distortion-rate bound) of locally-optimal alphabet-constrained (Lloyd-Max) quantization relative to uniform and locally-optimal entropy-constrained quantization, for a gamma source. Note that the distortion scale in this graph differs from those in Figures 2.12 and 2.13.

2.9 Summary and Conclusions

This chapter has considered the problem of designing entropy-constrained scalar quantizers for Gaussian, Laplacian, and gamma sources. Necessary conditions for optimality have been provided, and procedures for satisfying those conditions — i.e., for designing locally-optimal quantizers — have been described for both the alphabet- and entropy-constrained cases. Quantizers having uniformly spaced decision levels have been considered in detail, and it has been shown that such quantizers perform well at all rates and for a range of values of the distortion exponent ν. It has also been shown that in the case of the Laplacian source and mean-square error criterion, uniform quantizers are suboptimal at all rates when an entropy constraint is imposed.

This finding is contrary to previously reported results.

The advantage of entropy-constrained quantization over alphabet-constrained quantization for the sources of interest has been demonstrated over a range of rates and for several values of ν.

In an effort to gain insight into the nature of global optima, a random-search technique has been employed. This technique appears to be most effective at low rates, and the results of this technique provide support of the hypothesis that, in most cases, the locally-optimal quantizers designed in this study are, in fact, globally-optimal. It must be emphasized, however, that the results of the randomized search in no way constitute a proof.

Perhaps the most surprising finding of the present study is the apparent asymmetry of the optimal entropy-constrained quantizers in the low-rate region for a Gaussian source under the mean-absolute error criterion. Also surprising is the apparent breakpoint in the symmetry of decision levels at roughly 1.3 bits/sample for a Gaussian source and MSE criterion. Both of these phenomena require further study.