Local Rewriting under Assumed Erasure

Napassorn Litchiowong

pleng@u.nus.edu

Abstract. We study local rewriting after an assumed erasure step in an ideal oblivious-transfer protocol. Bob transforms his retained record while Alice's actual record remains fixed. Exact rewriting between shared-source and independent-source records is possible in both directions precisely when the retained records of the shared source are independent. For balanced deterministic maps retaining k and ℓ bits from an n-bit source, the optimal error over the maps is $\max\{0, 1 - 2^{n-k-\ell}\}$. For fixed full-row-rank linear maps, it is $1-2^{-d}$, where d is their row-space intersection dimension. A nonlinear example shows that an optimal approximate rewrite may change Bob's marginal distribution.

Keywords: Oblivious transfer · Erasures · Local rewriting
Information-theoretic security · Statistical distance

1 Introduction

The information available after a protocol execution determines which alternative executions a party can plausibly describe. Erasing part of a local record changes this information. To study the effect of erasure, one must specify both the retained state and the transformations allowed on it.

We consider two joint distributions. In the first, Alice and Bob retain functions $f(S)$ and $g(S)$ of the same uniform string. In the second, they retain $f(S)$ and $g(R)$, where S and R are independent. Bob is allowed to apply a randomized transformation to his own record. Alice's actual record remains unchanged. The question is whether the resulting pair has the distribution of retained records in the other experiment, and, when exact equality is impossible, how small the statistical error can be.

This question already has distinct answers for two bits. If $S = (S_1, S_2)$ and the parties retain different coordinates, the records are independent in either experiment. If both retain the first coordinate, the shared-source records agree with probability one, whereas a local transformation of Bob's independent-source record agrees with Alice's bit with probability only $1/2$. Thus the number of retained bits alone does not determine the rewriting error. Figure 1 displays the two joint distributions.

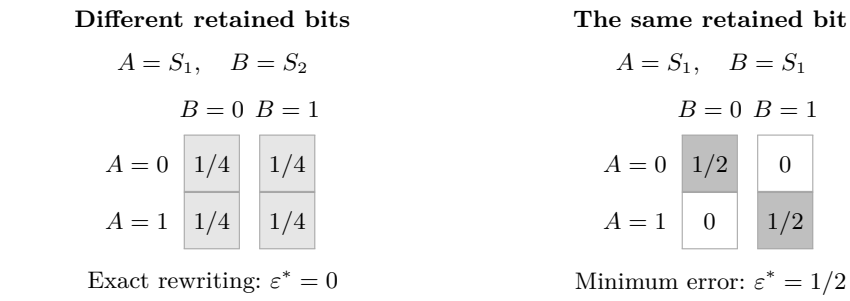

Fig. 1. Two retention rules for a uniform two-bit source. The cells show the shared-source joint distribution; the independent-source distribution is uniform on all four pairs in both cases. Each party retains one bit, but only the left-hand rule permits exact rewriting. Bob cannot create agreement with an independent bit held by Alice using a local transformation alone.

1.1 Results

We first reduce bidirectional rewriting to approximation of the shared-source law by a product distribution with Alice's marginal fixed (Proposition 1). The direction from shared to independent records is exact: Bob can resample his marginal. The reverse direction determines the optimum. In particular, exact rewriting in both directions is equivalent to independence of the retained records.

For balanced deterministic maps, Theorem 2 gives the minimum error over all maps with prescribed output lengths. This bound is attained by coordinate projections. Proposition 3 determines the error for each pair of full-row-rank linear maps in terms of their row-space intersection. The distinction between optimizing the maps and analyzing fixed maps is essential. The nonlinear example in Section 5 shows that changing Bob's claimed marginal can reduce the error from 5/8 to 1/2. Appendix C separates the joint average error from a worst-record guarantee. We also analyze how additional retained state changes the experiment.

The paper is a finite-distribution analysis motivated by erasure. The ideal-OT experiment supplies the two source laws; the results concern local transformations of their retained records. Erasure is assumed to remove all state outside that rule, including external copies and access to the old resource. The model does not provide a mechanism for enforcing deletion. The proofs are distributional; they do not establish a real-protocol simulation theorem or security against parties that deviate from the prescribed erasure.

1.2 Related work

Secure non-interactive reductions provide a general framework for local conversion of correlations. Agarwal et al. [2], Definition 7 and Theorem 3, characterize such reductions by stochastic-map identities. Exact bidirectional rewriting here is the specialization in which Alice's protocol map and simulator are both the identity. Appendix D gives all three identities and the substitution, including the support conventions. This connection concerns exact rewriting. We do not identify our approximate minimax with unrestricted approximate secure non-interactive reduction. The cited full version uses full L^1 distance, whereas we use half-L^1 total variation.

Statistical comparison of experiments also studies approximation by stochastic postprocessing. Raginsky [3], Section III-A, uses total variation to measure such approximation. Our objective averages over Alice's marginal; it need not agree with a maximum over her possible records. For linear images of a uniform source, the independence criterion follows from the rank expression for mutual information of Silva and Kschischang [4], Lemma 6. We use a support-counting argument to obtain the corresponding total-variation error. These connections place the calculations within established correlation-conversion and information-theoretic methods.

Erasure-based security requires a separate model. Hazay, Lindell, and Patra [5], Definitions 3.6–3.7 in the corrected full version, study adaptive corruption with partial erasures. Müller-Quade, Röhrich, and Unruh [6], Lemma 1, study undoable commitments between non-deleting parties, including an actively deviating receiver. Neither experiment coincides with the fixed honest retention rule considered here.

2 Model and Definitions

2.1 Protocol and retained state

We use the coordinatewise form of the construction in [1, Proposition 2], whose complete execution and retention rule we specify below. An ideal bit OT takes a pair (u_0, u_1) from Alice and a choice b from Bob, delivers u_b only to Bob, and gives Alice no output. Fix $n \geq 1$. Alice has an unused input label $x \in \{0,1\}$, and Bob has input $b \in \{0,1\}$. Independently of the inputs, Alice samples $S \leftarrow \mathrm{Unif}(\mathbb{F}_2^n)$ and Bob samples $R \leftarrow \mathrm{Unif}(\mathbb{F}_2^n)$. In each coordinate i, Alice supplies $(S_i, 0)$ to an ideal bit OT. Bob chooses b and receives O_i. Both declare the constant output zero. Thus

$$O = \begin{cases} S & b = 0, \\ 0^n & b = 1, \end{cases} \qquad T = \begin{cases} O & b = 0, \\ R & b = 1. \end{cases} \tag{1}$$

The resources are deterministic given these queries. There is no correlated setup, public randomness, or hidden resource randomness. Initial execution is honest. Up to deterministic fields, the full records are (x, S) and (b, R, O); explicitly, they also include all sender vectors, receiver choices and responses, the output and local work. In particular Bob retains R even when it is unused. Alice's marginal does not depend on b, and Bob's does not depend on x, so this is ordinary perfect privacy for the constant-output functionality.

Before execution, the parties publicly fix deterministic maps

$$f : \mathbb{F}_2^n \longrightarrow \mathcal{A}, \qquad g : \mathbb{F}_2^n \longrightarrow \mathcal{B}.$$

These maps do not depend on b or the realized tapes. At a prescribed boundary they retain exactly

$$(x, A, 0) = (x, f(S), 0), \qquad (b, B, 0) = (b, g(T), 0). \tag{2}$$

All other state is erased: source tapes, unused coordinates, OT vectors and responses, intermediate copies and any ability to query the old resource again. The model assumes that erasure succeeds and that no external copy exists. Inputs, map descriptions, a fixed session identifier and the constant output survive. In each comparison these are fixed labels or constants; omitting the common constant fields preserves TV. Source-dependent map selection or random public metadata would require their own joint-law analysis.

State component	Execution state	Retained state
Inputs and output	Alice x; Bob b; output 0	Kept; claimed Bob input becomes the target label
Source tapes	Alice S; Bob R, even if unused	Only information encoded by A, B survives
OT records	Sender $(S_i, 0)$; queries b; responses O_i	No accessible copies of the logs
Intermediate state	T, work fields and all copies	Only A, B remain available
Public descriptions	Fixed f, g, schedule and session label	Fixed constants remain available
Old resource; backups	Original execution interface	No later access or external copies
Rewriting randomness	Not part of the execution record	Fresh independent Bob-private randomness

Table 1. State retained after the prescribed erasure. Inputs, outputs and fixed public descriptions remain available; all other information is restricted to A and B.

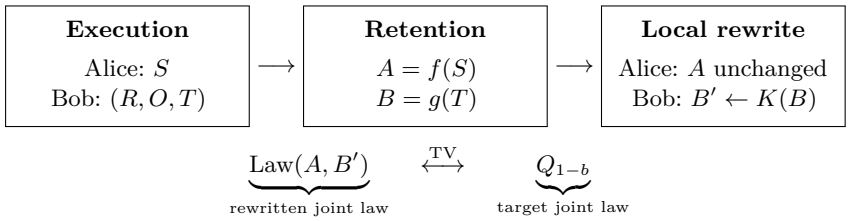

Fig. 2. The rewriting experiment. Retention removes access to the execution state except through A and B. The kernel acts only on Bob's retained symbol and uses independent randomness. Comparison always includes Alice's actual unchanged record. Fixed input and output labels are omitted.

2.2 Local rewriting and statistical error

As shown in Figure 2, Bob's rewriter receives only his retained record, the desired input label and fresh independent private randomness. It cannot consult T, R, O, an old log, or Alice's record. A later comparison combines Bob's claimed state with Alice's actual state. Only Bob is rewritten in this experiment. Alice's unused input can always be relabelled with her retained data fixed.

Omitting the fixed input and output labels, put

$$P = \mathrm{Law}(f(S), g(S)), \quad \mu = \mathrm{Law}(f(S)), \quad \nu = \mathrm{Law}(g(S)).$$

The honest retained joint laws are

$$Q_0 = P, \qquad Q_1 = \mu \otimes \nu. \tag{3}$$

Both have the same marginals. These are unconditional laws because the sole output has probability one. No impossible conditioning event occurs.

For two separately chosen stochastic kernels K_{01} and K_{10} acting only on Bob's retained symbol, define

$$\varepsilon^*(f,g) = \inf_{K_{01},K_{10}} \max\left\{\mathrm{TV}((\mathrm{Id},K_{01})_*Q_0, Q_1), \mathrm{TV}((\mathrm{Id},K_{10})_*Q_1, Q_0)\right\}. \tag{4}$$

Here $\mathrm{TV}(U,V) = \frac{1}{2}\sum_z |U(z) - V(z)|$. Returned records contain the target input label. The comparison averages over the actual retained pair. It is not a per-record conditional guarantee, and kernels need not invert individual executions. We allow arbitrary finite stochastic kernels with exact categorical sampling; efficiency is discussed separately.

3 Characterization of Local Rewriting

The reverse transformation starts from a product distribution. A local kernel on Bob's record preserves independence from Alice. The following proposition makes this restriction explicit.

Proposition 1 (Product-source calculation). *For every pair of finite deterministic retention maps,*

$$\varepsilon^*(f,g) = \delta(P) := \min_{r \in \mathcal{P}(\mathcal{B})} \mathrm{TV}(P, \mu \otimes r). \tag{5}$$

In particular, exact rewriting in both directions is possible if and only if $f(S)$ and $g(S)$ are independent.

Proof. For $0 \to 1$, Bob discards his retained symbol and independently samples from ν. Alice's actual marginal remains μ, so the pair has law Q_1 exactly.

For $1 \to 0$, any real stochastic kernel K gives

$$((\mathrm{Id},K)_*(\mu \otimes \nu))(a,c) = \sum_v \mu(a)\nu(v)K(v,c) = \mu(a)r(c),$$

$$r(c) = \sum_v \nu(v)K(v,c).$$

Conversely the constant-row kernel $K(v,c) = r(c)$ realizes every probability vector r. The two directions use separate kernels, establishing (5). The finite

probability simplex is nonempty and compact, and the TV objective is continuous, so its minimum is attained. If it is zero, $P = \mu \otimes r$; summing over a forces $r = \nu$. No division by a marginal is needed: zero-mass labels are included, and kernel rows at unseen labels may be assigned the same constant distribution. The converse follows by independent sampling. □

For a positive-mass Alice symbol, define $W_a(v) = P(a, v)/\mu(a)$. Factoring $\mu(a)$ out of each row of the half-L^1 sum gives

$$\delta(P) = \min_r \sum_{a:\, \mu(a)>0} \mu(a)\, \mathrm{TV}(W_a, r).$$

Zero-mass rows contribute zero and need no conditional law. This is a prior-weighted approximation. The worst-input radius in Raginsky [3], Theorem 3, uses

$$\min_r \max_{a:\, \mu(a)>0} \mathrm{TV}(W_a, r).$$

Appendix C gives a two-row example with different optimum values. Both comparisons keep the same family of conditional laws; changing the aggregation changes the objective.

The proof locates the inaccessible information. At input one Bob starts with a product distribution relative to Alice. Fresh coins cannot turn it into a nonproduct distribution while Alice remains fixed. At input zero he can always destroy correlation. The reverse direction therefore determines the bidirectional error.

For the two-bit example in the introduction, write $S = (S_1, S_2)$. Keeping $A = S_1$ and $B = S_2$ gives four pairs, each of probability $1/4$, under both inputs. Bob can change his input label and keep his retained bit. Keeping $A = B = S_1$ instead gives only 00 and 11 at input zero. From input one any local claim agrees with Alice's fair bit with probability exactly $1/2$. The test for agreement witnesses error at least $1/2$; a fresh fair claimed bit attains it. Both constructions retain one bit per party, but only the first removes the dependence between their records.

4 Bounds for Balanced and Linear Retention

4.1 Optimal error over balanced maps

Output length need not measure information about the source, since independent random bits can be appended to a record. We first restrict to deterministic balanced maps. A map to $\{0,1\}^k$ is balanced when every output has exactly 2^{n-k} preimages; its retained output consists of k uniform bits. The following optimizes the retention maps as well as the rewriter.

Theorem 2 (Balanced retention frontier). *For integers $n \geq 1$ and $0 \leq k, \ell \leq n$, among all balanced deterministic maps $f : \mathbb{F}_2^n \to \mathbb{F}_2^k$ and $g : \mathbb{F}_2^n \to \mathbb{F}_2^\ell$,*

$$\min_{f,g} \varepsilon^*(f,g) = \max\{0,\, 1 - 2^{\,n-k-\ell}\}. \tag{6}$$

Proof. Let $N = 2^n$, $m = 2^k$ and $h = 2^\ell$. For each retained Bob value v, its fiber $g^{-1}(v)$ has N/h elements. Consequently the target support $\operatorname{supp} P$ contains at most N/h different Alice values in that column. Since μ is uniform on m values, every candidate product law satisfies

$$(\mu \otimes r)(\operatorname{supp} P) \leq \sum_v r(v) \frac{N/h}{m} = \frac{N}{mh}. \tag{7}$$

The honest law puts probability one on that support. The support test and Proposition 1 give error at least $1 - N/(mh)$, as well as the trivial lower bound zero. The argument permits nonlinear maps and arbitrary biased fake marginals.

For attainment, let f keep the first k source coordinates and g keep the last ℓ. If $k + \ell \leq n$ these coordinates are disjoint, so the outputs are independent. Otherwise the two retained strings share $d = k+\ell-n$ coordinates. Their joint law is uniform on the N compatible pairs, and the product of the uniform marginals assigns those pairs total mass $N/(mh) = 2^{-d}$. On the support the honest mass is at least the product mass at every atom. Hence their TV distance is $1 - 2^{-d}$. Resampling a uniform Bob string attains the lower bound. □

Equation (6) is an existence frontier. It does not say that all maps using at most n retained bits in total permit exact rewriting. The identical-bit example has $k+\ell = n = 2$ and still has error $1/2$. When $k = 0$ or $\ell = 0$, the corresponding output is the unique empty string and the value is zero. At $k = \ell = n$ the value is $1 - 2^{-n}$. No balanced map onto the advertised alphabet exists for $k > n$ or $\ell > n$. Separately, extending the setup to $n = 0$ gives only $k = \ell = 0$ and value zero; this is a singleton extension, not a change to the theorem's $n \geq 1$ domain.

4.2 Fixed linear maps

The row-space intersection determines the error for fixed linear maps. Matrices below have entries in $\mathbb{F}_2$, and an empty matrix is allowed for a zero-bit output.

Proposition 3 (Fixed linear retention). *Let $f(s) = Ms$ and $g(s) = Ls$, with full row ranks k and ℓ. If*

$$d = \dim(\operatorname{row} M \cap \operatorname{row} L) = k + \ell - \operatorname{rank}\begin{pmatrix} M \\ L \end{pmatrix},$$

then $\varepsilon^*(f,g) = 1 - 2^{-d}$.

Proof. Write $\rho = \operatorname{rank}\binom{M}{L}$ and $W = \operatorname{im}\binom{M}{L}$. The stack's row space is row M + row L, so the dimension formula gives $d = k + \ell - \rho$. Every nonempty source fiber is a coset of the stack's kernel, of size $2^{n-\rho}$. Thus the shared-input pair is uniform on W. Projection from W onto Bob's coordinates is surjective because L has full row rank. Each fixed Bob value has exactly 2^{k-d} compatible Alice values: the stacked image has dimension $k+\ell-d$, and the projection has rank ℓ. Alice's marginal is uniform, so under every product fake the support test accepts with probability 2^{-d}. This gives the lower bound.

Uniform Bob resampling gives a uniform product on $2^{k+\ell}$ pairs. The target is uniform on $2^{k+\ell-d}$ pairs. The two laws' overlap mass is 2^{-d}, proving attainment. □

Figure 3 illustrates the formula for three pairs of coordinate projections with equal output lengths.

		Source coordinate				
Retention rule		1	2	3	4	Error
Disjoint	A	S_1	S_2			$d = 0$
coordinates	B			S_3	S_4	$\varepsilon^* = 0$
One shared	A	S_1	S_2			$d = 1$
coordinate	B		S_2	S_3		$\varepsilon^* = 1/2$
Two shared	A	S_1	S_2			$d = 2$
coordinates	B	S_1	S_2			$\varepsilon^* = 3/4$

Fig. 3. Fixed coordinate projections with $n = 4$ and $k = \ell = 2$. A cell labelled S_i denotes a retained coordinate; an empty cell denotes an omitted coordinate. Equal output lengths can give different row-space intersections and rewriting errors. Under the independent-source input, Bob's coordinates come from R.

Full row rank is essential to using the declared output lengths as ranks. Repeating a row adds a stored bit but no independent coordinate; mere nonconstancy does not justify replacing rank by length. Zero-row matrices have singleton output and are already covered.

In this linear family, an error strictly below $1/2$ is possible only when the error is zero: any positive overlap dimension contributes at least one shared constraint. This discrete conclusion is specific to uniform linear retention. It is not a general statistical rewriting threshold.

4.3 Randomized retention and source information

Here entropy and mutual information use base-two logarithms, with zero-mass terms interpreted as zero. The exact criterion also applies when each party uses a fixed local randomized retention rule, with finite output alphabets and coins mutually independent and independent of the source tapes. At input one the output law is still the product of the same marginals. At input zero, writing the local channels as $p(a \mid s)$ and $q(v \mid s)$ gives $\Pr[A = a, B = v \mid S = s] = p(a \mid s)q(v \mid s)$. Hence the Markov relation $A - S - B$ holds and $I(A; B \mid S) = 0$. If rewriting is exact, $I(A; B) = 0$, and the chain rule gives

$$I(S; A) + I(S; B) = I(S; A, B) + I(A; B) - I(A; B \mid S) = I(S; A, B) \leq n. \quad (8)$$

Thus independent padding does not evade a bound on information retained about the source. For deterministic maps it reads $H(A) + H(B) \leq n$. The derivation uses the standard chain rule for mutual information. The sharp approximate frontier (6) is asserted only for balanced deterministic maps. Shared retention coins can invalidate the conditional-independence equality or the product law at input one. They do not necessarily invalidate the final inequality: whenever A and B are independent, the general identity still gives a bound by $H(S)$ because $I(A; B \mid S) \geq 0$. Appendix C gives examples separating these conditions.

5 Optimal Rewriting for a Nonlinear Example

Resampling Bob's marginal is optimal for the linear maps above. This property does not extend to all balanced deterministic maps. The following example gives a strictly smaller error by changing the claimed marginal.

Take a uniform three-bit source, indexed by $s = 0, \ldots, 7$. Table 2 specifies two retention maps with output symbols $0, 1, 2, 3$.

s	0	1	2	3	4	5	6	7
$f(s)$	0	0	1	1	2	2	3	3
$g(s)$	0	0	1	1	2	3	2	3

Table 2. Balanced nonlinear retention maps on an eight-element source. Each output symbol has two preimages.

Both marginals are uniform, and the shared-input target table is

$$P = \frac{1}{8} \begin{pmatrix} 2\,0\,0\,0 \\ 0\,2\,0\,0 \\ 0\,0\,1\,1 \\ 0\,0\,1\,1 \end{pmatrix}. \quad (9)$$

Honest-marginal resampling gives the uniform distribution on all sixteen pairs. The overlap mass $\sum_{a,v} \min\{P(a,v), Q(a,v)\}$ is $6/16$, so its TV error is $5/8$.

Instead Bob can claim only values 2 and 3, each with probability $1/2$. The resulting product law has mass $1/8$ in those two columns and zero in the other columns. Its overlap with (9) is $4/8$, giving error $1/2$. No local rewriting kernel can do better: each column of (9) contains at most two supported rows among four equiprobable Alice values. Every product approximation therefore passes the honest-support test with probability at most $1/2$. This is also the lower bound (7).

The improvement comes from concentrating the claimed marginal on columns with greater overlap with the target law. The joint TV objective (4) already includes the resulting marginal discrepancy. Exact rewriting would still force the honest marginal. Adding an independent requirement that the approximate claimed marginal equal the honest one changes the problem and removes this particular improvement.

This example also explains why (6) is a frontier, rather than a claim that coordinate retention is the unique optimizer. Here some source states have the same retained pair, yet biased resampling still attains the frontier.

6 Effect of Additional Retained State

The retained-state specification is part of the experiment. Consider the two-bit construction with $A = S_1$, $B = S_2$, and suppose Alice also retains her old OT sender vectors. These vectors reconstruct all of S. Her actual available record is then equivalent to S, not just S_1. At input zero Bob's retained bit is a known coordinate of that full string; at input one his record is independent of it. Proposition 3, with M the identity, gives error $1/2$ again. More generally, if Alice keeps S and Bob keeps ℓ independent linear coordinates, the error is $1 - 2^{-\ell}$.

Bob's unused tape alone behaves differently. In the two-bit complementary construction, suppose Bob also keeps R but Alice keeps only S_1. At input zero his record is (S_2, R); at input one it is (R_2, R). Each is independent of Alice's fair S_1. Bob can sample the complete target marginal in either direction, so rewriting is still exact.

Keeping Bob's response log O restores an obstruction even when Alice keeps only $A = S_1$. Bob's retained pair (B, O) is (S_2, S) at input zero and $(R_2, 00)$ at input one. In the latter law his entire record is independent of the fair bit A. Any local claimed response O' therefore satisfies $O'_1 = A$ with probability $1/2$, whereas every honest input-zero pair satisfies it. The reverse-direction error is at least $1/2$. To attain it, sample a uniform two-bit O' and set $B' = O'_2$: the rewritten law assigns half the honest mass to each compatible atom, giving overlap $1/2$. For the other direction, sample a fair B' and set $O' = 00$, obtaining

the target law exactly. If Bob also keeps R, the same bound holds: append an independent uniform R' when claiming input zero; when claiming input one sample R' and set $B' = R'_2$, $O' = 00$. The agreement test originates in [1, Proposition 2]; the calculation above establishes both the lower bound and an attaining rewrite for the enlarged record.

The larger record defines a different experiment. Our proof uses an explicit projection of the entire ideal-OT execution onto (2), and its guarantee begins only after the projection has actually removed all other accessible information. Access to an old source tape changes the available information and is excluded from the rewriting kernel in (4).

7 Computational Considerations

The protocol uses n ideal bit OTs, $2n$ independent fair source bits and no public communication beyond fixed scheduling and an optional constant output announcement. Under fixed scheduling, its semantic complete execution records have $O(n)$-bit encodings, excluding the separately counted public map descriptions. A dense linear retention description needs $n(k + \ell)$ bits; matrix-vector evaluation and Gaussian elimination are polynomial in these dimensions. Uniform resampling uses ℓ fair bits. The nonlinear witness uses a fixed eight-entry table and one fair bit for its improved rewrite.

For arbitrary f, g, a complete truth table may have size exponential in n. Appendix A gives the explicit finite LP for (5) and proves rational attainment when P is rational. The LP has $|\mathcal{B}| + |\mathcal{A}||\mathcal{B}|$ variables and a comparable number of constraints. This table-based representation supplies no polynomial-time compiler from succinct nonlinear maps. Exact categorical sampling admits its rational optimum; bounded fair-bit exactness does not follow for every rational denominator. An eight-state example in Appendix B has a unique optimal claimed distribution with probability $1/5$ despite a dyadic honest law.

The coordinate constructions and linear algorithms are explicit. General attainment follows from compactness and does not by itself imply efficient optimization or sampling. Public-description length, execution-record length and rewriting complexity must therefore be accounted for separately.

References

1. Anonymous manuscript. *Local View Rewriting and Information Resale after Secure Computation*. Unpublished manuscript, 2026.
2. Pratyush Agarwal, Varun Narayanan, Shreya Pathak, Manoj Prabhakaran, Vinod M. Prabhakaran, and Mohammad Ali Rehan. *Secure Non-Interactive Reduction and Spectral Analysis of Correlations*. Cryptology ePrint Archive, Paper 2022/262, 2022. https://eprint.iacr.org/2022/262.
3. Maxim Raginsky. *Shannon meets Blackwell and Le Cam: channels, codes, and statistical experiments*. IEEE International Symposium on Information Theory, 2011. https://maxim.ece.illinois.edu/pubs/raginsky_ISIT11.pdf.
4. Danilo Silva and Frank R. Kschischang. *Universal Secure Network Coding via Rank-Metric Codes*. arXiv:0809.3546v2, 27 April 2010. https://arxiv.org/abs/0809.3546v2.
5. Carmit Hazay, Yehuda Lindell, and Arpita Patra. *Adaptively Secure Computation with Partial Erasures*. Cryptology ePrint Archive, Paper 2015/450, 2015; corrected full version revised 22 October 2017. https://eprint.iacr.org/2015/450.
6. Jörn Müller-Quade, Stefan Röhrich, and Dominique Unruh. *Oblivious Transfer is Incomplete for Deniable Protocols*. Workshop on The Past, Present and Future of Oblivious Transfer, 2005. https://kodu.ut.ee/~unruh/publications/ot_incomplete.pdf.

A Finite optimization and rational attainment

Let P_{av} be any rational probability table and $\mu_a = \sum_v P_{av}$. For probability laws P, Q, the identity $\mathrm{TV}(P,Q) = 1 - \sum_{a,v} \min(P_{av}, Q_{av})$ follows by expanding $|u-v| = u+v-2\min(u,v)$. Consequently (5) equals one minus the optimum of

$$\begin{aligned} \text{maximize} \quad & \sum_{a,v} z_{av} \\ \text{subject to} \quad & r_v \geq 0, \qquad \sum_v r_v = 1, \\ & 0 \leq z_{av} \leq P_{av}, \qquad z_{av} \leq \mu_a r_v \quad \text{for every } a, v. \end{aligned} \tag{10}$$

For fixed r, setting $z_{av} = \min(P_{av}, \mu_a r_v)$ proves equivalence, including zero-mass rows and columns. The feasible set is a nonempty bounded closed polytope: any r with $z = 0$ is feasible, r lies in a simplex, and each z is bounded by P. A linear objective attains its maximum on a face containing a vertex. At a vertex, a full-rank set of active constraints determines the coordinates uniquely; otherwise a nonzero feasible direction preserving all active constraints would express the point as the midpoint of two nearby feasible points. These active equations have rational coefficients, so Gaussian elimination gives a rational solution. Thus an optimal rational r exists for the original optimization over *all real* r.

This proof concerns an explicitly supplied rational table. The number of source states to enumerate from arbitrary maps can be 2^n, before the LP is even formed. For arbitrary real table entries, compactness still gives attainment, but rationality is not asserted.

B Exact Sampling with Fair Random Bits

Consider

$$P = \frac{1}{8}\begin{pmatrix} 2 & 1 \\ 1 & 4 \end{pmatrix}, \qquad \mu = (3/8, 5/8).$$

It is realized by eight uniform source labels with retained pairs $(0,0)$ twice, $(0,1)$ once, $(1,0)$ once, and $(1,1)$ four times. For $r = (t, 1-t)$, rowwise binary TV gives

$$\begin{aligned} D(t) = \mathrm{TV}(P, \mu \otimes r) &= \tfrac{3}{8}|t - \tfrac{2}{3}| + \tfrac{5}{8}|t - \tfrac{1}{5}| \\ &= \begin{cases} 3/8 - t & 0 \leq t \leq 1/5, \\ 1/8 + t/4 & 1/5 \leq t \leq 2/3, \\ t - 3/8 & 2/3 \leq t \leq 1. \end{cases} \end{aligned} \tag{11}$$

The slopes are $-1, 1/4, 1$, so the unique minimum over all real $t \in [0,1]$ is $D(1/5) = 7/40$.

A local kernel may use its random input as well as fresh coins. Here Bob's input marginal is $\nu = (3/8, 5/8)$. Suppose the kernel uses at most m fresh independent fair bits on every input and every branch. Padding each branch to length m gives $K(v,0) = j_v/2^m$ for integers $0 \le j_v \le 2^m$. Its claimed marginal therefore satisfies

$$r(0) = \sum_{v=0}^{1} \nu(v)K(v,0) = \frac{3j_0 + 5j_1}{2^{m+3}}.$$

This probability is dyadic and cannot equal $1/5$. Thus no such bounded-fair-bit kernel attains the optimum, even when it uses Bob's random input. Exact categorical sampling allows this optimum; fair-bit rejection sampling may instead use unbounded time.

C Alternative Error Criteria and Shared Randomness

C.1 Average and worst-record error

Let $P(0,0) = 1-p$, $P(1,1) = p$ with $0 < p < 1/2$ and all other masses zero. Writing t for the rewriter's probability of one gives $\mathrm{TV}(P, \mu \otimes r) = (1-p)t + p(1-t) = p + (1-2p)t$. The optimum is p, attained by always returning zero. Conditional on the rare actual Alice value one, that rewrite is always wrong. Thus small joint average TV does not give a worst-record guarantee. For the same two conditional rows, the worst-input objective is $\min_{0 \le t \le 1} \max\{t, 1-t\} = 1/2$: the maximum is at least the mean $1/2$, attained at $t = 1/2$. Its optimum therefore differs from the prior-weighted value p. Dyadic p realizes this example by deterministic retention of uniform source bits.

C.2 Shared retention randomness

Let both retention rules ignore their sources and output the same fair coin C, independent of S, R. Then $A = B = C$ under either input, and both joint laws assign probability $1/2$ to each of $(0,0)$ and $(1,1)$. The identity kernel gives exact rewriting despite the dependence between the records. Proposition 1 does not apply because the input-one law is not the product of its marginals.

Shared coins also affect the entropy calculation. For independent fair S, C, set $A = S \oplus C$ and $B = C$. Then A, B are independent and together determine S, so $I(S;A) + I(S;B) = 0$ while $I(S;A,B) = I(A;B \mid S) = 1$. Thus the equality in (8) need not hold with shared coins; the inequality remains valid.

D Exact Specialization of Secure Non-Interactive Reduction

We spell out the connection used in the related-work discussion. Represent a joint distribution by a matrix whose rows index Alice's symbols and whose columns index Bob's symbols. A local kernel is row-stochastic: its rows are probability distributions on output symbols.

At zero error, Theorem 3 of Agarwal et al. [2] characterizes a perfect secure non-interactive reduction of a target correlation D to a source correlation C by

$$A^{\mathrm{T}}CB = D, \qquad A^{\mathrm{T}}C = DV, \qquad CB = U^{\mathrm{T}}D.$$

Here A, B are the parties' protocol matrices and U, V their simulator matrices. These letters are local to this appendix; A, B do not denote retained random variables.

For the laws and kernels of Section 2, set

$$C = Q_0, \quad D = Q_1, \quad A = U = I, \quad B = K_{01}, \quad V = K_{10}.$$

The three identities then read, in the same order,

$$Q_0K_{01} = Q_1, \qquad Q_0 = Q_1K_{10}, \qquad Q_0K_{01} = Q_1.$$

The third repeats the first. The first two require precisely the two exact rewrites while leaving Alice's actual record unchanged. Conversely, kernels satisfying the two rewriting identities provide these protocol and simulator matrices. Thus exact bidirectional rewriting is equivalent to this specialization with Alice's protocol and simulator both fixed to the identity.

The source convention requires positive marginal supports. Since Q_0 and Q_1 have the same marginals, delete their common zero-mass rows and columns before applying the theorem. An exact kernel cannot send positive source mass to a zero-mass target symbol. Its rows at unseen input symbols can be extended arbitrarily when restoring the original alphabets. This restriction therefore does not change exact feasibility. The argument concerns one copy of each correlation and zero error; it makes no identification of approximate optima or of reductions using multiple copies.