

Stealing Bits From a Quantized Source

Aaron S. Cohen, Stark C. Draper, *Member, IEEE*, Emin Martinian, *Member, IEEE*, and Gregory W. Wornell, *Fellow, IEEE*

Abstract—We consider “bit stealing” scenarios where the rate of a source code must be reduced without prior planning. We first investigate the efficiency of source requantization to reduce rate, which we term successive degradation. We focus on finite-alphabet sources with arbitrary distortion measures as well as the Gaussian-quadratic and high-resolution scenarios. We show an achievable rate–distortion tradeoff and prove that this is the best guaranteeable tradeoff for any good source code. This tradeoff is in general different from the rate-distortion tradeoff with successive refinement, where there is prior planning. But, we show that with quadratic distortion measures, for all sources with finite differential entropy and at least one finite moment, the gap is at most 1/2 bit or 3 dB in the high-resolution limit. In the Gaussian-quadratic case, the gap is at most 1/2 bit for all resolutions. We further consider bit stealing in the form of information embedding, whereby an embedder acts on a quantized source and produces an output at the same rate and in the original source codebook. We develop achievable distortion–rate tradeoffs. Two cases are considered, corresponding to whether or not the source decoder is informed of the embedding rate. In the Gaussian-quadratic case, we show the informed decoder need only augment the regular decoder with simple post-reconstruction distortion compensation in the form of linear scaling for the resulting system to be as efficient as bit stealing via successive degradation. Finally, we show that the penalty for uninformed versus informed decoders is at most 3 dB or 0.21-bit in the Gaussian-quadratic case and that their performance also lies within the 1/2-bit gap to that of successive refinement.

Index Terms—Coding with side information, data hiding, digital watermarking, information embedding, quantization, rate-distortion theory, successive degradation, successive refinement, transcoding.

I. INTRODUCTION

THERE are a variety of engineering scenarios where the rate of a source code must be decreased. This paper considers achievable, and guaranteeable, distortion rate tradeoffs when no

Manuscript received September 12, 2003; revised January 3, 2006. This work was supported in part by the National Science Foundation under Grant CCR-0073520, by Hewlett-Packard through the MIT/HP Alliance, and by a grant from Microsoft Research. The material in this paper was presented in part at the Data Compression Conference, Salt Lake, UT, April 2002, and at the IEEE International Symposium on Information Theory, Lausanne, Switzerland, June/July 2002.

A. S. Cohen is with Knight Capital Group, Purchase, NY 10577 USA (e-mail: cohen@alum.mit.edu).

S. C. Draper is with the Department of Electrical Engineering and Computer Science, University of California, Berkeley, CA 94720 USA (e-mail: sdraper@eecs.berkeley.edu).

E. Martinian is with Mitsubishi Electric Research Labs, Cambridge, MA 02139 USA (e-mail: emin@alum.mit.edu).

G. W. Wornell is with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: gww@mit.edu).

Communicated by M. Effros, Associate Editor for Source Coding.

Digital Object Identifier 10.1109/TIT.2006.876237

provisions have been made for the rate reduction. Because it is not planned for, we term this “bit stealing.”

To illustrate an application of bit stealing, consider transporting a compressed source through a congested multihop network. Suppose that the source code is of rate R_o and achieves distortion d_o . We discuss two strategies for alleviating the congestion that occurs when the source packet (which can be further compressed) and a second digital data packet (which cannot be compressed) both arrive at a common link that cannot support their combined rate.

The first is to split the link rate into two independent data streams. With this strategy, the source codeword is transcoded into a lower rate codebook of rate R , which increases the distortion to $d > d_o$, and the remaining (stolen) rate is used to transmit the second data packet. Transcoding is efficient if the pairs (R_o, d_o) and (R, d) both lie on the rate-distortion curve. If the source was originally encoded in a successively refinable manner (see, e.g., [1] and the references therein), efficient transcoding is sometimes possible by discarding least significant descriptions. In [1], Equitz and Cover give a necessary and sufficient Markov condition for such efficiency. We show, however, that near-efficiency in “successive degradation” can be possible even without such special codebook structure.

An alternative to splitting the rate into two data streams is to inject the data bits into the source bits via information embedding. This scenario differs from other investigations that jointly consider quantization and information embedding (see, e.g., [2]–[4]) because the source is quantized before embedding occurs. Therefore, the host signal at the information embedder is not a source vector, but rather a quantization index. If the embedded (stolen) rate is r , then the residual rate is $R = R_o - r$. As with successive degradation, the information embedding approach is efficient if (R_o, d_o) and (R, d) both lie on the rate-distortion curve.

Each strategy has its advantages. In rate-splitting, the data packet is easily decoded as it is transmitted independently of the reduced-rate source description. However, source decoding is more involved because the decoder must be informed of what lower rate codebook was used during transcoding. In embedding, on the other hand, message decoding is more involved because data bits are now intertwined with source bits. However, while the embedding operation changes which codeword is transmitted, the codebook can be kept the same. In certain applications, this can be an advantage since source decoders do not necessarily have to be informed that any bit stealing has taken place, for example, when we are trying to communicate the embedded information secretly.

From a broader perspective, one can view the original successive refinement problem as one of transcoding with *informed* encoders, i.e., encoders aware of the possibility that the source may




Fig. 1. Bit stealing via successive degradation: the rate R_0 of the code for source \mathbf{x} is reduced to R , incurring an increase in reconstruction distortion from d_0 to d . Source decoder I is the decoder for the original encoding, which produces reconstruction $\hat{\mathbf{x}}^0$ at distortion d_0 . Source decoder II is the decoder for the transcoded source, which produces reconstruction $\hat{\mathbf{x}}$ at distortion d .

subsequently be transcoded, and the rate of that transcoding. In contrast, the bit stealing problem is one of transcoding with *uninformed* encoders, i.e., encoders that either do not know that transcoding may take place, or do but do not know the residual rate. Within this taxonomy, there are two natural bit stealing sub-problems, corresponding to whether a source *decoder* is also informed or not about whether transcoding has taken place. We ultimately explore both cases in our development of the successive degradation and information embedding approaches.

An outline of the paper is as follows. Section III poses the successive degradation problem and characterizes its solution. The proofs are developed in Section IV, and Section V characterizes the results for the case of a binary source and Hamming distortion measure for the purpose of illustration. Section VI develops the corresponding results for a Gaussian source and quadratic distortion measure, and Section VII discusses continuous sources more generally in the high-resolution limit. Finally, Section VIII develops aspects of the behavior of information embedding strategies in the corresponding scenarios, for both informed and uninformed decoders, and Section IX contains some concluding comments. Many of the more technical derivations are deferred to the Appendix.

II. NOTATION

We use $I(\cdot;\cdot)$, $H(\cdot)$, $h(\cdot)$, and $D(\cdot\|\cdot)$ to denote mutual information, entropy, differential entropy, and relative entropy (divergence), respectively. The argument to $H(\cdot)$ and $h(\cdot)$ can be either a random variable or its distribution; we use both interchangeably. In addition, $H_B(\cdot)$ denotes the entropy of a Bernoulli source with the specified parameter. We use $R(\cdot)$ to denote generically the rate-distortion function for a source, and $d(\cdot)$ the corresponding distortion rate function. We further use T to denote the type (i.e., empirical distribution of the elements) of its vector subscript, and $\mathcal{T}(\cdot)$ to denote the type class of its argument, i.e., the set of vectors with empirical distribution given by the argument. Joint types are defined similarly. The superscript c applied to an event denotes its complement, and $|\cdot|$ applied to a set denotes its cardinality. Finally, \leftrightarrow is used to denote Markov chain relationships, and $E[\cdot]$ denotes expectation.

III. SUCCESSIVE DEGRADATION

Fig. 1 depicts the successive degradation scenario. The independent and identically distributed (i.i.d.) source \mathbf{x} , to which we restrict our attention to simplify the exposition, is encoded at rate R_0 giving source reconstruction $\hat{\mathbf{x}}^0$ at distortion d_0 . The transcoder either re-encodes the source codeword into a second codebook of residual rate $R \leq R_0$ giving source reproduction $\hat{\mathbf{x}}$ at distortion d or does not do any transcoding and, therefore,

does not change the rate or the distortion. In our problem model, all codebooks are known to their respective decoders.

Given the possibility of an informed source encoder, the successive degradation problem is a special case of branching communication systems investigated by Yamamoto in [5]. Yamamoto considered the joint encoding, but sequential decoding, of a pair of correlated sources. Setting the two sources equal gives the same rate-distortion region as generalized successive refinement [6]. Therefore, given control over the design of the source encoder, an optimal successive degradation approach is to encode into two codewords, per successive refinement. Decoder I receives both codewords, while Decoder II receives only the most significant.

In this paper, by restricting our attention to uninformed encoders, we preclude the possibility of having control over the design of the original source code. Instead we only assume knowledge of source statistics and that the source code achieves a point on the rate-distortion curve. Under these assumptions, we determine guaranteeable transcoder performance that is independent of detailed knowledge of source code design. From such analysis, we bound the gap between the best case performance that can be achieved when we can pick the best design for the original source quantizer (the informed encoder case), and the best performance that we can guarantee when we are not allowed to pick the detailed structure of the original source code (the uninformed encoder case). The former situation is equivalent to the successive refinement problem, while the latter is successive degradation.

Specifically, in Section IV, we bound the performance of successive degradation through a converse and an achievability result by introducing a dithered encoding rule. When used for the original source quantization, the dithering rule is a worst case choice, the most difficult to transcode, and leads to a lower bound on achievable rate-distortion tradeoffs. When used for the transcoding, the rule gives a robustly achievable rate-distortion tradeoff—one that can be achieved regardless of the detailed design of the original source code. The rate achieved equals the lower bound.

The successive degradation scenario is of interest, e.g., when the transcoder must be backward (or future) compatible and function correctly with encoders not designed to anticipate transcoding. Furthermore, since a refinable structure is an additional design requirement, we want to determine how important such structure is. In some prominent cases, the performance lost by not imposing such structure is not great—the natural structure of any source code can be nearly refinable.

An informal sense of the difference between successive refinement and successive degradation is illustrated in Fig. 2. A successive refinement code has a nested structure as indicated




Fig. 2. Voronoi regions are designed to nest efficiently in a successive refinement code (Fig. 2(a)), but such a packing does not come naturally to all codes (Fig. 2(b)). To degrade a code, a clustering of Voronoi regions that form a lower rate, higher distortion, description of the source must be found.

in Fig. 2(a). The base index specifies a rough quantization region (the bold circles in Fig. 2(a)), and refinement regions (the small circles in Fig. 2(a)) are designed conditionally, given a base index and the source vector. This design is most efficient when fine regions nest in rough ones.¹

On the other hand, not all high-rate codes have a nested structure. In general, to lower the description rate of a source, a clustering of fine quantization regions must be chosen. The clusters describe the source at lower fidelity and can be enumerated at a lower rate. This is depicted in Fig. 2(b). Degradation performance will therefore be dependent on the design of the original source code. By determining guaranteeable rate–distortion tradeoffs that are independent of source code design, we determine the degree to which a nested code structure is important.

A. Rate–Distortion Functions

We begin by defining an information rate–distortion function, which we subsequently show to be operational in a rather natural sense.

Definition 1: Let $p_X(x)$ be the distribution of the i.i.d. source and let $d(\cdot, \cdot)$ be the corresponding per-letter distortion measure. Let $p_{\hat{x}^\circ | x}(\hat{x}^\circ | x)$ be a conditional distribution that uniquely achieves² a point on the rate–distortion function with

$$d_o = \sum_{x, \hat{x}^\circ} p_{\hat{x}^\circ | x}(\hat{x}^\circ | x) p_X(x) d(x, \hat{x}^\circ). \quad (1)$$

¹The nesting character of successive refinement codes is developed formally by Rimoldi [6].

²Achievability means that (1) implies $I(x; \hat{x}^\circ) = R_{\text{SC}}(d_o) \triangleq R_o$, where $R_{\text{SC}}(\cdot)$ is the rate–distortion function for regular source coding. Uniqueness means that no other conditional distribution $q_{\hat{x}^\circ | x}(\hat{x}^\circ | x)$ can achieve the same point on the rate–distortion function. We assume the sufficient conditions for such uniqueness (see, e.g., [7, Lemma 7]) are met. The dependence of the successive degradation rate–distortion function on $p_{\hat{x}^\circ | x}(\hat{x}^\circ | x)$, or d_o , is suppressed for notational convenience, just as is the dependence of the standard rate–distortion function on the source statistics.

Then, the *information successive degradation rate–distortion function* $R_{\text{SD}}(d)$ is defined for distortions $d \geq d_o$ as

$$R_{\text{SD}}(d) = \min\{\inf I(\hat{x}; \hat{x}^\circ), R_o\} \quad (2)$$

where the infimum is over all conditional distributions $p_{\hat{x} | \hat{x}^\circ}(\hat{x} | \hat{x}^\circ)$ such that the Markov constraint $x \leftrightarrow \hat{x}^\circ \leftrightarrow \hat{x}$ is satisfied, $E[d(x, \hat{x})] \leq d$, and $R_o \triangleq R_{\text{SC}}(d_o) = I(x; \hat{x}^\circ)$.

The corresponding information successive degradation distortion–rate function is defined for rates $R \leq R_o$ as

$$d_{\text{SD}}(R) = \inf\{d \geq 0 : R_{\text{SD}}(d) \leq R\}. \quad (3)$$

Some initial insights into successive degradation rate–distortion performance are obtained by relating it to three well known rate–distortion functions.

We first compare it to the regular rate–distortion function and then to the successive refinement rate–distortion function. In general, it does not coincide with either. Following that we build intuition by showing how it can be expressed in terms of the indirect (or “noisy”) rate–distortion function [8], [9]. For reference, typical successive degradation, successive refinement, indirect coding, and the regular rate–distortion functions are depicted in Fig. 3.

The successive degradation rate–distortion function generally is different from the regular rate–distortion function $R_{\text{SC}}(d)$. To see this, note that because of the Markov constraint in its definition, we can rewrite the mutual information being minimized in (2) as

$$I(\hat{x}; \hat{x}^\circ) = I(\hat{x}; x, \hat{x}^\circ) = I(x; \hat{x}) + I(\hat{x}^\circ; \hat{x} | x). \quad (4)$$

Then since the first term on the right-hand side of (4) corresponds to $R_{\text{SC}}(d)$, the successive degradation rate–distortion function equals $R_{\text{SC}}(d)$ when the second term is zero, i.e., when




Fig. 3. Typical rate–distortion functions corresponding to an original operating point (R_o, d_o) . The successive degradation rate–distortion function is indicated by the solid curve and is defined for distortions $d \geq d_o$. The indirect rate–distortion function for $p_{y|x} = p_{\hat{x}^\circ|x}$ is indicated by the dotted curve, is also defined for distortions $d \geq d_o$, and overlaps the successive degradation curve for $d \geq d_*$. The successive refinement rate–distortion function is indicated by the dash-dotted curve and is defined for distortions $d \leq d_o$. Finally, the regular rate–distortion function is indicated by the lower dashed curve.

the additional Markov constraint $\hat{x}^\circ \leftrightarrow x \leftrightarrow \hat{x}$ is satisfied. Under fairly general conditions,³ this Markov constraint and that in Definition 1 can only hold simultaneously if \hat{x} is independent of x and \hat{x}° . From this we conclude that, in most cases, $R_{SD}(d) > R_{SC}(d)$ unless $R_{SC}(d) = 0$ or $d = 0$.

Similarly, $R_{SD}(d)$ also generally differs from the successive refinement rate–distortion function $R_{SR}(d)$. In particular, [6] shows that a source description \hat{x}° at distortion d_o using rate $R_o = R_{SC}(d_o)$ is *refinable* to a distortion $d \leq d_o$ if and only if the total rate of the refined description is at least

$$R_{SR}(d) = \inf I(x; \hat{x}, \hat{x}^\circ) \quad (5)$$

where the infimum is taken over all conditional distributions $p_{\hat{x}|\hat{x}^\circ, x}(\hat{x}|\hat{x}^\circ, x)$ such that $E[d(x, \hat{x})] \leq d$. Expanding the mutual information in (5) as

$$I(x; \hat{x}, \hat{x}^\circ) = I(x; \hat{x}) + I(x; \hat{x}^\circ | \hat{x}) \quad (6)$$

we see, as established in [1], that $R_{SR}(d)$ differs from $R_{SC}(d)$ unless a still different Markov constraint $x \leftrightarrow \hat{x} \leftrightarrow \hat{x}^\circ$ is satisfied.

We gain additional insight into the structure of the successive degradation rate–distortion function by expressing it in terms of the noisy or “indirect coding” rate–distortion function [8], [9], which we denote as $R_{IC}(d)$. The indirect coding rate–distortion function is the minimum rate needed to describe a memoryless source x , observed by the encoder through the discrete memoryless channel $p_{y|x}(y|x)$. Depending on the channel law, certain nonzero distortions d are unachievable, regardless of encoding rate. In particular, [8] and [9] show that a source description \hat{x} at distortion d can be achieved if and only if the total rate of the description is at least

$$R_{IC}(d) = \inf I(y; \hat{x}). \quad (7)$$

The infimum is taken over all conditional distributions $p_{\hat{x}|y}(\hat{x}|y)$ such that $E[d(x, \hat{x})] \leq d$ where $d \geq d_o$. Note

³A sufficient condition is that $p_{x, \hat{x}^\circ}(x, \hat{x}^\circ) > 0$ for all x, \hat{x}° . This condition is satisfied, for example, in the binary–Hamming and Gaussian–quadratic cases unless $d_o = 0$.

that this choice of conditional distribution means that the Markov constraint $\hat{x} \leftrightarrow y \leftrightarrow x$ is satisfied. The value of d_o is determined by the minimum distortion estimate of x given full knowledge y . A basic result for a white Gaussian source observed in additive white Gaussian noise is that minimum mean-squared estimation followed by rate–distortion encoding of the estimate yields the optimum mean-squared distortion performance. In this case, d_o equals the variance of the error in the estimate $E[x|y]$.

To express $R_{SD}(d)$ in terms of $R_{IC}(d)$, it suffices to compare (2) with (7). In particular, if we let

$$p_{y|x}(y|x) = p_{\hat{x}^\circ|x}(y|x) \quad (8)$$

then the successive degradation rate–distortion function (2) can be re-expressed in terms of the indirect rate–distortion function (7) as

$$R_{SD}(d) = \begin{cases} R_{IC}(d), & \text{if } d \geq d_* \\ R_{IC}(d_*), & \text{if } d_o \leq d \leq d_* \end{cases} \quad (9)$$

where d_* is the distortion such that $R_{IC}(d_*) = R_{SC}(d_o)$.

To further interpret these relationships, we need to formally define the successive degradation problem for which we will subsequently show the information rate–distortion function (2) to be an operational one.

B. The Successive Degradation Game

To determine the best guaranteeable rate–distortion tradeoff for the successive degradation problem, consider the following zero-sum game. The first player picks a transcoder anticipating the worst possible encoder. The second player picks an encoder designed to be as difficult as possible to transcode while meeting given rate and distortion constraints. In Section III-D we show that the information successive degradation rate–distortion function $R_{SD}(d)$ defined in Section III-A is the operational rate–distortion function for this game.

Specifically, any of a broad class of rate- R_o source codes for x that achieve average distortion d_o with R_o arbitrarily close to $R_{SC}(d_o)$ can be *degraded* to distortion arbitrarily close to $d > d_o$ if and only if the transcoder rate satisfies $R > R_{SD}(d)$. To develop this result, we first introduce the formal problem setting.

An instance of the successive degradation game consists of the tuple

$$\{\mathcal{X}, p(x), d(\cdot, \cdot), R_o, R\}. \quad (10)$$

The source (and reconstruction) alphabet \mathcal{X} is finite unless otherwise indicated, $p(x), p(\hat{x}^\circ|x)$ and $d: \mathcal{X} \times \mathcal{X} \mapsto \mathbb{R}^+$ are as given in Definition 1, and R_o and R are rate constraints on the encoder and transcoder that are the inputs to the game. Note, too, that by the uniqueness in Definition 1, R_o specifies $p(\hat{x}^\circ|x)$ in the game.

An encoder $\Phi_n^\circ = (C_n^\circ, W^\circ(\cdot|\cdot))$ consists of a codebook C_n° of cardinality $|C_n^\circ| = 2^{nR_o}$ and a potentially randomized encoding rule $W^\circ(i|x)$. The latter denotes the probability that the n -length input vector x is mapped to codebook index $i \in \{1, 2, \dots, |C_n^\circ|\}$. The reconstruction corresponding to this

encoding is simply the codeword corresponding to the chosen index, i.e., $\hat{\mathbf{x}}^\circ = \mathcal{C}_n^\circ(i)$.

A transcoder $\Phi_n = (\mathcal{C}_n, W(\cdot|\cdot, \cdot))$ is a codebook \mathcal{C}_n of cardinality $|\mathcal{C}_n| = 2^{nR}$ and a potentially randomized rule $W(j|i, \Phi_n^\circ)$. The latter denotes the probability that a given index i produced by the (known) original encoder Φ_n° is mapped to index $j \in \{1, 2, \dots, |\mathcal{C}_n|\}$. The reconstruction corresponding to this transcoding is simply the codeword corresponding to the chosen index, i.e., $\hat{\mathbf{x}} = \mathcal{C}_n(j)$. Finally, we also allow the transcoder to output the original encoding followed by a special flag indicating that no transcoding has been done.

Notationally, we use $d_n(\cdot, \cdot)$ for the average distortion between length- n sequences, i.e., for any a^n and b^n

$$d_n(a^n, b^n) = \frac{1}{n} \sum_{i=1}^n d(a_i, b_i)$$

where $d(\cdot, \cdot)$ is the per-letter distortion measure, which is bounded unless otherwise indicated. Hence, the distortion in the reconstruction $\hat{\mathbf{x}}^\circ$ associated with encoding is $d_n(\hat{\mathbf{x}}^\circ, \mathbf{x})$, and likewise for the reconstruction $\hat{\mathbf{x}}$ associated with transcoding it is $d_n(\hat{\mathbf{x}}, \mathbf{x})$. We omit the subscript n when there is no risk of confusion.

We now describe the payoffs in the successive degradation game for a fixed n . The encoder and transcoder players choose randomized strategies P_{Φ° and P_Φ for generating encoder/decoder and transcoder/decoder pairs, respectively. Given these strategies and a reference distortion level d , the payoff to the transcoder is given by

$$\pi(P_{\Phi^\circ}, P_\Phi, d) = \Pr\{d(\mathbf{x}, \hat{\mathbf{x}}) \leq d\}. \quad (11)$$

The payoff to the encoder is simply $-\pi(P_{\Phi^\circ}, P_\Phi, d)$. The probability in (11) is evaluated as follows. First, the source \mathbf{x} is drawn according to $p_{\mathbf{x}}(\mathbf{x}) = \prod_{i=1}^n p(x_i)$. Then, an encoder Φ° is chosen according to P_{Φ° and used to generate $\hat{\mathbf{x}}^\circ$ from \mathbf{x} . Finally, a transcoder Φ is chosen according to P_Φ and used to generate $\hat{\mathbf{x}}$ from $\hat{\mathbf{x}}^\circ$ and Φ° . The random choices in these steps are mutually independent, except for when the transcoder does no transcoding, i.e., when it forwards the original source encoding unchanged. In this particular case, the same codebook is used by both.

We consider the asymptotic value of $\pi(P_{\Phi_n^\circ}, P_{\Phi_n}, d)$ for sequences of strategies $\{P_{\Phi_n^\circ}\}$ and $\{P_{\Phi_n}\}$ where the sequence of encoders must achieve the point (R_\circ, d_\circ) in a sense to be defined later and the sequence of transcoders must use rate R . We demonstrate a saddle point for this asymptotic game so that the order of play does not effect the equilibrium payoffs. Note, however, that by restricting our attention to the case in which the random encoder and transcoder choices are made independently, our results do not address scenarios in which the encoder codebook \mathcal{C}° is a function of the transcoder codebook \mathcal{C} , nor *vice versa*. An exception is made for the single case when transcoding is not performed, i.e., when the original encoding is left unmodified.

Before proceeding to develop a coding theorem establishing that the operational rate–distortion function for this game is given by (2), we discuss some initial aspects of players' behavior that are revealed by the structure of this function.

C. Preliminary Interpretation of the Solution

To begin, the expression (9) for $R_{\text{SD}}(d)$ emphasizes that the successive degradation rate–distortion function consists of two distinct regions. As we will develop, each region corresponds to a particular strategy. In one region, the transcoder does nothing, it simply leaves the original encoding unperturbed. In the other, the transcoder treats the original source reconstruction as a new source and requantizes it using an independently generated codebook. The boundary between the two regions is demarcated by the distortion $d = d_*$, as shown in Fig. 3.

The bi-modal character of the successive degradation rate–distortion function is ultimately caused by the quantized nature of the source space seen by the transcoder. In particular, suppose the distortion were increased only slightly, to $d_\circ + \Delta$, where $0 < \Delta \ll d_\circ$. Because the source reconstructions corresponding to the codeword indices are spaced roughly $2d_\circ$ apart, in order to requantize to increase the distortion by this small Δ , we would have to cover the original source code with a finer code. Many quantization regions of the new code would fall within a single old region. This would require a code of higher rate than the original code. Indeed, the required rate is specified by indirect coding tradeoff. But, by the data processing inequality, this new code contains no more information about the source than did the original, lower rate, source code. Thus, to achieve a d in the range $d_\circ \leq d \leq d_*$, it suffices to leave the original encoding alone, and simply add an appropriate amount of source-independent noise to the source reconstruction.

If, instead, the distortion were increased by a more substantial amount—to $d \geq d_*$ —we could group codewords as in Fig. 2(b). In particular, the analysis of our game will reveal that in this distortion region $R_{\text{SD}}(d)$ equals $R_{\text{IC}}(d)$ for a channel law $p_{Y|X}$ of the form (8). That there should be a discontinuity is also apparent from the analysis of a rate perturbation. Specifically, given any rate decrease to $R < R_\circ$, the ratio of new codewords to old is $2^{-n(R_\circ - R)}$, which is vanishingly small in the block length. Thus, we no longer have the option of leaving the original encoding unperturbed. Hence, unless the code has a special nested character, the nonnested packing (depicted in Fig. 2(b)) results in a finite jump in distortion.

Readers familiar with the many classical rate–distortion results may find the discontinuity disturbing nevertheless, since, as Fig. 3 reveals, it results in a rate–distortion function that is not convex. Indeed, one might be tempted to conclude that it is the convex envelope of $R_{\text{SD}}(d)$ that is relevant. In fact, such intuition would be incorrect, as we now discuss.

The key point is that convexification of $R_{\text{SD}}(d)$ corresponds to the adoption of time-sharing techniques by at least one of the players in the game. However, the fact that the coding theorem we develop in the Section III-D determines (9) to be the operational rate–distortion function establishes that it is in neither player's interest to use a time-sharing strategy. While this will become clear in the proofs of both the direct and converse parts of our coding theorem, it is worth emphasizing a couple of points in advance.

First, remember that the encoder chooses an encoding that is the most difficult to transcode, while meeting the specified rate and distortion requirements, and that the transcoder chooses a transcoding function guaranteed to work for any choice the encoder makes. Were the encoder to use a time-sharing-compatible encoding, i.e., one consisting of a sequence of encoded

source blocks, the transcoder would then certainly choose to use time sharing. It would requantize some of the codewords and leave the rest unperturbed. The resulting distortion, averaged across the codewords, would be decreased, lying somewhere on the convex hull of $R_{SD}(d)$. In terms of the game, this would increase the payoff to the transcoder and decrease the payoff to the encoder. However, such a transcoding strategy would not be robust. Were the encoder to switch to some non-time-sharing-compatible encoding, the transcoding scheme would break, resulting in an increase in distortion—a lowered payoff to the transcoder and an increased payoff to the encoder.

Fundamentally, the reason time sharing cannot be used to improve the successive degradation tradeoff is that via our game-theoretic formulation we are effectively restricting our attention to encoders unaware that transcoding may occur. Therefore, the encoder cannot be relied upon to produce a time-sharing-friendly encoding consisting of a sequence of encoded source blocks. If the encoder were aware of the transcoding then, as discussed above, distortion could be even further decreased by requiring the encoder to produce a successively refinable encoding.

Phrased differently, recall that the game-theoretic formulation of the problem was motivated by the goal of examining the impact of various design constraints that may be placed on the encoder (beyond the requirement to produce an encoding that meets the rate and distortion targets). For example, we ultimately use our results to bound the loss that the robust transcoder incurs relative to one tuned to an encoding with a successively refinable structure. We show that in some important cases the loss is not large. From this, one may infer that the gain from using a refinable structure may not always reward the extra design cost. From this perspective, our results establish that a similar message holds for encodings employing structure in the form of time sharing. In particular, it follows that the loss of a robust transcoder relative to one tuned to an encoding constrained to be time-sharing-compatible is bounded by the loss relative to a successively refinable structure.

We now proceed to formally establish the optimal actions of the two players, and the resulting rate–distortion tradeoff that is achieved.

D. A Coding Theorem

Intuitively, we expect that just as the familiar source coding rate–distortion function depends on the source distribution, the successive degradation rate–distortion function should depend on properties of the original source code. We focus on efficient encoders with performance close to the rate–distortion bound.

To formalize this notion, we define a class of admissible encoders. In the sequel, we use standard definitions of (and notation for) empirical distributions or types [10, Sec. 1.2], [11, Ch. 12]. For example, if $T_{\hat{\mathbf{x}}^\circ, \mathbf{x}}$ denotes the joint type of $(\hat{\mathbf{x}}^\circ, \mathbf{x})$ then $T_{\hat{\mathbf{x}}^\circ, \mathbf{x}}(\hat{x}^\circ, x)$ is the relative frequency of occurrences of the sample-pair (\hat{x}°, x) in the sequence-pair $(\hat{\mathbf{x}}^\circ, \mathbf{x})$.

Definition 2: Let $T_{\hat{\mathbf{x}}^\circ, \mathbf{x}}$ be the joint type of encoder output $\hat{\mathbf{x}}^\circ$ and source \mathbf{x} . A sequence of encoders $\{\Phi_n^\circ\}$ is said to be *admissible* if $D(T_{\hat{\mathbf{x}}^\circ, \mathbf{x}} \| p_{\hat{\mathbf{x}}^\circ, \mathbf{x}}) \rightarrow 0$ in probability as $n \rightarrow \infty$, where $D(\cdot \| \cdot)$ denotes relative entropy and

$$p_{\hat{\mathbf{x}}^\circ, \mathbf{x}}(\hat{x}^\circ, x) = p_{\mathbf{x}}(x) p_{\hat{\mathbf{x}}^\circ | \mathbf{x}}(\hat{x}^\circ | x).$$

An admissible sequence of rate- R_\circ encoders achieves the point (R_\circ, d_\circ) on the rate–distortion function. This is because the probability that $\hat{\mathbf{x}}^\circ$ and \mathbf{x} are asymptotically strongly typical according to $p_{\hat{\mathbf{x}}^\circ, \mathbf{x}}$ approaches one, whence

$$d(\mathbf{x}, \hat{\mathbf{x}}^\circ) \rightarrow E[d(x, \hat{x}^\circ)] = d_\circ$$

in probability. To verify the strong typicality claim, it suffices to note that via [11, Lemma 12.6.1, p. 300] we have, for all (\hat{x}°, x)

$$|T_{\hat{\mathbf{x}}^\circ, \mathbf{x}}(\hat{x}^\circ, x) - p_{\hat{\mathbf{x}}^\circ, \mathbf{x}}(\hat{x}^\circ, x)| \leq \sqrt{2 \ln 2 \cdot D(T_{\hat{\mathbf{x}}^\circ, \mathbf{x}} \| p_{\hat{\mathbf{x}}^\circ, \mathbf{x}})} \quad (12)$$

and, moreover, that for all (\hat{x}°, x) with $p_{\hat{\mathbf{x}}^\circ, \mathbf{x}}(\hat{x}^\circ, x) = 0$ we have $T_{\hat{\mathbf{x}}^\circ, \mathbf{x}}(\hat{x}^\circ, x) = 0$ (otherwise, $D(T_{\hat{\mathbf{x}}^\circ, \mathbf{x}} \| p_{\hat{\mathbf{x}}^\circ, \mathbf{x}})$ would be infinite).

Note that the set of admissible encoders is reasonably broad. For example, it includes the familiar strong typicality encoders.⁴ Furthermore, extensions to source coding (see [12] and particularly [13, Sec. 2.6]) of analogous results in channel coding [14] tell us that the probability that source sequences are jointly nontypical with their codewords decays exponentially in block length. In this statement, joint typicality is with respect to the joint distribution $p_{\hat{\mathbf{x}}^\circ, \mathbf{x}}$ induced by the source probability mass function and the (assumed unique) rate–distortion achieving channel.

Our main theorem is as follows.

Theorem 1: For the successive degradation game

$$\begin{aligned} & \inf_{\{P_{\Phi_n^\circ}\}} \sup_{\{P_{\Phi_n}\}} \lim_{n \rightarrow \infty} E[\pi(P_{\Phi_n^\circ}, P_{\Phi_n}, d)] \\ &= \sup_{\{P_{\Phi_n}\}} \inf_{\{P_{\Phi_n^\circ}\}} \lim_{n \rightarrow \infty} E[\pi(P_{\Phi_n^\circ}, P_{\Phi_n}, d)] \\ &= \begin{cases} 1, & \text{if } R > R_{SD}(d) \\ 0, & \text{if } R \leq R_{SD}(d) \end{cases} \end{aligned} \quad (13)$$

where the infima are taken over admissible sequences of rate- R_\circ encoders and the suprema are taken over sequences of rate- R transcoders.

Theorem 1 implies that the information rate–distortion function of (2) gives the best possible worst case successive degradation tradeoff. If the transcoder’s rate is below (2), then there exists at least one encoder that causes the transcoder to fail. If a rate higher than (2) is used, then there exists a transcoder that almost always wins.

⁴To see that strong typicality implies small divergence, consider a vector \mathbf{x} that is strongly typical with respect to the distribution $p_{\mathbf{x}}(x)$. That is, $|T_{\mathbf{x}}(x) - p_{\mathbf{x}}(x)| < \epsilon$ for all x and some ϵ . Without loss of generality, $p_{\mathbf{x}}(x) \geq p_{\min}$ for all x . Thus,

$$\begin{aligned} D(T_{\mathbf{x}} \| p_{\mathbf{x}}) &= \sum_x T_{\mathbf{x}}(x) \log \frac{T_{\mathbf{x}}(x)}{p_{\mathbf{x}}(x)} \\ &\leq \sum_x T_{\mathbf{x}}(x) \frac{T_{\mathbf{x}}(x) - p_{\mathbf{x}}(x)}{p_{\mathbf{x}}(x)} \\ &\leq \frac{\epsilon}{p_{\min}} \end{aligned}$$

where the first inequality follows since $\log x \leq x - 1$ for $x > 0$ and the second inequality follows by strong typicality and the assumption $p_{\mathbf{x}}(x) \geq p_{\min}$. Through the choice of ϵ , the divergence can be made as small as desired.




Fig. 4. The role of dithered quantization in the successive degradation rate–distortion coding theorem. Dithering is used at the encoder and transcoder in the converse and achievability arguments, respectively. (a) Achievability. (b) Converse.

IV. PROOF OF SUCCESSIVE DEGRADATION RATE–DISTORTION THEOREM

Intuitively, one can view the encoder output $\hat{\mathbf{x}}^\circ$ as a noisy source observation. If the quantization noise were i.i.d. then results on encoding noisy sources (see, e.g., [9] and the references therein) would give the transcoding rate–distortion region. However, the joint distribution of $(\mathbf{x}, \hat{\mathbf{x}}^\circ)$ for good vector quantizers is generally not i.i.d.

Accordingly, we prove Theorem 1 by using a special form of dithered quantization. The joint input/output distribution of quantizers in this class is essentially indistinguishable from an i.i.d. relationship, and yet as we will see the performance of quantizers in this class approaches the rate–distortion bound. The forward part of the theorem is proved using this dithered quantization at the transcoder. This induces an i.i.d.-like joint distribution on the transcoder inputs and outputs, allowing us to use the Markov lemma [15] to guarantee that as long as the source \mathbf{x} and encoder output $\hat{\mathbf{x}}^\circ$ are strongly typical, then the transcoder output $\hat{\mathbf{x}}$ and the source \mathbf{x} will be also. The converse is shown in a complementary manner. In this case, our dithered quantization is used at the source encoder. No transcoder can do better in this situation than the rate–distortion results for quantizing noisy sources. The position of dithered quantization in the achievability and converse halves of the proof is indicated in Fig. 4(a) and (b), respectively.

A. Dithered Quantization

The design of a dithered quantizer is governed by an input distribution $p_u(u)$, a quantization distribution $p_{v|u}(v|u)$, a quantization rate R^δ , and a parameter $\delta > 0$ that can be arbitrarily small. The resulting quantizer $\Phi_n^\delta = (\mathcal{C}_n^\delta, W_n^\delta(\cdot|\cdot))$ consists of a codebook \mathcal{C}_n^δ of cardinality $|\mathcal{C}_n^\delta| = 2^{nR^\delta}$ and a quantization rule $W_n^\delta(\cdot|\cdot)$ that avoids joint typicality encoding. We let \mathbf{u} and \mathbf{v} denote the quantizer input and output, respectively.

Dithered Quantization Codebook \mathcal{C}_n^δ Construction:

- 1) Generate 2^{nR^δ} sequences of length n in an i.i.d. manner according to

$$p_v(v) = \sum_u p_u(u) p_{v|u}(v|u).$$

These are the codewords in the codebook.

- 2) Label these codewords $\mathbf{v}(1), \mathbf{v}(2), \dots, \mathbf{v}(2^{nR^\delta})$.

Dithered Quantization Rule $W_n^\delta(\cdot|\cdot)$:

- 1) Generate a noisy observation \mathbf{w} from the input \mathbf{u} according to the conditional distribution

$$p_{\mathbf{w}|\mathbf{u}}(\mathbf{w}|\mathbf{u}) = \prod_{i=1}^n p_{v|u}(w_i|u_i). \quad (14)$$

Denote by $T_{\mathbf{w},\mathbf{u}}$ the joint type of (\mathbf{w}, \mathbf{u}) . Denote by $p_{v,u}$ the joint probability distribution $p_u(u)p_{v|u}(v|u)$.

- 2) If $D(T_{\mathbf{w},\mathbf{u}}||p_{v,u}) > \delta$ the quantization fails, so choose a codeword at random from the codebook.
- 3) Otherwise, list all sequences \mathbf{v} in the codebook such that $T_{\mathbf{v},\mathbf{u}} = T_{\mathbf{w},\mathbf{u}}$. If no such sequences exist, the quantization fails, so choose a codeword at random from the codebook.
- 4) Otherwise, choose a codeword at random from this list. In this case the quantization succeeds.

When they succeed, dithered quantizers have the property that their outputs “look” like the output of a memoryless noisy channel with a channel law given by the quantization distribution. This property is formalized by the following lemma, which establishes that any theorem regarding an appropriately chosen memoryless observation of a source vector (not necessarily a memoryless source vector) also holds for the dither-quantized representation of that source. Furthermore, the lemma continues to hold even when there is encoder side information. A proof is provided in Appendix A.

Lemma 1: Consider any binary-valued test $\theta[\cdot] \in \{0, 1\}$, and a random map $g(\cdot)$ with domain \mathcal{X} and a finite but arbitrary range. Furthermore, let \mathbf{u} and \mathbf{v} be the input and output, respectively, of a dithered quantizer, and let \mathbf{w} be generated from \mathbf{u} in an i.i.d. manner according to the distribution (14). Then

$$|E\{\theta[\mathbf{v}, \mathbf{u}, g(\mathbf{v})]\} - E\{\theta[\mathbf{w}, \mathbf{u}, g(\mathbf{w})]\}| < 4\Pr\{\mathcal{E}\}$$

provided that $\Pr\{\mathcal{E}\} < 1/2$ and that the map $g(\cdot)$ is independent of \mathbf{u}, \mathbf{v} , and \mathbf{w} , where \mathcal{E} denotes the event that the dithered quantization fails.

B. Design and Properties of Dithered Encoders and Transcoders

We now develop specific dithered encoders and dithered transcoders, and some of their key properties that will be useful in the sequel.

In the design of our dithered encoder, we choose the input distribution to be $p_x(x)$, that of the i.i.d. source. Furthermore, we choose the quantization rate to be R_o and the quantization distribution to be the conditional distribution $p_{\hat{x}^\circ|x}(\hat{x}^\circ|x)$ in Definition 1 that achieves the distortion rate function at the target rate, i.e., $d_{\text{SC}}(R_o) = d_o$ for the distortion measure $d(\cdot, \cdot)$ of interest.⁵

Propositions 1 and 2, whose proofs are provided in Appendices B and C, respectively, establish that dithered encoders designed in this way are good, i.e., they can perform within ϵ of the rate–distortion function, and are admissible in the sense of Definition 2.

⁵The distortion rate function is defined in terms of the rate–distortion function in the usual way: $d_{\text{SC}}(R_o) = \inf\{d_o \geq 0 : R_{\text{SC}}(d_o) \leq R_o\}$.

Proposition 1: The distortion for the dithered encoder satisfies

$$\Pr[d(\mathbf{x}, \hat{\mathbf{x}}^\circ) > E[d(x, \hat{x}^\circ)] + d_{\max} \sqrt{2\delta \ln 2}] \leq \Pr[\mathcal{E}] \quad (15)$$

where δ is the encoder parameter and

$$d_{\max} = \max_{a,b \in \mathcal{X}} d(a,b).$$

Proposition 2: There exists a $\delta_{\max} > 0$ such that for every $R_o > I(\hat{\mathbf{x}}^\circ; \mathbf{x})$ and $\delta \in (0, \delta_{\max})$, $\Pr[\mathcal{E}] \rightarrow 0$ exponentially as $n \rightarrow \infty$ when the input is an i.i.d. sequence \mathbf{x} . Furthermore, δ_{\max} depends only on R_o and $p_{\hat{\mathbf{x}}^\circ, \mathbf{x}}(\hat{\mathbf{x}}^\circ, \mathbf{x})$.

We turn next to our dithered transcoders. For their design, we choose the input distribution to be

$$p_{\hat{\mathbf{x}}^\circ}(\hat{\mathbf{x}}^\circ) = \sum_x p_{\hat{\mathbf{x}}^\circ | \mathbf{x}}(\hat{\mathbf{x}}^\circ | x) p_{\mathbf{x}}(x) \quad (16)$$

where $p_{\mathbf{x}}(x)$ is the original source distribution and $p_{\hat{\mathbf{x}}^\circ | \mathbf{x}}(\hat{\mathbf{x}}^\circ | x)$ is again the conditional distribution associated with the original source code in Definition 1 corresponding to the distortion rate function operating point $d_{\text{SC}}(R_o) = d_o$. Furthermore, we choose the quantization rate to be R , and the quantization distribution to be the optimizing $p_{\hat{\mathbf{x}} | \hat{\mathbf{x}}^\circ}$ in (2). The following proposition, whose proof is provided in Appendix D, establishes that dithered transcoders designed in this way are successful.

Proposition 3: Let $\hat{\mathbf{x}}^\circ$ be the reconstruction corresponding to an admissible rate- R_o encoder. Then there exists a $\delta_{\max} > 0$ such that for every $R > I(\hat{\mathbf{x}}^\circ; \hat{\mathbf{x}})$ and $\delta \in (0, \delta_{\max})$, $\Pr[\mathcal{E}] \rightarrow 0$ as $n \rightarrow \infty$ when the input is $\hat{\mathbf{x}}^\circ$. Furthermore, δ_{\max} only depends on R and $p_{\hat{\mathbf{x}}^\circ, \hat{\mathbf{x}}}(\hat{\mathbf{x}}^\circ, \hat{\mathbf{x}})$.

C. Successive Degradation Converse

To prove a converse, in this subsection we show that if the initial source encoder is a dithered quantizer, then $R_{\text{SD}}(d)$ from (2) gives a lower bound on the best rates that can be asymptotically guaranteed. A dithered quantizer is a worst case quantizer in the sense that the resultant encoding cannot be distinguished from a noisy observation of the source. That is, from the transcoder's point of view, the posterior $p_{\mathbf{x} | \hat{\mathbf{x}}^\circ}(\mathbf{x} | \hat{\mathbf{x}}^\circ) = \prod_i p_{x_i | \hat{x}_i^\circ}(x_i | \hat{x}_i^\circ)$.⁶ Therefore, the best any sort of requantization can do equals the corresponding indirect rate-distortion function.

To develop this result, it is convenient to express the payoff (11) for the successive degradation game in the form

$$\Pr\{d(\mathbf{x}, \hat{\mathbf{x}}) \leq d \mid \Phi^\circ, \Phi\} = E\{\theta_{d,\Phi}[\hat{\mathbf{x}}^\circ, \mathbf{x}, \Phi^\circ]\}$$

where

$$\theta_{d,\Phi}[\hat{\mathbf{x}}^\circ, \mathbf{x}, \Phi^\circ] \triangleq \begin{cases} 1, & \text{if } d(\mathbf{x}, \hat{\mathbf{x}}) \leq d \\ 0, & \text{otherwise.} \end{cases}$$

Then since a min-max upper bounds a max-min, to establish the converse of Theorem 1, it suffices to show the following.

⁶Note that this means that the Voronoi regions of a dithered quantizer overlap, while for a traditional nearest neighbor quantizer they are disjoint.

Proposition 4: For every $\epsilon > 0$ and rate $R \leq R_{\text{SD}}(d)$ there exists an admissible sequence of randomized encoders (in particular, a sequence of dithered encoders, $\{\Phi_n^\delta\}$) with

$$\sup_{\{\Phi_n\}} \lim_{n \rightarrow \infty} E\{\theta_{d,\Phi_n}[\hat{\mathbf{x}}^\circ, \mathbf{x}, \Phi_n^\delta]\} \leq \epsilon. \quad (17)$$

The expectation is over the source, codebooks, and encoding rules, and the maximization is over all sequences of rate- R transcoders.

Proof: First, consider the case when $R = R_o$. In this case, the transcoder has the option to output the index of the original source encoding along with a special flag to indicate that no transcoding has occurred. In this case, the rate and distortion are unchanged and the converse holds by the data processing inequality. This gives the converse to the expression for $R_{\text{SD}}(d)$ in (9) for the range $d_o \leq d \leq d_*$, which also corresponds to the second argument of the min in (2).

Next, consider the case when $R < R_o$. In this case, the original codebook cannot be used, so a new one must be generated. The key to our proof is showing that the transcoder input can be made by the dithered quantizer to “look” like a noisy observation of the source. Therefore, the converse results of noisy source coding can be applied directly to analyze the transcoder.

First, let $\hat{\mathbf{x}}^\circ$ be an encoding of \mathbf{x} created by the dithered quantizer $\Phi_n^\delta = (C_n^\delta, W_n^\delta(\cdot | \cdot))$. Then, we can express the original (dithered) quantizer in the form $\Phi_n^\delta = g(\hat{\mathbf{x}}^\circ)$ where $g(\cdot)$ is a random map. In particular, for $\mathbf{v} \in \mathcal{X}^n$, $g(\mathbf{v})$ consists of a codebook $C_n^\delta(\mathbf{v})$ with \mathbf{v} at a random entry and the remaining $(2^{nR_o} - 1)$ entries generated in a i.i.d. manner independently of \mathbf{v} , and a rule $W_n^\delta(\cdot | \cdot)(\mathbf{v})$ that is independent of \mathbf{v} . Hence, $g(\cdot)$ is independent of \mathbf{x} and $\hat{\mathbf{x}}^\circ$.

Next, let \mathbf{y} be a random vector generated from \mathbf{x} in an i.i.d. manner according to the distribution $p_{\hat{\mathbf{x}}^\circ | \mathbf{x}}$, independently of $g(\cdot)$. Then, we can use Lemma 1 along with Proposition 2 (to ensure that the original encoding succeeds) to infer that for all n sufficiently large

$$E\{\theta_{d,\Phi}[\hat{\mathbf{x}}^\circ, \mathbf{x}, g(\hat{\mathbf{x}}^\circ)]\} \leq E\{\theta_{d,\Phi}[\mathbf{y}, \mathbf{x}, g(\mathbf{y})]\} + \epsilon/2 \quad (18)$$

for any choice of transcoder Φ .

The expectation on the right-hand side of (18) is the probability that the distortion in quantizing a noisy observation of \mathbf{x} given the associated side information—i.e., the other codewords in the original codebook—is smaller than d . We note that such encoder side information does not help in source coding problems (see, e.g., [16], [17]). Therefore, if this $d < d_{\text{IC}}(R)$, the noisy distortion rate function defined by the channel law $p_{\hat{\mathbf{x}}^\circ | \mathbf{x}}$, then the converse of noisy rate-distortion coding [8], [9] tells us that no matter how Φ is chosen

$$E\{\theta_{d,\Phi}[\mathbf{y}, \mathbf{x}, g(\mathbf{y})]\} \leq \epsilon/2. \quad (19)$$

Substituting (19) into the right-hand side of (18) yields

$$E\{\theta_{d,\Phi}[\hat{\mathbf{x}}^\circ, \mathbf{x}, g(\hat{\mathbf{x}}^\circ)]\} \leq \epsilon$$

from which we obtain (17) as desired. This gives the converse to the first part of $R_{\text{SD}}(d)$ as expressed in (9) where $d \geq d_*$, corresponding to the first argument of the min in (2). \square

The inherent nonconvexity of the successive degradation rate–distortion function is reflected in the two separate cases of the converse: the transcoder must either leave the entire index unchanged or transcode the entire index.

D. Successive Degradation Direct Part

In this subsection, we establish that if the transcoding rate is larger than $R_{\text{SD}}(d)$ from (2), then the output of any admissible encoder can be transcoded successfully.

Since a max–min lower-bounds a min–max, to establish this forward part of Theorem 1 it suffices to show the following.

Proposition 5: For every $\epsilon > 0$ and rate $R > R_{\text{SD}}(d)$ there exists a sequence of dithered transcoders $\{\Phi_n^\delta\}$ with rates R such that

$$\inf_{\{\Phi_n^\delta\}} \lim_{n \rightarrow \infty} E \{ \theta_{d, \Phi_n^\delta} [\hat{\mathbf{x}}^\circ, \mathbf{x}, \Phi_n^\circ] \} \geq 1 - \epsilon.$$

The expectation is over the source, codebooks, and encoding and transcoding rules, and the infimum is taken over all admissible sequences of encoders with rate R_o .

Proof: In the range $d_o < d \leq d_*$, $R_{\text{SD}}(d)$ is easily achieved by simply outputting the index of the original source encoder along with a special flag to indicate that no transcoding has occurred. This achieves $R = R_{\text{SD}}(d)$. Distortions in the range $d_o < d \leq d_*$ are easily achieved by adding some amount of independent noise to $\hat{\mathbf{x}}^\circ$.

In the range $d > d_*$, let the transcoder be a dithered transcoder as designed in Sections IV-A and IV-B. Fix $\epsilon > 0$, then according to Proposition 3, this dithered transcoder succeeds in transcoding $\hat{\mathbf{x}}^\circ$ to $\hat{\mathbf{x}}$ with probability at least $1 - \epsilon$.

Therefore, \mathbf{x} and $\hat{\mathbf{x}}^\circ$ are strongly typical by the definition of admissible encoders and (12). Furthermore, the distribution of $(\hat{\mathbf{x}}^\circ, \hat{\mathbf{x}})$ is indistinguishable from an i.i.d. distribution according to Lemma 1, so $\hat{\mathbf{x}}^\circ$ and $\hat{\mathbf{x}}$ are also strongly typical. Thus, according to the Markov lemma [11, p. 436, Lemma 14.8.1], \mathbf{x} and $\hat{\mathbf{x}}$ are strongly typical according to the distribution

$$p(\hat{x}, x) = \sum_{\hat{x}^\circ} p(\hat{x} | \hat{x}^\circ) p(\hat{x}^\circ | x) p(x).$$

Hence the distortion will be close to $E[d(\mathbf{x}, \hat{\mathbf{x}})]$. \square

V. BINARY-HAMMING CASE

In this section, we derive the successive degradation rate–distortion function for the binary–Hamming case. As discussed in Section III-A, and we now show by example, the successive degradation rate–distortion function is generally different from the regular rate–distortion function.

Consider a binary (i.e., Bernoulli- p) source and a Hamming distortion measure

$$d(u, v) = \begin{cases} 1, & u \neq v \\ 0, & u = v \end{cases}$$

so the (average) distortion range of interest is the interval $[0, 1/2]$. In this case, it is well known that the optimal reverse test channel for the original quantization is a binary-symmetric channel with crossover probability d_o . As we outline later, based upon results in [9], it is straightforward to verify that

to optimize the first term in (2), i.e., $I(\hat{\mathbf{x}}; \hat{\mathbf{x}}^\circ)$, the successive degradation reverse test channel $p_{\hat{\mathbf{x}}^\circ | \hat{\mathbf{x}}}(\hat{\mathbf{x}}^\circ | \hat{\mathbf{x}})$ should also take the form of a binary-symmetric channel, but with crossover probability $0 \leq \gamma \leq 1/2$. For this choice, we quickly derive the corresponding rate and overall distortion. The rate is

$$R = H(\hat{\mathbf{x}}^\circ) - H(\hat{\mathbf{x}}^\circ | \hat{\mathbf{x}}) = H_B(q) - H_B(\gamma)$$

with $H_B(\cdot)$ denoting the binary entropy function and $q = p_{\hat{\mathbf{x}}^\circ}(1)$. Since

$$p_x(1) = p = \sum_{\hat{x}} p_{x|\hat{x}}(1|\hat{x}) p_{\hat{x}}(\hat{x}) = (1 - d_o)q + d_o(1 - q)$$

we get

$$q = \frac{p - d_o}{1 - 2d_o}. \quad (20)$$

By evaluating the overall transcoding distortion d , we now solve for γ . The concatenation of the two symmetric reverse test channels gives a distortion

$$\begin{aligned} d &= \Pr[x \neq \hat{x}] \\ &= \Pr[x \neq \hat{x}^\circ, \hat{x}^\circ = \hat{x}] + \Pr[x = \hat{x}^\circ, \hat{x}^\circ \neq \hat{x}] \\ &= \Pr[x \neq \hat{x}^\circ] \Pr[\hat{x}^\circ = \hat{x}] + \Pr[x = \hat{x}^\circ] \Pr[\hat{x}^\circ \neq \hat{x}] \\ &= d_o(1 - \gamma) + (1 - d_o)\gamma = d_o + \gamma(1 - 2d_o) \end{aligned}$$

where the second line follows since the test channels are independent. Solving for γ gives

$$\gamma = \frac{d - d_o}{1 - 2d_o}. \quad (21)$$

Substituting the values of γ and q into $R = H_B(q) - H_B(\gamma)$ gives the rate–distortion tradeoff achieved by this choice of test channel

$$R = H_B\left(\frac{p - d_o}{1 - 2d_o}\right) - H_B\left(\frac{d - d_o}{1 - 2d_o}\right). \quad (22)$$

By comparison, the corresponding regular rate–distortion bound for this source and distortion measure is

$$R_{\text{SC}}(d) = H_B(p) - H_B(d). \quad (23)$$

Since, without loss of generality we may assume $0 \leq d \leq p \leq 1/2$, $R_{\text{SD}}(d) \geq H_B(p) - H_B((d - d_o)/(1 - 2d_o))$, and since $H_B(\cdot)$ is strictly monotonic on $[0, 1/2]$, it follows that $R_{\text{SD}}(d) > R_{\text{SC}}(d)$ whenever $d_o \neq 0$.

As already discussed, for some range of distortions above d_o , requantization via the test channel giving rise to (22) does not yield rate savings and transcoding is better avoided. In this case, the second term of (2) is minimal and the general successive degradation rate–distortion function is

$$R_{\text{SD}}(d) = \min \left\{ H_B\left(\frac{p - d_o}{1 - 2d_o}\right) - H_B\left(\frac{d - d_o}{1 - 2d_o}\right), R_o \right\} \quad (24)$$

where $d \geq d_o$ and with $R_o = R_{\text{SC}}(d_o)$ given by (23).

By equating the two terms in (24) and solving for d , we obtain the distortion threshold d_* below which transcoding should

be avoided. A simple lower bound on d_* is found by lower-bounding the first term of (24) with $H_B(p)$ and setting the result equal to R_o . This gives

$$d_* \geq 2d_o(1 - d_o).$$

It remains only to confirm that there is no better choice of successive degradation reverse test channel than the binary-symmetric channel that leads to (22). Since a dithered encoder produces an output that is indistinguishable from a noisy version of the source, all we need to verify is that the proposed reverse test channel corresponds to the indirect source coding test channel. In [9], the indirect rate–distortion function was developed for a binary source observed through a binary-symmetric channel. A symmetric reverse test channel was shown to be optimal in that setting. If the source distribution $p_x(1) = p = 1/2$, then the symmetric reverse test channel of the original quantization can be re-expressed as a symmetric forward test channel (shown via Bayes' rule, using the prior and the crossover probabilities). Since this matches the binary-symmetric observation channel of [9], that result confirms that a symmetric reverse channel is the best choice for $p_{\hat{x}^\circ | \hat{x}}(\hat{x}^\circ | \hat{x})$. However, if $p \neq 1/2$, the symmetric reverse test channel of the original quantization gives rise to an asymmetric forward test channel $p_{\hat{x}^\circ | x}(\hat{x}^\circ | x)$. Asymmetric observations are not considered in [9], so those results do not immediately confirm that a symmetric $p_{\hat{x}^\circ | \hat{x}}(\hat{x}^\circ | \hat{x})$ is the correct choice. However, the results of [9] can be generalized to asymmetric observations in a straightforward manner. The only change in the derivation is that the two types of bit-error events (cf. [9, eq. (38)]) are now weighted by the (nonequal) prior probabilities of the observation symbols to be encoded. In our setting, the observation symbols correspond to the source symbols of the original reconstruction. One can check that the error weighting for the indirect coding of an asymmetric source corresponds to the same uneven likelihoods of bit errors induced by a symmetric successive degradation reverse test channel with crossover probability γ . This means that even for $p \neq 1/2$, a symmetric successive degradation reverse test channel is the best choice for this source and (24) is the rate–distortion function.

The associated rate gap for this example takes the form depicted in Fig. 5 for the symmetric case ($p = 1/2$). The successive degradation and regular rate–distortion functions are depicted with the solid and dashed curves, respectively, and intersect at distortion $d = 1/2$.

VI. GAUSSIAN-QUADRATIC CASE

In this section, we develop the rate–distortion expression (2) for Gaussian sources under a quadratic distortion measure: $d(a, b) = (a - b)^2$. We also extend the class of admissible encoders for achievability results in the Gaussian quadratic case to all encoders that achieve d_o and not just those encoders that satisfy the divergence condition of Definition 2. This shows that Gaussian-quadratic transcoding is robust in the sense that any good source code for the Gaussian-quadratic case can be successfully transcoded.




Fig. 5. Common form of the rate–distortion function for the symmetric ($p = 1/2$) binary-Hamming and Gaussian-quadratic cases. The successive degradation rate–distortion function is indicated by the solid curve, and d_* is the threshold above which transcoding should be performed. The dashed curve indicates the lower bound corresponding to the regular rate–distortion function. The distortion at which both the dashed and solid curves intersect the axis ($R = 0$) is $1/2$ in the binary-Hamming case, and σ_x^2 in the Gaussian-quadratic case.

The proofs in Section IV assumed that the signals were drawn from finite alphabets and that all distortion measures were bounded. This simplified the development, but the results can be generalized to continuous alphabets with unbounded distortion. Among other subtleties, care must be taken to preserve the Markov relationship $x \leftrightarrow \hat{x}^\circ \leftrightarrow \hat{x}$. Such techniques are developed in, e.g., [18]. However, in the sequel, we establish the achievable transcoding tradeoff for the Gaussian-quadratic scenario more directly.

A. Rate–Distortion Function

For an i.i.d. zero-mean Gaussian source x of variance σ_x^2 , and an original source code with average mean-square reconstruction distortion d_o , the information successive degradation rate–distortion function given by (2) is, for $d \geq d_o$

$$R_{SD}(d) = \min \left\{ \frac{1}{2} \log \left(\frac{\sigma_x^2 - d_o}{d - d_o} \right), R_o \right\}. \quad (25)$$

where, according to the familiar Gaussian-quadratic rate–distortion function

$$R_o = \frac{1}{2} \log \left(\frac{\sigma_x^2}{d_o} \right). \quad (26)$$

This successive degradation rate–distortion function also takes the form shown in Fig. 5. Again, the successive degradation and regular rate–distortion functions are depicted with the solid and dashed curves, respectively, and intersect at distortion $d = \sigma_x^2$.

Equation (25) is obtained as follows. The usual conditional distribution for achieving the rate–distortion bound for Gaussian sources, which we assume to be unique, corresponds to $\hat{x}^\circ = \alpha x + e_o$, where

$$\alpha = (1 - d_o/\sigma_x^2) \quad (27)$$

and where e_o is a zero-mean Gaussian random variable with variance αd_o that is independent of x . Since \hat{x}° is Gaussian, we know from [19, Lemma A.3] that the optimum \hat{x} and \hat{x}° must also be jointly Gaussian. Optimizing the specific moments we conclude that the second conditional distribution is of the form $\hat{x} = \gamma \hat{x}^\circ + e$, where $\gamma = (1 - \Delta/\sigma_{\hat{x}^\circ}^2)$, and where e is a

zero-mean Gaussian random variable with variance $\gamma\Delta$ that is independent of \hat{x}° . Combining the two conditional distributions results in a rate

$$I(\hat{x}; \hat{x}^\circ) = \frac{1}{2} \log \left(\frac{\sigma_x^2 - d_o}{\Delta} \right), \quad \Delta \geq 0 \quad (28)$$

with overall mean-square distortion

$$d = E[(\hat{x} - x)^2] = d = d_o + \Delta. \quad (29)$$

Substituting for Δ in (28) using (29), and we obtain the first term in (25). As in the binary-Hamming case, reference to [9] confirms that this is the best choice of test channel.

The minimization in (25) suggests that there is a range of distortions above d_o for which requantization will not save any rate, i.e., it is better to leave the original source code as is. For example, if we set $\Delta = 0$, we get $d = d_o$ at the cost of infinite rate. A better strategy is to leave the original description unmodified. Again, this is inherently a result of the quantized nature of the input to the transcoder. To identify the distortion threshold d_* at which successive degradation becomes effective, we substitute (26) into (25), equate the two resulting terms, and solve for d , yielding

$$d_* = d_o \left(2 - \frac{d_o}{\sigma_x^2} \right). \quad (30)$$

In the Gaussian-quadratic scenario, the successive degradation rate loss is at most 1/2 bit, which we see as follows. To begin, note that $1 \leq d_*/d_o \leq 2$, with $d_*/d_o \rightarrow 2$ as $R_o \rightarrow \infty$. Thus, in the high-resolution limit, each time the source is requantized with an independently generated source code of the same rate, the overall distortion increases by roughly d_o . Since the rate loss in successive degradation is largest in this high-resolution limit, an upper bound on the loss occurs at $d = d_*$

$$\begin{aligned} R_{SD}(d_*) - R_{SC}(d_*) &\leq \frac{1}{2} \log \left(\frac{\sigma_x^2}{d_o} \right) - \frac{1}{2} \log \left(\frac{\sigma_x^2}{d_o(2 - d_o/\sigma_x^2)} \right) \\ &= \frac{1}{2} \log \left(2 - \frac{d_o}{\sigma_x^2} \right) < \frac{1}{2} \end{aligned} \quad (31)$$

where the last inequality follows from the fact that $d_o \geq 0$. Recall that, by comparison, there is no rate loss with successive refinement in this Gaussian-quadratic scenario [1]. This is because the full original source signal—not just a quantized version of it—is available when selecting the second codeword. Finally, from Fig. 5 and (30), we see that the corresponding distortion gap is at most 3 dB

$$\frac{d_*}{d_o} = 2 - \frac{d_o}{\sigma_x^2} < 2.$$

B. Converse

Since the dithered encoding rule of Section IV-A only applied to discrete alphabets, we turn to a different type of dithering to prove the converse for the Gaussian case. We use entropy coded dithered quantization (ECDQ) [20] for the original source encoder. For the Gaussian-quadratic case, this is the natural counterpart of the finite-alphabet dithered quantization strategy introduced in Section IV-A. With this source encoder, the quantized source becomes asymptotically indistinguishable

in an appropriate sense [20] from an observation of the original source seen through the channel that corresponds to the forward test channel of Gaussian quantization, i.e., an attenuating channel with additive white Gaussian noise. Therefore, the best achievable requantization rate–distortion tradeoff is the indirect rate–distortion function for this source and channel. This function was shown in [9] to be given by the first term in (25). The overall expression (25) is, therefore, the successive degradation rate–distortion function.

C. Achievability

We now generalize our achievability results in the Gaussian case by showing that for *any* good source coder that operates close to its optimal distortion level, (25) is an upper bound on the operational successive degradation rate–distortion function. We extend our achievability beyond the admissible set of encoders introduced in Definition 2 by making use of a fundamental result of [21]. In this sense, our results in this case are inherently more robust.

Specifically, we show that for any rate- R_o original encoder operating close to its optimal distortion level

$$d_{SC}(R_o) = \sigma_x^2 2^{-2R_o}$$

a rate- R transcoder where $R < R_o$ can be designed so that it operates close to the information successive degradation distortion–rate function

$$d_{SD}(R) = d_{SC}(R_o) + (\sigma_x^2 - d_{SC}(R_o))2^{-2R} \quad (32)$$

obtained from (3) with (25).

We quantify this argument in the following theorem, which makes use of a basic result from [21].

Theorem 2: Let \mathbf{x} be a length- n i.i.d. sequence of Gaussian random variables with variance σ_x^2 . For any $\epsilon > 0$ and any rate- R_o original source code with

$$\frac{1}{n} E[\|\mathbf{x} - \hat{\mathbf{x}}^\circ\|^2] \leq d_{SC}(R_o) + \epsilon \quad (33)$$

there exists a rate- R transcoder for $R < R_o$ with

$$\frac{1}{n} E[\|\mathbf{x} - \hat{\mathbf{x}}\|^2] \leq d_{SD}(R) + \epsilon + \delta(n) \quad (34)$$

where $\delta(n) \rightarrow 0$ as $n \rightarrow \infty$.

Proof: To obtain this result, we first bound the variance of a processed version of the output of the original coder. We then use a result from [21] on quantizing sources given only second-order statistics.

Let $\hat{\mathbf{x}}'$ be the minimum mean-square error (MMSE) estimate of \mathbf{x} given $\hat{\mathbf{x}}^\circ$. Since $\hat{\mathbf{x}}'$ is a reconstruction of \mathbf{x} based on nR_o bits and since $\hat{\mathbf{x}}'$ estimates \mathbf{x} (in the mean-square sense) at least as well as $\hat{\mathbf{x}}^\circ$, we know that

$$d_{SC}(R_o) \leq \frac{1}{n} E[\|\mathbf{x} - \hat{\mathbf{x}}'\|^2] \leq \frac{1}{n} E[\|\mathbf{x} - \hat{\mathbf{x}}^\circ\|^2]. \quad (35)$$

Furthermore, the error in the MMSE estimate is uncorrelated with the reconstruction itself, i.e.,

$$E[\langle (\mathbf{x} - \hat{\mathbf{x}}'), \hat{\mathbf{x}}' \rangle] = 0 \quad (36)$$

whence, by Pythagoras' theorem

$$\frac{1}{n}E[\|\hat{\mathbf{x}}'\|^2] = \frac{1}{n}E[\|\mathbf{x}\|^2] - \frac{1}{n}E[\|\mathbf{x} - \hat{\mathbf{x}}'\|^2] \quad (37)$$

$$\leq \sigma_x^2 - d_{\text{SC}}(R_o) \quad (38)$$

where (37) follows from (36), and (38) follows from the left-hand inequality in (35) and the fact that \mathbf{x} is an i.i.d. variance- σ_x^2 Gaussian sequence.

Consider a random transcoder which encodes $\hat{\mathbf{x}}'$ by mapping it to the element $\hat{\mathbf{x}}$ of a rate- R random Gaussian codebook that is closest in Euclidean distance. Regardless of the distribution of $\hat{\mathbf{x}}$, via [21, Theorem 3] we know there exists a deterministic transcoder with output $\hat{\mathbf{x}}$ such that

$$\frac{1}{n}E[\|\hat{\mathbf{x}}' - \hat{\mathbf{x}}\|^2] \leq \frac{1}{n}E[\|\hat{\mathbf{x}}'\|^2]2^{-2R} + \delta(n) \quad (39)$$

where $\delta(n) \rightarrow 0$ as $n \rightarrow \infty$. The underlying result [21, Theorem 3] basically says that any source $\hat{\mathbf{x}}'$ of a given second moment $E[\|\hat{\mathbf{x}}'\|^2]/n$ can be encoded into a codebook designed for a Gaussian source with the same second moment, and the distortion achieved will equal that achieved if the source $\hat{\mathbf{x}}'$ were Gaussian. Note that since a Gaussian source is the hardest to encode, we would do better if we took into account more detailed information about the source statistics.

Due to the structure of the two encodings we have the Markov chain $\mathbf{x} \leftrightarrow \hat{\mathbf{x}}' \leftrightarrow \hat{\mathbf{x}}$ and hence, $(\mathbf{x} - \hat{\mathbf{x}}') \leftrightarrow \hat{\mathbf{x}}' \leftrightarrow \hat{\mathbf{x}}$ as well. With this latter Markov chain we conclude from the optimality properties of MMSE estimates that (36) yields

$$E[\langle (\mathbf{x} - \hat{\mathbf{x}}'), \hat{\mathbf{x}} \rangle] = 0. \quad (40)$$

To complete the proof, it suffices to observe that

$$\begin{aligned} \frac{1}{n}E[\|\mathbf{x} - \hat{\mathbf{x}}\|^2] &= \frac{1}{n}E[\|(\mathbf{x} - \hat{\mathbf{x}}') + (\hat{\mathbf{x}}' - \hat{\mathbf{x}})\|^2] \\ &= \frac{1}{n}E[\|\mathbf{x} - \hat{\mathbf{x}}'\|^2] + \frac{1}{n}E[\|\hat{\mathbf{x}}' - \hat{\mathbf{x}}\|^2] \end{aligned} \quad (41)$$

$$\leq d_{\text{SC}}(R_o) + (\sigma_x^2 - d_{\text{SC}}(R_o))2^{-2R} + \epsilon + \delta(n) \quad (42)$$

where to obtain (41) we have used the orthogonality implied by (36) and (40), and where to obtain (42) we have used the right-hand inequality in (35) with (33), and (38) with (39). Combining (42) with (32) give (34). \square

VII. CONTINUOUS SOURCES IN THE HIGH-RESOLUTION LIMIT

In this section, we show that the 1/2-bit gap bound (31) in the high-resolution limit with a quadratic distortion measure holds not just for Gaussian sources as developed in Section VI, but in fact for all sources with finite differential entropy and at least one finite moment. In fact, we believe that the high-resolution limit is the worst case and that the gap is at most 1/2-bit for all distortions.

First, from [22], we know that the successive refinement rate-distortion function is within 1/2-bit of the regular rate-distortion function for this scenario. Thus, it remains only to show that the successive degradation rate-distortion function also

lies within 1/2-bit of the regular rate-distortion function as well, which we develop in the sequel.

Our proof considers two separate regimes. For $d < 2d_o$, we stay within the 1/2-bit bound by avoiding transcoding altogether, whereby $R = R_{\text{SC}}(d_o)$. To see this, note that the rate loss in this case is largest as $d_o \rightarrow 0$, given by

$$R_{\text{SC}}(d_o) - R_{\text{SC}}(d) \rightarrow \frac{1}{2} \log \left(\frac{d}{d_o} \right) \leq \frac{1}{2} \quad (43)$$

where to obtain the limit in (43) we have used the asymptotic tightness of the Shannon lower bound [23]

$$\lim_{\delta \rightarrow 0} R_{\text{SC}}(\delta) = h(x) - \frac{1}{2} \log 2\pi e \delta \quad (44)$$

with $h(\cdot)$ denoting differential entropy.

For the regime $d \geq 2d_o$, we use the following argument. First, we let $\hat{x} = \hat{x}^\circ + e$ where e is a zero-mean Gaussian random variable with variance $d - d_o$ that is independent of \hat{x}° and x , so the distortion is, as required

$$\begin{aligned} E[(x - \hat{x})^2] &= E[(x - \hat{x}^\circ + \hat{x}^\circ - \hat{x})^2] \\ &= E[(x - \hat{x}^\circ)^2] + E[(\hat{x}^\circ - \hat{x})^2] \\ &\quad + 2E[(x - \hat{x}^\circ)(\hat{x}^\circ - \hat{x})] \\ &= d_o + E[e^2] - 2E[(x - \hat{x}^\circ)e] \end{aligned} \quad (45)$$

$$\begin{aligned} &= d_o + (d - d_o) + 0 \\ &= d \end{aligned} \quad (46)$$

where to obtain (45) and (46) we have used the definition and properties, respectively, of e . Then from (2) we see the associated rate is

$$\begin{aligned} R_{\text{SD}}(d) &\leq I(\hat{x}; \hat{x}^\circ) \\ &= h(\hat{x}) - h(\hat{x} | \hat{x}^\circ) \\ &= h(\hat{x}^\circ + e) - h(\hat{x}^\circ + e | \hat{x}^\circ) \\ &= h(\hat{x}^\circ + e) - h(e) \end{aligned} \quad (47)$$

where to obtain (47) we have used that e and \hat{x}° are independent. In turn, we can bound the rate loss as $d \rightarrow 0$ via

$$\begin{aligned} R_{\text{SD}}(d) - R_{\text{SC}}(d) &\leq h(\hat{x}^\circ + e) - h(e) - R_{\text{SC}}(d) \\ &= [h(\hat{x}^\circ + e) - h(\hat{x}^\circ)] + [h(\hat{x}^\circ) - h(x)] \\ &\quad + [h(x) - R_{\text{SC}}(d)] - h(e) \end{aligned} \quad (48)$$

$$\rightarrow 0 + 0 + \frac{1}{2} \log 2\pi e d - h(e) \quad (49)$$

$$\begin{aligned} &= \frac{1}{2} \log \left(\frac{d}{d - d_o} \right) \\ &\leq \frac{1}{2} \end{aligned} \quad (50)$$

where to obtain (48) we have used (47), and where to obtain the first and third terms in (49) we have used, respectively, the continuity of $h(\cdot)$ in [23, Theorem 1], and the high-resolution rate-distortion function [23], and where to obtain (50) we have used that the differential entropy of a Gaussian random variable z of variance σ^2 is $h(z) = \frac{1}{2} \log 2\pi e \sigma^2$.

To obtain that the second term in (49) is zero, it suffices to let $\hat{x}^\circ = x + e_o$ where e_o is a zero-mean Gaussian random variable

with variance d_o that is independent of x , and again exploit the continuity of $h(\cdot)$ [23, Theorem 1] to conclude

$$h(\hat{x}^\circ) - h(x) = h(x + e_o) - h(x) \rightarrow 0 \text{ as } d_o \rightarrow 0.$$

Thus, it remains only to check that the conditional distribution associated with this definition of e_o asymptotically achieves the rate–distortion function $R_{\text{SC}}(d_o)$. To see this, as $d \rightarrow 0$, whence $d_o \rightarrow 0$ since $d \geq d_o \geq 0$, we have

$$\begin{aligned} R_{\text{SC}}(d) &\leq I(\hat{x}^\circ, x) \\ &= h(\hat{x}^\circ) - h(\hat{x}^\circ | x) \\ &= h(x + e_o) - h(x + e_o | x) \end{aligned} \quad (51)$$

$$= h(x + e_o) - h(e_o) \quad (52)$$

$$= \left[h(x) - \frac{1}{2} \log 2\pi e d_o \right] + [h(x + e_o) - h(x)] \quad (53)$$

$$\rightarrow \left[h(x) - \frac{1}{2} \log 2\pi e d_o \right] + 0 \quad (54)$$

$$= \lim_{d_o \rightarrow 0} R_{\text{SC}}(d_o) \quad (55)$$

where to obtain (51), (52), and (53) we have used the definition and properties of e_o , where to obtain (54) we have used the continuity of differential entropy, and where to obtain (55) we have used (44).

VIII. EMBEDDING IN A QUANTIZED SOURCE

Up to this point, we have considered bit stealing approaches that take the form of rate splitting and successive degradation. In this paradigm, the transcoder is given the freedom to design an entirely new codebook. This requires that the ultimate destination(s) for the quantized source be informed that transcoding has taken place so that the destination(s) learn to decode using the new codebook.

However, in a number of scenarios it may be either impractical or inconvenient to inform the decoder when bit stealing has taken place. Such is the case, for example, when there is an installed base of legacy source decoders in a network, or, as another example, when the bit stealing is to be covert, in which case no cooperation between transcoders and source decoders is possible. In these and other such cases, there is a need for bit stealing techniques in which the transcoder output lies in the same codebook as its input.

One natural approach to bit stealing with this constraint is based on the use of information embedding ideas. In this section, we give an achievability result that shows that bit stealing systems of this type can, in fact, be as efficient as those implemented through rate splitting and successive degradation. From this we can conclude that the transcoder codebook constraint need not incur a loss in performance. To develop this result, it suffices to restrict our attention to the case where our original source \mathbf{x} has been encoded using a random codebook via a joint typicality encoding rule. While this restriction is for the simplicity and self-containment of the presentation, the results can be applied directly to the class of admissible encoders, Definition 2, which include dithered encoders. This is because the




Fig. 6. Bit stealing via information embedding: a message m of rate r is embedded into a source quantized to rate R_o . Source decoder I is the decoder used in the absence of transcoding, or equivalently, the uninformed decoder when transcoding has taken place, in which case it produces reconstruction \hat{x}^u at distortion d_u . Source decoder II is the decoder that is informed of the rate of any embedding that has taken place and produces reconstruction \hat{x} at distortion d . The message decoder produces a reliable estimate \hat{m} of the embedded message with high probability.

admissibility definition implies joint strong typicality between source and reconstruction according to the joint distribution induced by the source and the (assumed unique) test channel. Such typicality is what we require in the ensuing construction.

As depicted in Fig. 6, bit stealing via embedding is implemented as follows. The transcoder embeds a message m of rate r into the index corresponding to source reconstruction codeword \hat{x}° , which is the “host.” A source decoder generates the reconstruction \hat{x} from the received bits. Recall that informed source decoders are aware of any transcoding that has taken place—specifically, they know what rate $0 \leq r \leq R_o$ has been stolen. Informed decoders can exploit this information in reconstructing the source, while uninformed decoders operate as if no embedding had taken place. In the sequel, we consider both informed and uninformed source decoders. A message decoder recovers the embedded message bits from the received bits.

A naive embedding approach would treat the source reconstruction \hat{x}° as the host, i.e., as the “dirty paper” of [24], and use the associated random binning code or one of its constructive counterparts such as quantization index modulation (QIM) [25]–[27] to embed the message. Such approaches fail to exploit that \hat{x}° is a codeword of a finite-rate codebook. And further, they assume the use of a new quantization codebook, rather than embedding in the original codebook. However, for embedding approaches that do take such problem characteristics into account, the following rate–distortion tradeoff can be achieved.

In constructing our result, we continue to focus for simplicity on an i.i.d. source—as with successive degradation—so an instance of the problem continues to consist of the tuple (10), where as before R_o implicitly defines $p_{\hat{x}^\circ | x}(\hat{x}^\circ | x)$, but where, now, $R = R_o - r$. And as before $p_{\hat{x}^\circ | x}(\hat{x}^\circ | x)$ denotes the conditional distribution characterizing the rate–distortion function at distortion d_o where $R_o = R_{\text{SC}}(d_o)$.

Theorem 3: For a source \mathbf{x} quantized at rate R_o (via typicality encoding) to a codeword \hat{x}° in a codebook generated randomly according to $p_{\hat{x}^\circ | x}(\hat{x}^\circ | x)$, a distortion arbitrarily close to $d \geq d_o$ is achievable if $r \leq R_{\text{IE}}(d)$ where [cf. (2)]

$$R_{\text{IE}}(d) = \max\{R_o - \inf I(u; \hat{x}^\circ), 0\} \quad (56)$$

where the infimum is taken over all conditional distributions $p_{u | \hat{x}^\circ}(u | \hat{x}^\circ)$ and functions $f : \mathcal{X} \mapsto \mathcal{X}$ such that the Markov constraint $x \leftrightarrow \hat{x}^\circ \leftrightarrow u$ is satisfied, $p_{\hat{x}^\circ}(x) = p_u(x)$, and

$E[d(x, \hat{x})] \leq d$, where $\hat{x} = f(u)$ and u is an auxiliary random variable with alphabet \mathcal{X} .

This theorem applies when the decoder is informed of the embedding, and $f(\cdot)$ can be viewed as a distortion compensation function. Achievable rates for the case in which the decoder is uninformed are obtained by constraining $f(\cdot)$ to be the identity function. Just as with the successive degradation development, in the case of information embedding there is a range of distortions for which no positive rate embedding is possible.

Before proving Theorem 3, we introduce some additional notation. Specifically, $d_{\text{IE}}(r)$ is the distortion rate function corresponding to (56), i.e.,

$$d_{\text{IE}}(r) = \inf\{d \geq 0 : R_{\text{IE}}(d) \geq r\}. \quad (57)$$

Furthermore, we let $R_{\text{IE}}^u(\cdot)$ and $d_{\text{IE}}^u(\cdot)$ denote the corresponding rate–distortion and distortion–rate functions for *uninformed* decoders.

We establish Theorem 3 by construction. In particular, the original source is quantized using a randomly generated codebook. The source codebook \mathcal{C}° consists of 2^{nR_o} sequences of length- n generated according to $\prod_{i=1}^n p_{\hat{x}^\circ}(\hat{x}_i^\circ)$. These sequences are labeled $\hat{\mathbf{x}}^\circ(1), \hat{\mathbf{x}}^\circ(2), \dots, \hat{\mathbf{x}}^\circ(2^{nR_o})$. To encode, we find the index i such that $(\mathbf{x}, \hat{\mathbf{x}}^\circ(i)) \in T_{\mathbf{x}, \hat{\mathbf{x}}^\circ}$. If there is more than one such index, choose any one of them. Transmit that i . If there is no such index, declare an encoding failure.

Information embedding is implemented as follows. First, a message code is constructed whereby, for any $\epsilon > 0$, we randomly bin \mathcal{C}° into 2^{nr} subcodes \mathcal{C}_j° where $r = R_o - I(u; \hat{x}^\circ) - \epsilon$. Specifically, for each $\hat{\mathbf{x}}^\circ(i)$, we pick an index j uniformly distributed over $\{1, 2, \dots, 2^{nr}\}$ and assign $\hat{\mathbf{x}}^\circ(i)$ to subcode \mathcal{C}_j° . On average, there are $2^{n(R_o - r)}$ sequences $\hat{\mathbf{x}}^\circ$ in each \mathcal{C}_j° . We relabel the $\hat{\mathbf{x}}^\circ$ sequences in \mathcal{C}° as $\mathbf{u}(j, k)$ where $j \in \{1, 2, \dots, 2^{nr}\}$ and $k \in \{1, 2, \dots, |\mathcal{C}_j^\circ|\}$. This partitioning and labeling is then shared with the message decoder.

Message encoding is accomplished as follows. Given a source codeword $\hat{\mathbf{x}}^\circ(i)$, and a message $m = m$, find the index k such that $(\hat{\mathbf{x}}^\circ(i), \mathbf{u}(m, k)) \in T_{\hat{\mathbf{x}}^\circ, u}$. This can be done, e.g., using the dithered transcoder. If there is more than one such index pick any one. Transmit the index l such that $\hat{\mathbf{x}}^\circ(l) = \mathbf{u}(m, k)$. If there is no such index, declare an embedding failure.

With this encoding and embedding, decoding is straightforward: the source reconstruction is $\hat{\mathbf{x}}$ with $\hat{x}_i = f(\hat{x}_i^\circ(l))$, and the message estimate is $\hat{m} = m$ such that $\hat{\mathbf{x}}(m, k) = \hat{\mathbf{x}}^\circ(l)$. It remains only to ensure the error probability vanishes and the distortion constraint is met, which we verify in the sequel.

That the probability of a source encoding failure goes to zero follows from joint strong typicality, since $R_o > I(\hat{x}^\circ; \mathbf{x})$. The probability of an embedding failure also goes to zero with large n . To see this, first note that the probability that the original source-quantization vector $\hat{\mathbf{x}}^\circ(i)$ falls into the selected bin m is 2^{-nr} , which goes to zero for n large. Moreover, conditioned on the event that $\hat{\mathbf{x}}^\circ(i)$ is not in bin m , the codewords in bin m look like i.i.d. sequences generated independently of $\hat{\mathbf{x}}^\circ(i)$ according to $\prod_{i=1}^n p_u(x_i)$. Indeed, $\mathcal{C}_m^\circ \subset \mathcal{C}^\circ$, the entries of \mathcal{C}° are generated independently according to $\prod_{i=1}^n p_{\hat{x}^\circ}(x_i)$, and $p_{\hat{x}^\circ}(x) = p_u(x)$. The probability that at least one such sequence $\mathbf{u}(m, k)$ is jointly strongly typical with $\hat{\mathbf{x}}^\circ(i)$ approaches one because there are $2^{n(R_o - r)}$ codewords in bin m and, via (56),

$R_o - r > I(u; \hat{x}^\circ)$. The probability that $\hat{m} \neq m$ is zero because the decoder has direct access to the embedder output.

To see that the distortion constraint d is met, we first note that $(\mathbf{x}, \mathbf{u}) \in T_{\mathbf{x}, u}$ by the Markov lemma [11, p. 436, Lemma 14.8.1]. Indeed, $\mathbf{x} \leftrightarrow \hat{\mathbf{x}}^\circ \leftrightarrow u$, and we have both $(\mathbf{x}, \hat{\mathbf{x}}^\circ) \in T_{\mathbf{x}, \hat{\mathbf{x}}^\circ}$ and $(\hat{\mathbf{x}}^\circ, \hat{\mathbf{x}}^\circ(l)) \in T_{\hat{\mathbf{x}}^\circ, u}$. Hence, by choosing ϵ small enough and n large enough, and by exploiting that d is bounded, we can obtain, for any $\delta > 0$

$$\begin{aligned} E[d(\mathbf{x}, \hat{\mathbf{x}})] &= \frac{1}{n} \sum_{i=1}^n E[d(x_i, f(\hat{x}_i^\circ(l)))] \\ &= \sum_{x, u} E[T_{\mathbf{x}, \hat{\mathbf{x}}^\circ(l)}(x, u)] d(x, f(u)) \\ &\leq d + \delta \end{aligned}$$

which establishes the theorem. \square

Because of the transcoder output codebook constraint, we have, in general

$$R_{\text{SD}}(d) + R_{\text{IE}}(d) \leq R_o. \quad (58)$$

We now develop cases in which (58) holds with a strict inequality, and when it holds with equality.

A. Binary-Hamming Case

As we now show by example, constraining the transcoder output codebook to coincide with the input codebook according to Theorem 3 generally incurs a loss in performance.

To see this, consider again the case of a Bernoulli- p source. In this case, because of the codebook constraint $p_{\hat{x}^\circ}(x) = p_u(x)$, the information embedding reverse test channel can no longer be a binary-symmetric channel as it was for successive degradation, but rather must be an asymmetric channel

$$\begin{aligned} p_{\hat{x}^\circ | u}(0 | 1) &= \alpha \\ p_{\hat{x}^\circ | u}(1 | 0) &= \beta. \end{aligned}$$

Moreover, for this binary case we can, without loss of generality, skip the distortion compensation in Theorem 3 (i.e., let $f(\cdot)$ be the identity function so that $\hat{x} = u$). Thus, in a manner analogous to the way we obtained the successive degradation rate–distortion function for the binary-Hamming case in Section V, we obtain, via (56), that

$$\begin{aligned} r &= R_o - I(u; \hat{x}^\circ) = R_o - H(\hat{x}^\circ) + H(\hat{x}^\circ | u) \\ &= R_o - H_B(q) + p_u(1)H_B(\alpha) + p_u(0)H_B(\beta) \\ &= R_o - H_B(q) + qH_B(\alpha) + (1 - q)H_B(\beta) \end{aligned} \quad (59)$$

where

$$q \triangleq p_{\hat{x}^\circ}(1) = p_u(1) = \frac{p - d_o}{1 - 2d_o} \quad (60)$$

and the derivation of (60) is identical to that of (20) in Section V.

The marginal distribution constraint $p_u(1) = p_{\hat{x}^\circ}(1) = q$ gives a relationship between α and β

$$\begin{aligned} p_{\hat{x}^\circ}(1) &= p_{\hat{x}^\circ | u}(1 | 1)p_u(1) + p_{\hat{x}^\circ | u}(1 | 0)p_u(0) \\ q &= (1 - \alpha)q + \beta(1 - q) \\ \beta &= \alpha \frac{q}{1 - q}. \end{aligned}$$




Fig. 7. Rate–distortion loss for bit stealing via information embedding in the case of a Bernoulli-2/5 source and Hamming distortion measure, with $R_o = 1/4$. The progressively lower solid, dashed, and dotted curves correspond to $R_o - R_{IE}(d)$, $R_{SD}(d)$, and $R_{RD}(d)$, respectively, i.e., bit stealing by embedding with an informed decoder, bit stealing by successive degradation (which does not have the transcoder output codebook constraint), and successive refinement (informed encoding).

Finally, for this choice of test channel, we also derive the overall transcoder distortion achieved

$$\begin{aligned}
 d &= \Pr[x \neq u] = \sum_u \Pr[x \neq u | u = u] p_u(u) \\
 &= \sum_u \{ \Pr[x = \hat{x}^\circ, \hat{x}^\circ \neq u | u = u] \\
 &\quad + \Pr[x \neq \hat{x}^\circ, \hat{x}^\circ = u | u = u] \} p_u(u) \\
 &= \sum_u \{ \Pr[x = \hat{x}^\circ] \Pr[\hat{x}^\circ \neq u | u = u] \\
 &\quad + \Pr[x \neq \hat{x}^\circ] \Pr[\hat{x}^\circ = u | u = u] \} p_u(u) \\
 &= [(1 - d_o)\alpha + d_o(1 - \alpha)]q \\
 &\quad + [(1 - d_o)\beta + d_o(1 - \beta)](1 - q) \\
 &= d_o + (1 - 2d_o)[q\alpha + (1 - q)\beta] \\
 &= d_o + 2(1 - 2d_o)q\alpha.
 \end{aligned}$$

The simplification in the last line follows from using $\beta = \alpha q / (1 - q)$. From this result we derive the following value for α :

$$\alpha = \frac{d - d_o}{2q(1 - 2d_o)}.$$

Substituting the values for α , β , and q into (59), and the result into (56), we obtain

$$\begin{aligned}
 R_{IE}(d) &= \max \left\{ R_o - H_B \left(\frac{p - d_o}{1 - 2d_o} \right) \right. \\
 &\quad + \frac{p - d_o}{1 - 2d_o} H_B \left(\frac{d - d_o}{2(p - d_o)} \right) \\
 &\quad \left. + \frac{1 - p - d_o}{1 - 2d_o} H_B \left(\frac{d - d_o}{2(1 - p - d_o)} \right), 0 \right\}. \quad (61)
 \end{aligned}$$

To evaluate the relative efficiencies of rate splitting and embedding, compare $R_{SD}(d)$ in (24) with $R_o - R_{IE}(d)$ in (61).

The latter is the residual source coding rate in information embedding. By the convexity of entropy and applying Jensen's inequality, we can upper-bound the latter two terms in (61) by $H_B((d - d_o)/(1 - 2d_o))$. This means that in general the embedding strategy is less efficient than rate splitting because $R_{SD}(d) \leq R_o - R_{IE}(d)$. It therefore takes a higher residual rate to describe the source to a target distortion level: the transcoder output codebook constraint exacts a price in performance. The gap for the case $p = 2/5$ and $R_o = 1/4$ (for which $d_o \simeq 0.2$) is depicted in Fig. 7 over the relevant distortion range: $d_o \leq d \leq p$. Note the step discontinuity at $d = p$: when stealing all the rate, so that $R_o - R_{IE}(p) = 0$, it suffices for the decoder to ignore all the received data and reconstruct the all-zero sequence.

Note that in the special case $p = 1/2$ (whence $q = 1/2$), no loss of performance is realized, i.e., the transcoder codebook constraint is not limiting: $R_{SD}(d) + R_{IE}(d) = R_o$. We now consider another important situation in which this property holds.

B. Gaussian-Quadratic Case

In this subsection, we show not only that constraining the transcoder codebook to be the same as its input need not incur a loss in performance vis-à-vis rate splitting in the Gaussian-quadratic case, but also that the only additional decoder processing required to ensure there is no loss is simple distortion compensation in the form of (embedding-rate-dependent) attenuation of the source reconstruction. We further show that even an uninformed decoder suffers no more than a 3-dB distortion penalty—or equivalently a 0.21-bit rate penalty—relative to the informed decoder with distortion compensation.

As in Section VI, we let \mathbf{x} be a length- n i.i.d. sequence of Gaussian random variables with variance σ_x^2 and consider the quadratic distortion measure $d(a, b) = (a - b)^2$. In this case,

we obtain that the distortion rate function (57) for informed decoders when the embedding rate is positive, $r > 0$, takes the form

$$\frac{d_{\text{IE}}(r)}{\sigma_x^2} = 2^{-2R_o} + (1 - 2^{-2R_o})2^{-2(R_o-r)} \quad (62)$$

and is achieved using distortion compensation of the form

$$\hat{x} = f(u) = \beta u \quad (63)$$

where

$$\beta = \sqrt{1 - 2^{-2(R_o-r)}}. \quad (64)$$

In contrast, the corresponding distortion rate function for uninformed decoders takes the form

$$\frac{d_{\text{IE}}^u(r)}{\sigma_x^2} = 1 + (1 - 2^{-2R_o})(1 - 2\sqrt{1 - 2^{-2(R_o-r)}}). \quad (65)$$

Equations (62) and (65) are obtained from (56) and (57) as follows. First, as in Section VI-A we generate the usual conditional distribution for achieving the rate–distortion bound for Gaussian sources according to $\hat{x}^\circ = \alpha x + e_o$, where α is given by (27), i.e., $\alpha = (1 - d_o/\sigma_x^2)$, and where e_o is a zero-mean Gaussian random variable with variance αd_o that is independent of x . Using [19, Lemma A.3], we know that the optimizing distribution in (56) is Gaussian. When we further constrain the distribution so that $p_u(x) = p_{\hat{x}^\circ}(x)$, we obtain that it must be of the form (cf. [24]) $u = \gamma \hat{x}^\circ + e$, where $\gamma \geq 0$ is a parameter and e is a zero-mean Gaussian random variable with variance $(1 - \gamma^2)(\sigma_x^2 - d_o)$ that is independent of \hat{x}° .

Next, it is straightforward to confirm that the optimum distortion compensation $f(\cdot)$ must be the MMSE estimator for \mathbf{x} from u . In turn, since we have concluded these are jointly Gaussian random variables, this estimator is linear, whence (63).

It remains only to optimize over the remaining parameters γ and β . In terms of our parameterized distribution, we have

$$R_o - r = I(\hat{x}^\circ; u) = \frac{1}{2} \log \left(\frac{1}{1 - \gamma^2} \right)$$

whence

$$\gamma = \sqrt{1 - 2^{-2(R_o-r)}}. \quad (66)$$

Thus, the distortion takes the form

$$\begin{aligned} d &= E[(\hat{x} - x)^2] \\ &= (\alpha\beta\gamma - 1)^2\sigma_x^2 + \alpha\beta^2\gamma^2d_o + \beta^2(1 - \gamma^2)(\sigma_x^2 - d_o) \\ &= \sigma_x^2 \left[\beta^2 \left(1 - \frac{d_o}{\sigma_x^2}\right) - 2\beta \left(1 - \frac{d_o}{\sigma_x^2}\right) \sqrt{1 - \frac{d_o}{\sigma_x^2} 2^{2r}} + 1 \right]. \end{aligned} \quad (67)$$

where via (26) we have $d_o = \sigma_x^2 2^{-2R_o}$, and where to obtain (67) we have substituted for α and γ according to (27) and (66), respectively.




Fig. 8. Comparing the rate–distortion function for bit stealing via information embedding. The informed and uninformed decoder performances are depicted by the solid and dashed curves, respectively. Below the distortion thresholds d_* and d_*^u for the informed and uninformed decoders, respectively, no embedding should be used.

For uninformed decoders, it suffices to substitute $\beta = 1$ into (67) to obtain (65). For informed decoders, simple optimization of the quadratic (67) with respect to β yields (62) with β given by (64).

The corresponding rate–distortion functions are readily obtained from (62) and (65), and take the forms, respectively

$$R_{\text{IE}}(d) = \max \left\{ R_o - \frac{1}{2} \log \left(\frac{\sigma_x^2 - d_o}{d - d_o} \right), 0 \right\} \quad (68)$$

and

$$R_{\text{IE}}^u(d) = \max \left\{ R_o + \frac{1}{2} \log \left[1 - \frac{1}{4} \left(1 + \frac{\sigma_x^2 - d}{\sigma_x^2 - d_o} \right)^2 \right], 0 \right\}. \quad (69)$$

Comparing (25) and (68), we see that (58) holds with equality in the informed Gaussian case: the transcoder codebook constraint does not exact a price in performance provided the decoder is informed.

The embedding rate–distortion functions (68) and (69) take the form depicted in Fig. 8. As with bit stealing by successive degradation, there is a distortion threshold below which no embedding can be performed and still meet the distortion constraint. One can view this threshold as the minimum amount of distortion that must be incurred if any embedding is used.⁷ For informed decoders, this is given by (30) since successive degradation and embedding have identical performance characteristics in this case. For uninformed decoders, the threshold is found by setting $r = 0$ in (65), revealing that the point of discontinuity between positive and zero-rate embedding is

$$\frac{d_*^u}{\sigma_x^2} = 1 + (1 - 2^{-2R_o})(1 - 2\sqrt{1 - 2^{-2R_o}}). \quad (70)$$

We quantify the loss in performance suffered by an uninformed decoder relative to an informed one in terms of both distortion and rate. We look first at the large r regime, where

⁷This threshold is strictly positive because in order to embed the transcoder must replace codewords it receives with another codewords in the codebook and these codewords have some average minimum distance—in fact, $2d_o$ in the high rate limit—from one another.




Fig. 9. Distortion–rate tradeoffs for bit stealing in the Gaussian-quadratic case, with $R_o = 1$. The progressively lower dash-dotted, solid, and dashed curves correspond to $d_{\text{IE}}^u(R_o - R)$, $d_{\text{SD}}(R) = d_{\text{IE}}(R_o - R)$, and $d_{\text{SC}}(R)$, respectively, i.e., using information embedding without an informed decoder, using information embedding with an informed decoder or using successive degradation, and using successive refinement (informed encoding).

comparing (62) and (65), and consistent with Fig. 8 we see that the loss is largest. Accordingly, we obtain

$$\frac{d_{\text{IE}}^u(r)}{d_{\text{IE}}(r)} \leq \frac{d_{\text{IE}}^u(R_o)}{d_{\text{IE}}(R_o)} = 2 - 2^{-2R_o} < 2$$

which corresponds to a 3-dB gap in the large R_o limit.⁸ The corresponding rate loss comes from comparing (68) and (69), where we see the loss is again largest when r is largest. Thus,

$$\begin{aligned} R_{\text{IE}}(d) - R_{\text{IE}}^u(d) &\leq R_{\text{IE}}(\sigma_x^2) - R_{\text{IE}}^u(\sigma_x^2) \\ &= \frac{1}{2} \log \frac{4}{3} \approx 0.21 \text{ bit} \quad (71) \end{aligned}$$

which we note is independent of R_o .

Turning next to the performance losses in the small r regime, it is straightforward to verify the distortion gap $d_{\text{IE}}^u(r)/d_{\text{IE}}(r)$ is small as $r \rightarrow 0$. Indeed, it is at most 0.2834 dB, which occurs for $R_o \approx 0.2528$. At the two extremes in this small stolen rate regime, $R_o = 0$ and $R_o \rightarrow \infty$, the distortion gap is zero.

To compute the associated rate gap $R_{\text{IE}}(d) - R_{\text{IE}}^u(d)$ we begin by noting that there exists a rate threshold for R_o below which $d_*^u > \sigma_x^2$. In this region $R_{\text{IE}}^u(d) = 0$. In particular, it is straightforward to verify from (70) that $d_*^u > \sigma_x^2$ whenever $0 < R_o < R_o^*$, where [cf. (71)]

$$R_o^* = \frac{1}{2} \log \frac{4}{3} \approx 0.21. \quad (72)$$

Thus, the gap in the small stolen rate regime below this threshold is $R_{\text{IE}}(d) - R_{\text{IE}}^u(d) = R_{\text{IE}}(d) \leq R_{\text{IE}}(\sigma_x^2) = R_o$, which de-

⁸This gap arises because with $r = R_o$, the source is completely overwritten by the embedded message. An informed decoder will ignore the received source codeword, reproduce $\hat{x} = E[x]$ and experience distortion σ_x^2 . However, an uninformed decoder does not know to ignore what it receives, so it experiences an additional distortion of $\sigma_x^2 - d_o$, the variance of the received codeword.

creases monotonically to zero as R_o decreases from R_o^* to 0. Above $R_o = R_o^*$, the rate gap is largest and equals

$$R_{\text{IE}}(d_*^u) = \frac{1}{2} \log \left[2^{1+2R_o} (1 - \sqrt{1 - 2^{-2R_o}}) \right]$$

which decreases monotonically to zero with increasing R_o . Hence, regardless of R_o , the rate gap is, again, at most R_o^* or 0.21 bit.

The different rate–distortion tradeoffs for bit stealing in the Gaussian-quadratic case are summarized in Fig. 9, where normalized distortion is plotted versus residual source coding rate $R = R_o - r$, for $R_o = 1$. The common distortion for successive degradation and embedding with an informed decoder, i.e., $d_{\text{IE}}(R_o - R) = d_{\text{SD}}(R)$, appears as the solid middle curve. The distortion for embedding with an uninformed decoder, i.e., $d_{\text{IE}}^u(R_o - R)$, appears as the dash-dotted upper curve. Finally, the distortion for successive refinement, which corresponds to an informed encoder, is the regular Gaussian-quadratic distortion rate function $d_{\text{SC}}(R)$ and appears as the dashed lower curve.

IX. CONCLUSION

In this paper, we show a variety of results on transcoding. In particular, we show that transcoding with uninformed encoders and/or decoders need not incur significant losses relative to their informed counterpart for some meaningful models, at least if not applied repeatedly. In some sense, this means that for some sources, almost all good source codes have an almost nested structure even when not designed to have one. Thus, it may make sense for system designers to avoid imposing the nested constraint in code design and use the associated degrees of freedom in other ways. Of course, while we have argued that transcoder performance need not significantly suffer,

the complexity of transcoding may be increased. More generally, this raises interesting questions for future work regarding which transcoding complexity–performance tradeoffs are most attractive in different scenarios.

We further show that information embedding approaches to bit stealing that allow for both informed and uninformed source decoders can also be quite efficient. In particular, for reasonable models, they do as well as bit stealing approaches that are not so constrained. When a large fraction of the bits are being stolen, uninformed decoders can incur a substantial performance loss relative to informed decoders. However, in the important Gaussian-quadratic case, informed decoders differ only by incorporating distortion compensation in the form of simple post-reconstruction scaling. Thus, for many audiovisual-oriented sources of practical interest, and the associated gain-invariant distortion measures arising out of human perceptual characteristics, the uninformed and informed source decoder outputs are equivalently good. In this case, the price for enabling uninformed source decoders is increased complexity for extracting the stolen bits, at least relative to successive degradation.

APPENDIX I

A. Proof of Lemma 1

Before proceeding, we need the following lemma.

Lemma 2: For any event \mathcal{B} with $\Pr[\mathcal{B}] > 1/2$

$$|\Pr[\mathcal{A}|\mathcal{B}] - \Pr[\mathcal{A}]| < 2(1 - \Pr[\mathcal{B}]). \quad (73)$$

Proof: We begin by upper-bounding $\Pr[\mathcal{A}|\mathcal{B}]$ via

$$\begin{aligned} \Pr[\mathcal{A}|\mathcal{B}] &= \frac{\Pr[\mathcal{A} \cap \mathcal{B}]}{\Pr[\mathcal{B}]} \leq \frac{\Pr[\mathcal{A}]}{\Pr[\mathcal{B}]} \\ &< \Pr[\mathcal{A}] \cdot (3 - 2\Pr[\mathcal{B}]) \leq \Pr[\mathcal{A}] + 2(1 - \Pr[\mathcal{B}]) \end{aligned} \quad (74)$$

where (74) follows since $1/\Pr[\mathcal{B}] < 3 - 2\Pr[\mathcal{B}]$ if $\Pr[\mathcal{B}] > 1/2$.

Next, we lower-bound $\Pr[\mathcal{A}|\mathcal{B}]$ via

$$\begin{aligned} \Pr[\mathcal{A}|\mathcal{B}] &= \frac{\Pr[\mathcal{A}] + \Pr[\mathcal{B}] - \Pr[\mathcal{A} \cup \mathcal{B}]}{\Pr[\mathcal{B}]} \\ &\geq \Pr[\mathcal{A}] + \Pr[\mathcal{B}] - \Pr[\mathcal{A} \cup \mathcal{B}] \\ &> \Pr[\mathcal{A}] - (1 - \Pr[\mathcal{B}]) \geq \Pr[\mathcal{A}] - 2(1 - \Pr[\mathcal{B}]). \end{aligned}$$

Combining these upper and lower bounds establishes (73) \square

Returning to our proof of Lemma 1, let

$$\begin{aligned} \mathcal{G}_\delta &\triangleq \{(\mathbf{a}, \mathbf{b}) : D(T_{\mathbf{a}, \mathbf{b}} \| p_{\mathbf{v}, \mathbf{u}}) \leq \delta, \\ &\quad \text{and } T_{\mathbf{a}, \mathbf{b}} = T_{\mathbf{v}(j), \mathbf{u}} \text{ for some } j\} \end{aligned}$$

where $\mathbf{v}(j)$ is the j th codeword in the quantization codebook, for $j = 1, 2, \dots, 2^{nR^\delta}$. Then

$$\begin{aligned} &|E\{\theta[\mathbf{v}, \mathbf{u}, g(\mathbf{v})]\} - E\{\theta[\mathbf{w}, \mathbf{u}, g(\mathbf{w})]\}| \\ &= |\Pr\{\theta[\mathbf{v}, \mathbf{u}, g(\mathbf{v})] = 1\} - \Pr\{\theta[\mathbf{w}, \mathbf{u}, g(\mathbf{w})] = 1\}| \end{aligned} \quad (75)$$

$$\begin{aligned} &\leq |\Pr\{\theta[\mathbf{v}, \mathbf{u}, g(\mathbf{v})] = 1\} - \Pr\{\theta[\mathbf{v}, \mathbf{u}, g(\mathbf{v})] = 1|\mathcal{E}^c\}| \\ &\quad + |\Pr\{\theta[\mathbf{w}, \mathbf{u}, g(\mathbf{w})] = 1|(\mathbf{w}, \mathbf{u}) \in \mathcal{G}_\delta\} \\ &\quad - \Pr\{\theta[\mathbf{w}, \mathbf{u}, g(\mathbf{w})] = 1\}| \\ &\quad + |\Pr\{\theta[\mathbf{v}, \mathbf{u}, g(\mathbf{v})] = 1|\mathcal{E}^c\} \\ &\quad - \Pr\{\theta[\mathbf{w}, \mathbf{u}, g(\mathbf{w})] = 1|(\mathbf{w}, \mathbf{u}) \in \mathcal{G}_\delta\}| \quad (76) \\ &\leq 2\Pr[\mathcal{E}] + 2\Pr[\mathcal{E}] + 0 \quad (77) \end{aligned}$$

where (75) follows from the fact that $\theta \in \{0, 1\}$, where (76) follows from two applications of the triangle inequality, and where the first and second terms in (77) come from applications of Lemma 2. Note that to obtain the second term we have used that, in accordance with the dithered quantization rule, $\Pr[(\mathbf{w}, \mathbf{u}) \in \mathcal{G}_\delta] \geq 1 - \Pr[\mathcal{E}]$. To see that the third term in (77) is zero, note that

$$\begin{aligned} &\Pr\{\theta[\mathbf{v}, \mathbf{u}, g(\mathbf{v})] = 1|\mathcal{E}^c\} \\ &\quad - \Pr\{\theta[\mathbf{w}, \mathbf{u}, g(\mathbf{w})] = 1|(\mathbf{w}, \mathbf{u}) \in \mathcal{G}_\delta\} \\ &= \sum_{\mathbf{a}, \mathbf{b}, \mathbf{c}} \theta[\mathbf{a}, \mathbf{b}, \mathbf{c}] [p_{\mathbf{v}, \mathbf{u}, g(\mathbf{v})|\mathcal{E}^c}(\mathbf{a}, \mathbf{b}, \mathbf{c}) \\ &\quad - p_{\mathbf{w}, \mathbf{u}, g(\mathbf{w})|(\mathbf{w}, \mathbf{u}) \in \mathcal{G}_\delta}(\mathbf{a}, \mathbf{b}, \mathbf{c})] \\ &= \sum_{\mathbf{a}, \mathbf{b}, \mathbf{c}} \theta[\mathbf{a}, \mathbf{b}, \mathbf{c}] p_{g(\mathbf{v})|\mathbf{v}}(\mathbf{c}|\mathbf{a}) [p_{\mathbf{v}, \mathbf{u}|\mathcal{E}^c}(\mathbf{a}, \mathbf{b}) \\ &\quad - p_{\mathbf{w}, \mathbf{u}|(\mathbf{w}, \mathbf{u}) \in \mathcal{G}_\delta}(\mathbf{a}, \mathbf{b})] \end{aligned} \quad (78)$$

where to obtain (78) we have used the map independence property, which yields the Markov relationships $\mathbf{u} \leftrightarrow \mathbf{v} \leftrightarrow g(\mathbf{v})$ and $\mathbf{u} \leftrightarrow \mathbf{w} \leftrightarrow g(\mathbf{w})$. Finally, that the term in brackets in (78) is zero follows immediately from the way that \mathbf{v} is generated from a noisy observation \mathbf{w} in the dithered quantization rule when the quantization is successful.

B. Proof of Proposition 1

If \mathbf{y} is the output of the dithering (14) when \mathbf{x} is the input, and $\hat{\mathbf{x}}^\circ$ is the codeword to which \mathbf{y} is mapped, then when encoding succeeds we have

$$\begin{aligned} d(\mathbf{x}, \hat{\mathbf{x}}^\circ) &= d(\mathbf{x}, \mathbf{y}) \\ &= \sum_{\hat{\mathbf{x}}^\circ, x} d(x, \hat{\mathbf{x}}^\circ) T_{\mathbf{y}, \mathbf{x}}(\hat{\mathbf{x}}^\circ, x) \\ &= E[d(x, \hat{\mathbf{x}}^\circ)] \\ &\quad + \sum_{\hat{\mathbf{x}}^\circ, x} d(x, \hat{\mathbf{x}}^\circ) [T_{\mathbf{y}, \mathbf{x}}(\hat{\mathbf{x}}^\circ, x) - p_{\hat{\mathbf{x}}^\circ, \mathbf{x}}(\hat{\mathbf{x}}^\circ, x)] \\ &\leq E[d(x, \hat{\mathbf{x}}^\circ)] \\ &\quad + d_{\max} \sum_{\hat{\mathbf{x}}^\circ, x} |T_{\mathbf{y}, \mathbf{x}}(\hat{\mathbf{x}}^\circ, x) - p_{\hat{\mathbf{x}}^\circ, \mathbf{x}}(\hat{\mathbf{x}}^\circ, x)| \quad (79) \end{aligned}$$

$$\begin{aligned} &\leq E[d(x, \hat{\mathbf{x}}^\circ)] \\ &\quad + d_{\max} \sqrt{2 \ln 2 \cdot D(T_{\mathbf{y}, \mathbf{x}} \| p_{\hat{\mathbf{x}}^\circ, \mathbf{x}})} \quad (80) \end{aligned}$$

$$\leq E[d(x, \hat{\mathbf{x}}^\circ)] + d_{\max} \sqrt{2\delta \ln 2} \quad (81)$$

where (79) follows since $d(x, \hat{\mathbf{x}}^\circ) \leq d_{\max}$, (80) follows, like (12), from [11, p. 300, Lemma 12.6.1], and (81) is a consequence of dithered quantization step 2. The remaining steps follow from simple algebraic manipulations and the definition of types.

Finally, we obtain

$$\begin{aligned} & \Pr \left[d(\mathbf{x}, \hat{\mathbf{x}}^\circ) > E[d(x, \hat{x}^\circ)] + d_{\max} \sqrt{2\delta \ln 2} \right] \\ &= \Pr \left[d(\mathbf{x}, \hat{\mathbf{x}}^\circ) > E[d(x, \hat{x}^\circ)] + d_{\max} \sqrt{2\delta \ln 2} \mid \mathcal{E} \right] \Pr[\mathcal{E}] \\ & \quad + \Pr \left[d(\mathbf{x}, \hat{\mathbf{x}}^\circ) > E[d(x, \hat{x}^\circ)] + d_{\max} \sqrt{2\delta \ln 2} \mid \mathcal{E}^c \right] \Pr[\mathcal{E}^c] \\ & \leq 1 \cdot \Pr[\mathcal{E}] + 0 \cdot \Pr[\mathcal{E}^c] \end{aligned}$$

where to obtain the zero in the last line we have used (81), yielding (15) as desired.

C. Proof of Proposition 2

First, we need the following lemma, which establishes that for any pair of sequences (\mathbf{y}, \mathbf{x}) satisfying encoding step 2, the empirical type $T_{\mathbf{y}, \mathbf{x}}$ is about as close to $p_{\hat{\mathbf{x}}^\circ}^\circ p_{\mathbf{x}}$ as $p_{\hat{\mathbf{x}}^\circ, \mathbf{x}}$ is to $p_{\hat{\mathbf{x}}^\circ}^\circ p_{\mathbf{x}}$.

Lemma 3: For any empirical joint type $T_{\mathbf{y}, \mathbf{x}}$ where (\mathbf{y}, \mathbf{x}) satisfies the condition in dithered quantization step 2

$$D(T_{\mathbf{y}, \mathbf{x}} \| p_{\hat{\mathbf{x}}^\circ}^\circ p_{\mathbf{x}}) \leq I(\hat{\mathbf{x}}^\circ; \mathbf{x}) + \delta - \log \left[(p_{\hat{\mathbf{x}}^\circ, \mathbf{x}}^{\min})^3 \right] \sqrt{2\delta \ln 2} \quad (82)$$

where the superscript \min denotes the smallest nonzero value of the distribution that is its argument, i.e., for an arbitrary distribution p_w

$$p_w^{\min} = \min_{\{w | p_w(w) > 0\}} p_w(w)$$

Proof: Equation (82) is obtained via

$$\begin{aligned} & D(T_{\mathbf{y}, \mathbf{x}} \| p_{\hat{\mathbf{x}}^\circ}^\circ p_{\mathbf{x}}) = I(\hat{\mathbf{x}}^\circ; \mathbf{x}) \\ & \quad + D(T_{\mathbf{y}, \mathbf{x}} \| p_{\hat{\mathbf{x}}^\circ}^\circ p_{\mathbf{x}}) - D(p_{\hat{\mathbf{x}}^\circ, \mathbf{x}} \| p_{\hat{\mathbf{x}}^\circ}^\circ p_{\mathbf{x}}) \\ &= I(\hat{\mathbf{x}}^\circ; \mathbf{x}) + H(p_{\hat{\mathbf{x}}^\circ, \mathbf{x}}) - H(T_{\mathbf{y}, \mathbf{x}}) \\ & \quad - \sum_{\hat{x}^\circ, x} [T_{\mathbf{y}, \mathbf{x}}(\hat{x}^\circ, x) - p_{\hat{\mathbf{x}}^\circ, \mathbf{x}}(\hat{x}^\circ, x)] \log [p_{\hat{\mathbf{x}}^\circ}^\circ(\hat{x}^\circ) p_{\mathbf{x}}(x)] \\ & \leq I(\hat{\mathbf{x}}^\circ; \mathbf{x}) + |H(T_{\mathbf{y}, \mathbf{x}}) - H(p_{\hat{\mathbf{x}}^\circ, \mathbf{x}})| \\ & \quad - \sum_{\hat{x}^\circ, x} |T_{\mathbf{y}, \mathbf{x}}(\hat{x}^\circ, x) - p_{\hat{\mathbf{x}}^\circ, \mathbf{x}}(\hat{x}^\circ, x)| \log [p_{\hat{\mathbf{x}}^\circ}^\circ(\hat{x}^\circ) p_{\mathbf{x}}(x)] \quad (83) \\ & \leq I(\hat{\mathbf{x}}^\circ; \mathbf{x}) + D(T_{\mathbf{y}, \mathbf{x}} \| p_{\hat{\mathbf{x}}^\circ, \mathbf{x}}) \\ & \quad - \sum_{\hat{x}^\circ, x} |T_{\mathbf{y}, \mathbf{x}}(\hat{x}^\circ, x) - p_{\hat{\mathbf{x}}^\circ, \mathbf{x}}(\hat{x}^\circ, x)| \\ & \quad \quad \log [p_{\hat{\mathbf{x}}^\circ}^\circ(\hat{x}^\circ) p_{\mathbf{x}}(x) p_{\hat{\mathbf{x}}^\circ, \mathbf{x}}(\hat{x}^\circ, x)] \quad (84) \\ & \leq I(\hat{\mathbf{x}}^\circ; \mathbf{x}) + \delta - \log \left[(p_{\hat{\mathbf{x}}^\circ}^\circ p_{\mathbf{x}})^{\min} p_{\hat{\mathbf{x}}^\circ, \mathbf{x}}^{\min} \right] \\ & \quad - \sum_{\hat{x}^\circ, x} |T_{\mathbf{y}, \mathbf{x}}(\hat{x}^\circ, x) - p_{\hat{\mathbf{x}}^\circ, \mathbf{x}}(\hat{x}^\circ, x)| \quad (85) \end{aligned}$$

$$\leq I(\hat{\mathbf{x}}^\circ; \mathbf{x}) + \delta - \log \left[(p_{\hat{\mathbf{x}}^\circ, \mathbf{x}}^{\min})^3 \right] \sqrt{D(T_{\mathbf{y}, \mathbf{x}} \| p_{\hat{\mathbf{x}}^\circ, \mathbf{x}}) 2 \ln 2} \quad (86)$$

$$\leq I(\hat{\mathbf{x}}^\circ; \mathbf{x}) + \delta - \log \left[(p_{\hat{\mathbf{x}}^\circ, \mathbf{x}}^{\min})^3 \right] \sqrt{2\delta \ln 2} \quad (87)$$

where (83) follows from fact that $\log p_{\hat{\mathbf{x}}^\circ}^\circ p_{\mathbf{x}} \leq 0$, (84) follows from the inequality

$$\begin{aligned} & |H(T_{\mathbf{y}, \mathbf{x}}) - H(p_{\hat{\mathbf{x}}^\circ, \mathbf{x}})| = \left| D(T_{\mathbf{y}, \mathbf{x}} \| p_{\hat{\mathbf{x}}^\circ, \mathbf{x}}) \right. \\ & \quad \left. + \sum_{\hat{x}^\circ, x} [T_{\mathbf{y}, \mathbf{x}}(\hat{x}^\circ, x) - p_{\hat{\mathbf{x}}^\circ, \mathbf{x}}(\hat{x}^\circ, x)] \log [p_{\hat{\mathbf{x}}^\circ, \mathbf{x}}(\hat{x}^\circ, x)] \right| \\ & \leq D(T_{\mathbf{y}, \mathbf{x}} \| p_{\hat{\mathbf{x}}^\circ, \mathbf{x}}) \\ & \quad - \sum_{\hat{x}^\circ, x} |T_{\mathbf{y}, \mathbf{x}}(\hat{x}^\circ, x) - p_{\hat{\mathbf{x}}^\circ, \mathbf{x}}(\hat{x}^\circ, x)| \log [p_{\hat{\mathbf{x}}^\circ, \mathbf{x}}(\hat{x}^\circ, x)] \end{aligned}$$

and (85) follows from successful quantization in encoding step 2 and the following argument. If for some (\hat{x}°, x) the marginal product satisfies $p_{\hat{\mathbf{x}}^\circ, \mathbf{x}}(\hat{x}^\circ) p_{\mathbf{x}}(x) = 0$ then

$$p_{\hat{\mathbf{x}}^\circ, \mathbf{x}}(\hat{x}^\circ, x) = T_{\mathbf{y}, \mathbf{x}}(\hat{x}^\circ, x) = 0.$$

This is because $p_{\hat{\mathbf{x}}^\circ, \mathbf{x}}$ must have at least the same zeros as $p_{\hat{\mathbf{x}}^\circ}^\circ p_{\mathbf{x}}$. In this case, $T_{\mathbf{y}, \mathbf{x}}(\hat{x}^\circ, x) = 0$ since $D(T_{\mathbf{y}, \mathbf{x}} \| p_{\hat{\mathbf{x}}^\circ, \mathbf{x}})$ cannot be infinite because we know the dithered encoder succeeded in step 2. Hence, the largest term in the sum in (84) is $\log[(p_{\hat{\mathbf{x}}^\circ}^\circ p_{\mathbf{x}})^{\min} p_{\hat{\mathbf{x}}^\circ, \mathbf{x}}^{\min}]$, and (85) follows. The remaining two inequalities (86) and (87) follow from analogous arguments in the proof of Proposition 1, together with the fact that $(p_{\hat{\mathbf{x}}^\circ}^\circ p_{\mathbf{x}})^{\min} > (p_{\hat{\mathbf{x}}^\circ, \mathbf{x}}^{\min})^2$. \square

Next we need a lemma lower-bounding the probability that the encoder could encode to the i th codeword.

Lemma 4: For any empirical joint type $T_{\mathbf{y}, \mathbf{x}}$ (and in particular the ones where (\mathbf{y}, \mathbf{x}) satisfies the condition in the dithered encoding step 2)

$$\Pr[T_{\hat{\mathbf{x}}^\circ(i), \mathbf{x}} = T_{\mathbf{y}, \mathbf{x}}] \geq 2^{-n \left[D(T_{\mathbf{y}, \mathbf{x}} \| p_{\hat{\mathbf{x}}^\circ}^\circ p_{\mathbf{x}}) + |\mathcal{X}|^2 \cdot \frac{\log(n+1)}{n} \right]}. \quad (88)$$

Proof: The desired result follows from the chain of inequalities

$$\begin{aligned} & \Pr[T_{\hat{\mathbf{x}}^\circ(i), \mathbf{x}} = T_{\mathbf{y}, \mathbf{x}}] = \Pr[T_{\hat{\mathbf{x}}^\circ(i), \mathbf{x}} = T_{\mathbf{y}, \mathbf{x}} | \mathbf{x} = \mathbf{x}] \quad (89) \\ &= \frac{\Pr[T_{\hat{\mathbf{x}}^\circ(i), \mathbf{x}} = T_{\mathbf{y}, \mathbf{x}}] \cdot \Pr[\mathbf{x} = \mathbf{x} | T_{\hat{\mathbf{x}}^\circ(i), \mathbf{x}} = T_{\mathbf{y}, \mathbf{x}}]}{\Pr[\mathbf{x} = \mathbf{x}]} \end{aligned}$$

$$= \frac{\Pr[T_{\hat{\mathbf{x}}^\circ(i), \mathbf{x}} = T_{\mathbf{y}, \mathbf{x}}] \cdot 2^{-n \cdot H(T_{\mathbf{x}})}}{2^{-n(H(T_{\mathbf{x}}) + D(T_{\mathbf{x}} \| p_{\hat{\mathbf{x}}^\circ}^\circ))}} \quad (90)$$

$$\geq \Pr[T_{\hat{\mathbf{x}}^\circ(i), \mathbf{x}} = T_{\mathbf{y}, \mathbf{x}}] \quad (91)$$

$$\geq (n+1)^{-|\mathcal{X}|^2} \cdot 2^{-n D(T_{\mathbf{y}, \mathbf{x}} \| p_{\hat{\mathbf{x}}^\circ}^\circ p_{\mathbf{x}})}. \quad (92)$$

Equation (89) follows from the fact that $\hat{\mathbf{x}}^\circ$ and \mathbf{x} are independent since the codewords are generated independently of the source, and (90) follows from the definition of conditional probability. Equation (91) follows from [11, p. 281, Theorem 12.1.2], and the observation that $T_{\hat{\mathbf{x}}^\circ(i), \mathbf{x}} = T_{\mathbf{y}, \mathbf{x}}$ implies $T_{\mathbf{x}} = T_{\mathbf{x}}$. Equation (92) follows from the fact that relative entropy is nonnegative. Finally, (93) is a consequence of [11, p. 285, Theorem 12.1.4]. \square

Now we are ready to prove Proposition 2. First, we express the dithered quantization failure event \mathcal{E} as

$$\mathcal{E} = \mathcal{E}_D \cup \mathcal{E}_T \quad (94)$$

where \mathcal{E}_D denotes the event that the relative entropy $D(T_{\mathbf{y},\mathbf{x}}||p_{\hat{x}^\circ,\mathbf{x}})$ is too large in step 2 and \mathcal{E}_T denotes the event that no sequences exist in the codebook such that $T_{\hat{x}^\circ,\mathbf{x}} = T_{\mathbf{y},\mathbf{x}}$ in step 3.

From (94), we obtain

$$\Pr[\mathcal{E}] \leq \Pr[\mathcal{E}_D] + \Pr[\mathcal{E}_T|\mathcal{E}_D^c] \quad (95)$$

whose components we now bound in turn.

Via [11, p. 287, Theorem 12.2.1], we obtain

$$\Pr[\mathcal{E}_D] \leq 2^{-n(\delta - |\mathcal{X}|^2 \cdot \frac{\log(n+1)}{n})}.$$

Hence there exists an $\alpha_1 > 0$ and n_1 such that for all $n > n_1$

$$\Pr[\mathcal{E}_D] < \exp(-n\alpha_1). \quad (96)$$

For example, it suffices to choose $\alpha_1 = \delta/2$ and $n_1 = 4|\mathcal{X}|^4/\delta^2$.

The second term in (95) is bounded by the following chain of inequalities:

$$\begin{aligned} & \Pr[\mathcal{E}_T|\mathcal{E}_D^c] \\ &= \sum_{\mathbf{x},\mathbf{y}} \Pr[\mathcal{E}_T|(\mathbf{x},\mathbf{y}),\mathcal{E}_D^c] \Pr[(\mathbf{x},\mathbf{y}) = (\mathbf{x},\mathbf{y})|\mathcal{E}_D^c] \\ &\leq \sum_{\mathbf{x},\mathbf{y}} \max_{\mathbf{x},\mathbf{y}} \{ \Pr[\mathcal{E}_T|(\mathbf{x},\mathbf{y}),\mathcal{E}_D^c] \} \\ &\quad \times \Pr[(\mathbf{x},\mathbf{y}) = (\mathbf{x},\mathbf{y})|\mathcal{E}_D^c] \\ &= \max_{\mathbf{x},\mathbf{y}} \Pr[\mathcal{E}_T|(\mathbf{x},\mathbf{y}),\mathcal{E}_D^c] \\ &= \max_{\mathbf{x},\mathbf{y}} \prod_{i=1}^{2^{nR_o}} (1 - \Pr[T_{\hat{x}^\circ(i),\mathbf{x}} = T_{\mathbf{y},\mathbf{x}}|\mathcal{E}_D^c]) \\ &= \max_{\mathbf{x},\mathbf{y}} (1 - \Pr[T_{\hat{x}^\circ(1),\mathbf{x}} = T_{\mathbf{y},\mathbf{x}}|\mathcal{E}_D^c])^{2^{nR_o}} \quad (97) \end{aligned}$$

$$\leq \max_{\mathbf{x},\mathbf{y}} \left(1 - 2^{-n \left[D(T_{\mathbf{y},\mathbf{x}}||p_{\hat{x}^\circ p_x}) + |\mathcal{X}|^2 \cdot \frac{\log(n+1)}{n} \right]} \right)^{2^{nR_o}} \quad (98)$$

$$\leq \max_{\mathbf{x},\mathbf{y}} \exp \left\{ -2^{-n \left[D(T_{\mathbf{y},\mathbf{x}}||p_{\hat{x}^\circ p_x}) + |\mathcal{X}|^2 \cdot \frac{\log(n+1)}{n} \right]} \cdot 2^{nR_o} \right\} \quad (99)$$

$$= \max_{\mathbf{x},\mathbf{y}} \exp \left\{ -2^n \left[R_o - D(T_{\mathbf{y},\mathbf{x}}||p_{\hat{x}^\circ p_x}) - |\mathcal{X}|^2 \cdot \frac{\log(n+1)}{n} \right] \right\}$$

$$\leq \exp \left\{ -2^n \left[R_o - I(\hat{x}^\circ; \mathbf{x}) - \delta + \log \left[(p_{\hat{x}^\circ, \mathbf{x}}^{\min})^3 \right] \sqrt{2\delta \ln 2} - |\mathcal{X}|^2 \cdot \frac{\log(n+1)}{n} \right] \right\} \quad (100)$$

where (97) follows from symmetry, (98) follows from Lemma 4, (99) follows from the inequality [11, p. 353, Lemma 13.5.3], $(1-y)^n \leq \exp(-yn)$, and (100) follows from Lemma 3.

Let δ_{\max} be the largest value of δ such that

$$R_o > I(\hat{x}^\circ; \mathbf{x}) + \delta - \log \left[(p_{\hat{x}^\circ, \mathbf{x}}^{\min})^3 \right] \sqrt{2\delta \ln 2}.$$

Note that since $R_o > I(\mathbf{x}; \hat{x}^\circ)$ we have $\delta_{\max} > 0$. Hence, for every $\delta \in (0, \delta_{\max})$, there exists an n_2 such that for all $n > n_2$

$$\Pr[\mathcal{E}_T|\mathcal{E}_D^c] \leq \exp(-\alpha_2 n). \quad (101)$$

for some $\alpha_2 > 0$.

Finally, combining the exponential bounds (96) and (101) with (95) we conclude there exists an n_0 such that for all $n > n_0$

$$\Pr[\mathcal{E}] \leq 2 \exp[-n \min(\alpha_1, \alpha_2)].$$

D. Proof of Proposition 3

For our proof, we require the following lemma.

Lemma 5: If \hat{x}° is the output of an admissible encoder and $p_{\hat{x}^\circ}$ is its distribution, then for all n

$$\Pr[D(T_{\hat{x}^\circ}||p_{\hat{x}^\circ}) \leq \epsilon_0(n)] \geq 1 - \epsilon_0(n), \quad (102)$$

where $\epsilon_0(n)$ is a function satisfying $\lim_{n \rightarrow \infty} \epsilon_0(n) = 0$.

Proof: For admissible encoders we know $D(T_{\hat{x}^\circ, \mathbf{x}}||p_{\hat{x}^\circ, \mathbf{x}})$ converges to zero in probability (Definition 2) and therefore so does $D(T_{\hat{x}^\circ}||p_{\hat{x}^\circ})$, i.e., for every $\epsilon > 0$ there exists an $n_0(\epsilon)$ such that for all $n > n_0(\epsilon)$

$$\Pr[D(T_{\hat{x}^\circ}||p_{\hat{x}^\circ}) \leq \epsilon] \geq 1 - \epsilon.$$

Without loss of generality, we can assume that the function is a mapping of the form $n_0(\epsilon) : (0, 1) \mapsto \mathbb{R}$ and monotonically increasing as ϵ decreases. As such, it possesses an inverse that is the desired $\epsilon_0(n)$ in our lemma. \square

For our main result, we use proof by contradiction. Suppose that the probability that the dithered transcoder fails does not converge to 0 with n , but stays above some fixed $\epsilon_1 > 0$. This implies, when combined with Lemma 5, that there is a set of sequences \mathcal{S}_n that cannot be transcoded successfully and such that for all $\hat{x}^\circ \in \mathcal{S}_n$

$$D(T_{\hat{x}^\circ}||p_{\hat{x}^\circ}) \leq \epsilon_0(n). \quad (103)$$

Since by their construction dithered transcoders treat all inputs in a given type identically, the set \mathcal{S}_n must contain some number of whole type classes. Denote these type classes $\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_K$. Let \mathcal{T}_* denote the worst of these type classes, i.e.,

$$\mathcal{T}_* = \arg \max_{\mathcal{T} \in \{\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_K\}} \Pr[\mathcal{E}|\hat{x}^\circ \in \mathcal{T}] \quad (104)$$

where \mathcal{E} denotes the event that the dithered transcoding fails.

Then \mathcal{T}_* satisfies the following:

$$\begin{aligned} \epsilon_1 &= \Pr[\mathcal{E}] = \sum_{i=1}^K \Pr[\mathcal{E}|\hat{x}^\circ \in \mathcal{T}_i] \Pr[\hat{x}^\circ \in \mathcal{T}_i] \\ &\leq \sum_{i=1}^K \Pr[\mathcal{E}|\hat{x}^\circ \in \mathcal{T}_i] \\ &\leq \Pr[\mathcal{E}|\hat{x}^\circ \in \mathcal{T}_*] \cdot (n+1)^{|\mathcal{X}|} \quad (105) \end{aligned}$$

where to obtain the last inequality we have used (104) and that there are at most $(n+1)^{|\mathcal{X}|}$ type classes [11, p. 280, Theorem 12.1.1]. From (105), we obtain

$$\Pr[\mathcal{E}|\hat{\mathbf{x}}^\circ \in \mathcal{T}_*] \geq \frac{\epsilon_1}{(n+1)^{|\mathcal{X}|}}. \quad (106)$$

Now let us further suppose that the transcoder input is generated in an i.i.d. manner according to the distribution $p_{\hat{\mathbf{x}}^\circ}$, which is a valid output of an admissible source encoder. Then [11, p. 285, Theorem 12.1.4]

$$\Pr[\hat{\mathbf{x}}^\circ \in \mathcal{T}_*] \geq \frac{2^{-nD(\mathcal{T}_*||p_{\hat{\mathbf{x}}^\circ)}}}{(n+1)^{|\mathcal{X}|}} \geq \frac{2^{-n\epsilon_0(n)}}{(n+1)^{|\mathcal{X}|}}. \quad (107)$$

where to obtain the second inequality we have used (103).

From (106) and (107), we see that the probability of dithered transcoder failure does not decay exponentially in n , since

$$\Pr[\mathcal{E}] \geq \Pr[\mathcal{E}|\hat{\mathbf{x}}^\circ \in \mathcal{T}_*] \Pr[\hat{\mathbf{x}}^\circ \in \mathcal{T}_*] \geq \frac{\epsilon_1}{(n+1)^2|\mathcal{X}|}$$

where $\epsilon_0(n) \rightarrow 0$ as $n \rightarrow \infty$.

But this contradicts Proposition 2 which states that the probability of a dithered quantization failure must decrease exponentially in n when the quantizer input is i.i.d. Thus, we conclude that the probability of dithered transcoder failure must approach zero as $n \rightarrow \infty$.

ACKNOWLEDGMENT

The authors wish to thank Prof. Ioannis Kontoyiannis and Prof. Ram Zamir for helpful interactions, as well as the Associate Editors and reviewers whose comments and suggestions have greatly aided in the presentation of the ideas.

REFERENCES

- [1] W. H. Equitz and T. M. Cover, "Successive refinement of information," *IEEE Trans. Inf. Theory*, vol. 37, no. 2, pp. 269–275, Mar. 1991.
- [2] D. Karakos and A. Papamarcou, "A relationship between quantization and watermarking rates in the presence of additive Gaussian attacks," *IEEE Trans. Inf. Theory*, vol. 49, no. 8, pp. 1970–1982, Aug. 2003.
- [3] D. Kundur, "Implications for high capacity data hiding in the presence of lossy compression," in *Proc. IEEE Int. Conf. Information Technology: Coding and Computing*, Las Vegas, NV, Mar. 2000, pp. 16–21.
- [4] M. Ramkumar and A. N. Akansu, "Information theoretic bounds for data hiding in compressed images," in *Proc. IEEE Workshop on Multimedia Signal Processing*, Redondo Beach, CA, Dec. 1998, pp. 267–272.
- [5] H. Yamamoto, "Source coding theory for cascade and branching communications systems," *IEEE Trans. Inf. Theory*, vol. IT-27, no. 3, pp. 299–308, May 1981.
- [6] B. Rimoldi, "Successive refinement of information: Characterization of the achievable rates," *IEEE Trans. Inf. Theory*, vol. 40, no. 1, pp. 253–259, Jan. 1994.

- [7] E.-H. Yang and Z. Zhang, "The redundancy of source coding with a fidelity criterion—Part II: Coding at a fixed rate level with unknown statistics," *IEEE Trans. Inf. Theory*, vol. 47, no. 1, pp. 126–145, Jan. 2001.
- [8] R. L. Dobrushin and B. S. Tsybakov, "Information transmission with additional noise," *IEEE Trans. Inf. Theory*, vol. IT-8, no. 5, pp. 293–304, Sep. 1962.
- [9] J. K. Wolf and J. Ziv, "Transmission of noisy information to a noisy receiver with minimum distortion," *IEEE Trans. Inf. Theory*, vol. IT-16, no. 4, pp. 406–411, Jul. 1970.
- [10] I. Csiszár and J. Körner, *Information Theory, Coding Theorems for Discrete Memoryless Systems*. Budapest, Hungary: Akadémiai Kiadó, 1981.
- [11] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York: Wiley, 1991.
- [12] A. Kanlis, S. Khudanpur, and P. Narayan, "Typicality of a good rate-distortion code," *Probl. Inf. Transm.*, vol. 32, no. 1, pp. 112–121, Jan. 1996.
- [13] A. Kanlis, "Compression and Transmission of Information at Multiple Resolutions," Ph.D. dissertation, Univ. Maryland, College Park, MD, 1998.
- [14] S. Shamai (Shitz) and S. Verdú, "The empirical distribution of good codes," *IEEE Trans. Inf. Theory*, vol. 43, no. 3, pp. 836–846, May 1997.
- [15] T. Berger, "Multiterminal source coding," in *The Information Theory Approach to Communications*, G. Longo, Ed. New York: Springer-Verlag, 1977, ch. 4.
- [16] —, *Rate Distortion Theory*. Englewood Cliffs, NJ: Prentice-Hall, 1971.
- [17] T. M. Cover and M. Chiang, "Duality between channel capacity and rate-distortion with two-sided side information," *IEEE Trans. Inf. Theory*, vol. 48, no. 6, pp. 1629–1638, Jun. 2002.
- [18] A. D. Wyner, "The rate-distortion function for source coding with side information at the decoder-II: General sources," *Inf. Contr.*, vol. 38, pp. 60–80, 1978.
- [19] A. S. Cohen and A. Lapidith, "The Gaussian watermarking game," *IEEE Trans. Inf. Theory*, vol. 48, no. 6, pp. 1639–1667, Jun. 2002.
- [20] R. Zamir and M. Feder, "On universal quantization by randomized uniform/lattice quantizers," *IEEE Trans. Inf. Theory*, vol. 38, no. 2, pp. 428–436, Mar. 1992.
- [21] A. Lapidith, "On the role of mismatch in rate-distortion theory," *IEEE Trans. Inf. Theory*, vol. 43, no. 1, pp. 38–47, Jan. 1997.
- [22] L. Lastras and T. Berger, "All sources are nearly successive refinable," *IEEE Trans. Inf. Theory*, vol. 47, no. 3, pp. 918–926, Mar. 2001.
- [23] T. Linder and R. Zamir, "On the asymptotic tightness of the Shannon lower bound," *IEEE Trans. Inf. Theory*, vol. 40, no. 6, pp. 2026–2031, Nov. 1994.
- [24] M. H. Costa, "Writing on dirty paper," *IEEE Trans. Inf. Theory*, vol. IT-29, no. 3, pp. 439–441, May 1983.
- [25] B. Chen and G. W. Wornell, "Quantization index modulation: A class of provably good methods for digital watermarking and information embedding," *IEEE Trans. Inf. Theory*, vol. 47, no. 4, pp. 1423–1443, May 2001.
- [26] R. Zamir, S. Shamai (Shitz), and U. Erez, "Nested codes: An algebraic binning scheme for noisy multiterminal networks," *IEEE Trans. Inf. Theory*, vol. 42, no. 6, pp. 1250–1276, Jun. 2002.
- [27] R. J. Barron, B. Chen, and G. W. Wornell, "The duality between information embedding and source coding with side information and some applications," *IEEE Trans. Inf. Theory*, vol. 49, no. 5, pp. 1159–1180, May 2003.