Universal Contextual Composition Is Strictly Weaker than Neighbor-Preserving Differential Obliviousness

September 1, 2026

Abstract

For a randomized mechanism M, let P_x be the law of the first-stage execution $\mathrm{Exec}_M(x) = (V, Y)$ on a finite or countable space $\Omega = \mathcal{V} \times \mathcal{Y}$, where the input and output spaces have symmetric neighbor relations. Write aRb when two executions have the same view and neighboring outputs, and let $R(S) = \{b : \text{some } a \in S \text{ satisfies } aRb\}$. For $\eta \geq 0$ and $\gamma \in [0, 1]$, an (η, γ)-admissible context is a Markov kernel K satisfying $K(a, A) \leq e^{\eta} K(b, A) + \gamma$ for every measurable event A whenever aRb. We define universal contextual composability by quantifying over all such contexts. For $\varepsilon \geq 0$ and $\delta \in [0, 1]$, if the (ε, δ)-NPDO inequalities $P_x(S) \leq e^{\varepsilon} P_{x'}(R(S)) + \delta$ hold for every ordered pair $x \sim_{\mathcal{X}} x'$ and every $S \subseteq \Omega$, then, for every $\eta \geq 0$, $\gamma \in [0, 1]$, every (η, γ)-admissible K, and every contextual event A,

$$(P_x K)(A) \leq e^{\varepsilon+\eta} (P_{x'} K)(A) + \delta + \gamma - \delta\gamma \qquad (x \sim_{\mathcal{X}} x').$$

The converse fails with $\varepsilon = \delta = 0$. A finite mechanism with a fixed view exchanges masses $1/3$ and $2/3$ between the endpoints of a three-point output path. It is $(0, 0)$-universally contextually composable; in fact, for every $\eta \geq 0$, $\gamma \in [0, 1]$, (η, γ)-admissible kernel, and event, its two nontrivial contextual comparisons have additive loss at most

$$\frac{\max\{0, 2 + e^{\eta} - e^{2\eta}\}}{3} \gamma \leq \frac{2\gamma}{3}.$$

Nevertheless, it violates $(0, 0)$-NPDO. The mechanism is implemented in $O(1)$ time using one exact uniform sample from three outcomes and outputs an element of the three-point path with a fixed execution view.

Note. This paper was fully AI generated through a harness created by Richard Peng that discovers open problems and runs them through an automated research pipeline.

Contents

1 Introduction

Composability is particularly delicate when privacy protects what an execution reveals but not the intermediate value that it computes. Classical oblivious computation seeks to hide data-dependent access patterns, whereas *differential obliviousness* (DO) asks the observable access pattern to satisfy differential privacy on neighboring inputs [GO96, DMNS06, CCMS22]. This relaxation supports efficient algorithms for tasks such as database joins and relational operators [CZSC21, QJS+22]. In a pipeline, however, a first algorithm may expose a view V while passing an unobserved output Y to a second algorithm. Privacy of V alone does not ensure that neighboring first-stage inputs produce neighboring values of Y, so the second algorithm's privacy guarantee cannot simply be invoked.

Zhou, Shi, Chan, and Maimon introduced neighbor-preserving differential obliviousness (NPDO) to control precisely this interface [ZSCM23]. Their main composition theorem shows, on finite or countable execution spaces, that an NPDO first stage composes additively with a DO second stage. The paper also characterizes NPDO by a supported fractional matching between the two first-stage execution laws. Subsequent work develops advanced composition and divergence-based variants of NPDO [ZZCS24]. These results establish NPDO as a sufficient interface. They motivate a converse question: if a first stage remains private after every admissible downstream context, must it satisfy NPDO?

In the finite-or-countable statistical model, we show that NPDO suffices for a strong guarantee against every common downstream context but is not necessary: the converse already fails with zero first-stage error on a three-point output path. To state the result, let the input and output spaces carry symmetric neighbor relations, let $\mathrm{Exec}_M(x) = (V, Y)$ have law P_x on a finite or countable space $\Omega = \mathcal{V} \times \mathcal{Y}$, and write

$$(v, y)R(v', y') \iff v = v' \text{ and } y \sim_{\mathcal{Y}} y'.$$

For $S \subseteq \Omega$, let $R(S)$ be the set of executions related to some element of S. The (ε, δ)-NPDO condition is the setwise requirement

$$P_x(S) \le e^{\varepsilon} P_{x'}(R(S)) + \delta \qquad (x \sim_{\mathcal{X}} x')$$

for every $S \subseteq \Omega$. An (η, γ)-admissible context is a Markov kernel $K : \Omega \rightsquigarrow \mathcal{W}$ into an arbitrary measurable observation space such that

$$K(a, A) \le e^{\eta} K(b, A) + \gamma \qquad (aRb)$$

for every measurable event A. We call M (ε, δ)-*universally contextually composable* when, for every $\eta \ge 0$ and $\gamma \in [0, 1]$, every such common kernel applied in both neighboring worlds satisfies

$$(P_x K)(A) \le e^{\varepsilon+\eta} (P_{x'} K)(A) + \delta + \gamma \qquad (x \sim_{\mathcal{X}} x').$$

Our main theorem proves sufficiency and separates it from necessity.

Theorem 1.1 (Sufficiency and strict separation)**.** *Fix $\varepsilon \ge 0$ and $\delta \in [0, 1]$. If M is (ε, δ)-NPDO on a finite or countable execution space, then for every $\eta \ge 0$, $\gamma \in [0, 1]$, every (η, γ)-admissible Markov context K, and every measurable contextual event A,*

$$(P_x K)(A) \le e^{\varepsilon+\eta} (P_{x'} K)(A) + \delta + \gamma - \delta\gamma$$

for every ordered pair $x \sim_{\mathcal{X}} x'$. Hence M is (ε, δ)-universally contextually composable.

The converse fails with zero first-stage error. Let $\mathcal{X} = \{x_0, x_1\}$ with every pair related, let $\mathcal{Y} = \{0,1,2\}$ with $i \sim_{\mathcal{Y}} j$ exactly when $|i-j| \le 1$, and let the execution view always equal v. On input x_i, the algorithm ThreePoint outputs a vertex $y \in \mathcal{Y}$ and realizes the execution laws

	$(v,0)$	$(v,1)$	$(v,2)$
P_{x_0}	$1/3$	0	$2/3$
P_{x_1}	$2/3$	0	$1/3.$

It runs in $O(1)$ time assuming unit-cost exact uniform sampling from three outcomes. The resulting mechanism is $(0,0)$-universally contextually composable but is not $(0,0)$-NPDO. More strongly, for every $\eta \ge 0$, $\gamma \in [0,1]$, every (η,γ)-admissible K, and every measurable A, define

$$c_{\eta,\gamma} := \frac{\max\{0, 2 + e^{\eta} - e^{2\eta}\}}{3}\gamma \le \frac{2\gamma}{3}.$$

Then

$$(P_{x_0}K)(A) \le e^{\eta}(P_{x_1}K)(A) + c_{\eta,\gamma} \quad \text{and} \quad (P_{x_1}K)(A) \le e^{\eta}(P_{x_0}K)(A) + c_{\eta,\gamma}.$$

This common-kernel definition allows the context to depend on the full execution pair (V, Y) and therefore includes the usual output-fed serial composition as a special case. We use the phrase "universal contextual composability" only for this definition. Thus, under the common-context definition above, universal contextual composability is strictly weaker than NPDO already for a finite output graph and $(\varepsilon, \delta) = (0,0)$.

Relation to prior work. The direct NPDO results concern the same hidden-output interface as the present theorem: they prove sufficient composition theorems, while our result asks what is forced by quantification over all common downstream contexts. Structured forms of the interface also occur in guarantees for threshold binning, local-randomizer-to-DO-shuffler pipelines, and distance-preserving database operators [CCMS22, GKLX22, ZS22, QJS+22]. Those are restricted, mechanism-specific composition results rather than characterizations of arbitrary first-stage execution laws.

The setwise NPDO inequality also has a nearby, non-DO-specific interpretation. After aligning one-sided and symmetric conventions, it is the discrete form of a witness-free approximate relational lifting [Sat16]; on discrete spaces, approximate Strassen theorems turn this inequality into a supported star-lifting or fractional matching [BEH+19, ZSCM23]. Universal relational-testing formulations in the codensity-lifting literature quantify over a *pair* of related contexts [yKSU18, ABG+21]. Our property instead uses one common kernel in both worlds. It is this diagonal restriction, together with a nontransitive output adjacency relation, that leaves room for the converse to fail.

Proof ideas. For sufficiency, the NPDO neighborhood inequalities yield a fractional R-subcoupling that matches all but δ of P_x and whose right marginal is bounded by $e^{\varepsilon}P_{x'}$. Transporting the function $a \mapsto K(a, A)$ through this subcoupling keeps the exact unmatched mass rather than replacing it immediately by δ. The contextual error is then charged only on matched mass, giving the sharper bound $\delta + \gamma - \delta\gamma$ without exponential amplification. For strictness, the three-point path has no mass at its middle vertex. The endpoint laws fail the one-hop neighborhood comparison required by NPDO, while admissibility still relates the two endpoints through the middle vertex; averaging those two-hop constraints according to the displayed laws gives the sharper contextual bound in Theorem 1.1.

Organization. Section 2 formalizes execution laws, common contexts, NPDO, and fractional subcouplings. Section 3 explains the matching-and-transport argument and the three-point separation. Section 4 proves both parts of Theorem 1.1.

2 Preliminaries

Execution laws and relational neighborhoods. Let $(\mathcal{X}, \sim_{\mathcal{X}})$ and $(\mathcal{Y}, \sim_{\mathcal{Y}})$ be sets with symmetric neighbor relations. For a randomized mechanism M, write $\mathrm{Exec}_M(x) = (V, Y)$ for its execution record on input x, where $V \in \mathcal{V}$ is the observable view and $Y \in \mathcal{Y}$ is the output. Throughout, the execution space $\Omega := \mathcal{V} \times \mathcal{Y}$ is finite or countable and has the discrete σ-algebra. Its execution law is

$$P_x(a) := \Pr[\mathrm{Exec}_M(x) = a].$$

For any probability mass function p on Ω and $S \subseteq \Omega$, write $p(S) := \sum_{a \in S} p(a)$. For $a = (v, y)$ and $b = (v', y')$, define the symmetric execution relation

$$aRb \quad \iff \quad v = v' \;\text{ and }\; y \sim_{\mathcal{Y}} y', \tag{1}$$

and, for $S \subseteq \Omega$, its relational neighborhood

$$R(S) := \{b \in \Omega : \text{there exists } a \in S \text{ with } aRb\}. \tag{2}$$

All privacy comparisons below are required for each ordered neighboring pair $x \sim_{\mathcal{X}} x'$.

Contexts. Let $(\mathcal{W}, \mathcal{F})$ be an arbitrary measurable observation space. A Markov kernel $K : \Omega \rightsquigarrow \mathcal{W}$ assigns a probability measure $K(a, \cdot)$ to every $a \in \Omega$. For a law P on Ω, the contextual law PK is

$$(PK)(A) := \sum_{a \in \Omega} P(a) K(a, A), \qquad A \in \mathcal{F}. \tag{3}$$

Definition 2.1 (NPDO and universal contextual composability)**.** *Fix $\varepsilon \geq 0$ and $\delta \in [0, 1]$. The mechanism M is (ε, δ)-NPDO if, for every $x \sim_{\mathcal{X}} x'$ and every $S \subseteq \Omega$,*

$$P_x(S) \leq e^{\varepsilon} P_{x'}(R(S)) + \delta. \tag{4}$$

For $\eta \geq 0$ and $\gamma \in [0, 1]$, a kernel $K : \Omega \rightsquigarrow \mathcal{W}$ is (η, γ)-admissible if, for every aRb and $A \in \mathcal{F}$,

$$K(a, A) \leq e^{\eta} K(b, A) + \gamma. \tag{5}$$

The mechanism M is (ε, δ)-universally contextually composable if, for every $\eta \geq 0$, $\gamma \in [0, 1]$, every observation space, every (η, γ)-admissible K, every $x \sim_{\mathcal{X}} x'$, and every $A \in \mathcal{F}$,

$$(P_x K)(A) \leq e^{\varepsilon + \eta} (P_{x'} K)(A) + \delta + \gamma. \tag{6}$$

The NPDO clause in Definition 2.1 is the discrete setwise formulation of Zhou, Shi, Chan, and Maimon [ZSCM23]. The universal-composability clause fixes the meaning of that term throughout this paper. It includes the usual output-fed serial context. Indeed, if a downstream algorithm takes y as input and has view kernel L_y on a measurable space $\mathcal{Z}$, retain the first view and use observations in $\mathcal{V} \times \mathcal{Z}$. For an event A and a fixed v, set $A_v := \{z : (v, z) \in A\}$ and define

$$K((v, y), A) := L_y(A_v).$$

Whenever the kernels L_y and $L_{y'}$ satisfy the downstream privacy inequality for $y \sim_{\mathcal{Y}} y'$, this K is admissible by Equations (1) and (5).

Fractional subcouplings. Let p and q be probability mass functions on a finite or countable set Ω, let $R \subseteq \Omega \times \Omega$, and let $c \geq 1$. A c-dominated fractional R-subcoupling of p and q is a function $w : \Omega \times \Omega \to \mathbb{R}_{\geq 0}$ satisfying

$$w(a,b) > 0 \implies aRb, \tag{7}$$

$$\sum_b w(a,b) \leq p(a) \qquad \text{for every } a \in \Omega, \tag{8}$$

$$\sum_a w(a,b) \leq cq(b) \qquad \text{for every } b \in \Omega. \tag{9}$$

Its mass is $\operatorname{mass}(w) := \sum_{a,b} w(a,b)$. The row bound gives $\operatorname{mass}(w) \leq 1$; we call $1 - \operatorname{mass}(w)$ its unmatched mass.

Fact 2.2 (Countable approximation and finite flows)**.** *We use the following standard facts. If $h : D \to \mathbb{R}_{\geq 0}$ is summable on a countable set, then for every $\rho > 0$ there is a finite $F \subseteq D$ with $\sum_{z \notin F} h(z) < \rho$. If D is countable, every sequence in $[0,1]^D$ has a pointwise-convergent subsequence. Finally, in a finite network with nonnegative capacities, the maximum flow value equals the minimum capacity of a source–sink cut.*

3 Overview

The proof of Theorem 1.1 has two parts. For sufficiency, we turn the setwise NPDO inequalities into a fractional matching between neighboring execution laws and then transport an arbitrary contextual event through that matching. The central difficulty is the error accounting: NPDO controls indicators of sets, while a context tests the law against the $[0,1]$-valued function $a \mapsto K(a,A)$, and the desired conclusion must have additive loss at most $\delta + \gamma - \delta\gamma$ without multiplying either error by a privacy factor. For strictness, we construct a finite endpoint-exchange mechanism on a three-point path. Its laws fail the one-hop neighborhood comparison required by NPDO, although every admissible context can compare them through the two edges of the path.

The fractional matching lift. Fix an ordered pair $x \sim_{\mathcal{X}} x'$ and apply the NPDO inequality with $p = P_x$, $q = P_{x'}$, and $c = e^{\varepsilon}$. The countable fractional matching lift, Lemma 4.1, produces weights $w(a,b)$ satisfying

$$w(a,b) > 0 \implies aRb, \qquad \sum_b w(a,b) \leq P_x(a), \qquad \sum_a w(a,b) \leq e^{\varepsilon} P_{x'}(b),$$
$$\sum_{a,b} w(a,b) \geq 1 - \delta.$$

Thus w couples all but at most δ of the first law to related executions, while allowing the second law's capacity to be enlarged by the multiplicative factor e^{ε}. On a finite space, these weights are the flow in the bipartite graph induced by R: the NPDO set inequality is precisely the cut condition needed by max-flow/min-cut. For countable Ω, we build flows on finite truncations with mass $1 - \delta - o(1)$ and take a coordinatewise limit. A tightness argument prevents matched mass from escaping to infinity, which is the extra point needed beyond the finite fractional Hall argument.

Contextual transport. For an (η,γ)-admissible context, fix an event A, write $k(a) = K(a,A)$, and let $m = \sum_{a,b} w(a,b)$. The four properties above give the complete transport estimate

$$(P_x K)(A) \leq \sum_{a,b} w(a,b)k(a) + 1 - m$$

$$\le e^{\eta}\sum_{a,b} w(a,b)k(b) + \gamma m + 1 - m$$
$$\le e^{\varepsilon+\eta}(P_{x'}K)(A) + \delta + \gamma - \delta\gamma.$$

The support of w permits the pointwise admissibility comparison, its column bound supplies e^{ε}, and $m \ge 1-\delta$ gives $\gamma m + 1 - m \le \delta + \gamma - \delta\gamma$. Thus contextual error is paid only on matched mass, while the unmatched mass is not multiplied by e^{η}. Since the context, event, and ordered neighboring inputs were arbitrary, this proves the positive direction through Lemma 4.2.

The separating mechanism. For the converse, let the output relation be the path $0 \sim_{\mathcal{Y}} 1 \sim_{\mathcal{Y}} 2$ and let the execution view always be v. The two neighboring inputs have laws

$$\big(P_{x_0}((v,0)), P_{x_0}((v,1)), P_{x_0}((v,2))\big) = (1/3, 0, 2/3),$$
$$\big(P_{x_1}((v,0)), P_{x_1}((v,1)), P_{x_1}((v,2))\big) = (2/3, 0, 1/3).$$

The ThreePoint algorithm realizes these laws exactly: on input x_i it takes one uniform ternary sample, selects the corresponding endpoint, and writes it using a fixed constant-length access schedule whose view is v. Its output is the selected path vertex, its joint execution law is the displayed row, and it runs in $O(1)$ time under unit-cost exact ternary sampling. The zero mass at the middle vertex makes the NPDO failure immediate. For $S = \{(v,2)\}$,

$$R(S) = \{(v,1),(v,2)\}, \qquad P_{x_0}(S) = \frac{2}{3} > \frac{1}{3} = P_{x_1}(R(S)).$$

Why every admissible context still composes. The unused middle vertex nevertheless constrains a context at both endpoints. For an arbitrary contextual event, set $t = e^{\eta}$ and

$$a = K((v,0),A), \qquad b = K((v,1),A), \qquad c = K((v,2),A).$$

Chaining admissibility along the two edges gives

$$c \le tb + \gamma \le t^2 a + (t+1)\gamma, \qquad a \le tb + \gamma \le t^2 c + (t+1)\gamma.$$

A bare two-hop comparison would lose a second factor of t. The asymmetric endpoint masses are chosen so that their averaging removes this loss. Indeed, for the two contextual probabilities $\mu = (a+2c)/3$ and $\nu = (2a+c)/3$, the key identity is

$$3(\mu - t\nu) = (1-2t)a + (2-t)c.$$

Combining this identity with the appropriate two-hop bound, and then exchanging a and c for the reverse comparison, yields

$$\mu \le t\nu + \frac{\max\{0, 2+t-t^2\}}{3}\,\gamma.$$

The coefficient is at most $2/3$, so Lemma 4.3 proves the stronger additive guarantee $\max\{0, 2+e^{\eta}-e^{2\eta}\}\gamma/3 \le 2\gamma/3$, and hence $(0,0)$-universal contextual composability. Together with the one-hop NPDO violation, this completes the strict separation.

4 Proof of the Main Theorem

4.1 Sufficiency via fractional matching

The first ingredient turns the setwise expansion in Equation (4) into a fractional matching. Its countable form requires both a finite approximation and a tightness argument.

Lemma 4.1 (Countable fractional matching lift)**.** *Let Ω be finite or countable, let p and q be probability mass functions on Ω, let $R \subseteq \Omega \times \Omega$ be a relation, and let $c \geq 1$ and $\delta \in [0,1]$. Suppose that*

$$p(S) \leq cq(R(S)) + \delta \qquad \text{for every } S \subseteq \Omega. \tag{10}$$

Then there is a c-dominated fractional R-subcoupling w of p and q such that

$$\operatorname{mass}(w) \geq 1 - \delta. \tag{11}$$

Proof. Suppose first that Ω is finite. Form a flow network with a source, a left copy of Ω, a right copy of Ω, and a sink. Give the edge from the source to left vertex a capacity $p(a)$, each middle edge from a to b with aRb capacity $2 + c$, and the edge from right vertex b to the sink capacity $cq(b)$. The cut having only the source on its source side has capacity 1. Hence a minimum cut cannot cross a middle edge, whose capacity is greater than 1. If S is the set of left vertices on the source side of such a cut, all of $R(S)$ must also lie on its source side. The cut therefore has capacity at least

$$p(\Omega \setminus S) + cq(R(S)) \geq 1 - p(S) + p(S) - \delta = 1 - \delta,$$

where Equation (10) gives the inequality. By the finite max-flow/min-cut fact in Fact 2.2, there is a flow of value at least $1 - \delta$. The flows on the middle edges define the required subcoupling w.

Now suppose that Ω is countable. By the summable-tail fact in Fact 2.2, choose finite sets $F_n \subseteq \Omega$ and positive numbers $\rho_n \downarrow 0$ such that $p(F_n) \geq 1 - \rho_n$. For each of the finitely many subsets $S \subseteq F_n$, choose a finite set $G_{n,S} \subseteq R(S)$ satisfying

$$cq(R(S) \setminus G_{n,S}) < \rho_n,$$

and let G_n be the union of these sets. Apply the finite network construction with left vertex set F_n and right vertex set G_n. A cut determined by $S \subseteq F_n$ has capacity at least

$$\begin{aligned} p(F_n \setminus S) + cq(R(S) \cap G_n) &\geq p(F_n) - p(S) + cq(R(S)) - \rho_n \\ &\geq p(F_n) - \delta - \rho_n \\ &\geq 1 - \delta - 2\rho_n. \end{aligned}$$

Consequently there is a function w_n, extended by zero to Ω^2, that satisfies Equations (7) to (9) and has total mass at least $1 - \delta - 2\rho_n$.

Since $0 \leq w_n(a,b) \leq p(a) \leq 1$, the pointwise subsequence fact in Fact 2.2, applied to the countable set Ω^2, gives a subsequence along which every coordinate of w_n converges. Denote the pointwise limit by w. Passing to limits in finite partial sums proves Equations (7) to (9). To retain the total mass, fix $\rho > 0$ and choose finite $F, G \subseteq \Omega$ such that

$$p(F^{\mathsf{c}}) < \rho \qquad \text{and} \qquad cq(G^{\mathsf{c}}) < \rho.$$

The row and column bounds imply

$$\sum_{(a,b) \notin F \times G} w_n(a,b) \leq p(F^{\mathsf{c}}) + cq(G^{\mathsf{c}}) < 2\rho.$$

After passing to the limit in the finite sum over $F \times G$ and using $\rho_n \to 0$, we obtain

$$\sum_{a,b} w(a,b) \geq 1 - \delta - 2\rho.$$

Letting $\rho \downarrow 0$ proves Equation (11). $\square$

The matching lift is supported only on related executions. It therefore allows the admissibility inequality to be applied pointwise, while its row and column bounds control the two execution laws and its missing mass accounts for at most δ.

Lemma 4.2 (Contextual transport without error amplification)**.** *Every (ε, δ)-NPDO mechanism on a finite or countable execution space satisfies, for every $\eta \geq 0$, $\gamma \in [0, 1]$, every (η, γ)-admissible context K, every ordered pair $x \sim_{\mathcal{X}} x'$, and every contextual event A,*

$$(P_x K)(A) \leq e^{\varepsilon+\eta}(P_{x'}K)(A) + \delta + \gamma - \delta\gamma.$$

In particular, it is (ε, δ)-universally contextually composable.

Proof. Fix an ordered pair $x \sim_{\mathcal{X}} x'$. Apply Lemma 4.1 with

$$p = P_x, \qquad q = P_{x'}, \qquad c = e^{\varepsilon},$$

and the execution relation R. The hypothesis of that lemma is exactly Equation (4). For the resulting w, set

$$r(a) := \sum_b w(a, b), \qquad s(b) := \sum_a w(a, b), \qquad m := \sum_{a,b} w(a, b).$$

Fix arbitrary $\eta \geq 0$ and $\gamma \in [0, 1]$, an observation space, an (η, γ)-admissible kernel K, and an event A in that space. Write $k(a) := K(a, A)$. Since $0 \leq k(a) \leq 1$, the row and mass bounds give

$$\begin{aligned}(P_x K)(A) &= \sum_a r(a)k(a) + \sum_a (P_x(a) - r(a))k(a) \\ &\leq \sum_{a,b} w(a, b)k(a) + 1 - m.\end{aligned}$$

On the support of w, Equation (5) gives $k(a) \leq e^{\eta}k(b) + \gamma$. Thus, using the column bound and $m \geq 1 - \delta$,

$$\begin{aligned}(P_x K)(A) &\leq e^{\eta} \sum_b s(b)k(b) + \gamma m + 1 - m \\ &\leq e^{\varepsilon+\eta} \sum_b P_{x'}(b)k(b) + 1 - (1 - \gamma)m \\ &\leq e^{\varepsilon+\eta}(P_{x'}K)(A) + 1 - (1 - \gamma)(1 - \delta) \\ &= e^{\varepsilon+\eta}(P_{x'}K)(A) + \delta + \gamma - \delta\gamma.\end{aligned}$$

All choices were arbitrary. Since $\delta + \gamma - \delta\gamma \leq \delta + \gamma$, Equation (6) follows. □

It remains to prove that the implication is strict. The next construction places no mass at the middle vertex of a three-point path. The endpoint laws cannot satisfy a one-hop NPDO matching, but every admissible context relates the endpoints through the middle vertex.

4.2 The finite separating mechanism

Let

$$\mathcal{X} = \{x_0, x_1\}, \qquad \mathcal{V} = \{v\}, \qquad \mathcal{Y} = \{0, 1, 2\}.$$

Relate every pair of inputs, and define the output relation by

$$i \sim_{\mathcal{Y}} j \quad \iff \quad |i - j| \le 1.$$

Define a mechanism $M_\star$ by the execution laws

$$\begin{array}{c|ccc} & (v,0) & (v,1) & (v,2) \\ \hline P_{x_0} & 1/3 & 0 & 2/3 \\ P_{x_1} & 2/3 & 0 & 1/3 \end{array}. \tag{12}$$

Assume that an exact uniform sample from a three-element set is available at unit cost. Fix a constant-length access schedule σ, independent of the input and sampled value, that writes to a designated output location, and take v to be its access-pattern view. The following procedure implements the already defined law.

ThreePoint: Exact sampling of the endpoint-exchange mechanism.

Input: An input $x_i \in \mathcal{X}$, where $i \in \{0, 1\}$.
Output: An element $y \in \mathcal{Y}$ whose execution record is (v, y).
Guarantee: The execution law is the row indexed by x_i in Equation (12).

1. Sample U exactly and uniformly from $\{1, 2, 3\}$.
2. If $i = 0$, assign $y \leftarrow 2$ for $U \le 2$ and $y \leftarrow 0$ for $U = 3$. If $i = 1$, assign $y \leftarrow 0$ for $U \le 2$ and $y \leftarrow 2$ for $U = 3$.
3. Execute σ to write y to the designated output location, and return y; the execution view is v.

The implementation guarantee follows by counting the three equally likely samples. To prove the privacy claim, the essential additional step is to convert the two adjacent admissibility inequalities into a two-hop endpoint bound and then average according to Equation (12).

Lemma 4.3 (Exact realization and strict three-point separation)**.** *The* ThreePoint *algorithm realizes Equation (12) in $O(1)$ time under exact uniform ternary sampling. The resulting mechanism $M_\star$ is $(0,0)$-universally contextually composable. For every $\eta \ge 0$ and $\gamma \in [0,1]$, each nontrivial contextual comparison has additive term at most $\max\{0, 2 + e^\eta - e^{2\eta}\}\gamma/3$, and hence at most $2\gamma/3$. The mechanism is not $(0,0)$-NPDO.*

Proof. For input x_0, two values of U produce output 2 and one produces output 0; for input x_1, these probabilities are exchanged. The schedule σ always gives the fixed view v, so the joint laws are exactly Equation (12). The procedure uses one unit-cost sample and a constant number of other operations, hence runs in $O(1)$ time.

Fix $\eta \ge 0$, $\gamma \in [0,1]$, an (η,γ)-admissible kernel K, and an event A in its observation space. Set

$$t := e^\eta, \qquad a := K((v,0), A), \qquad b := K((v,1), A), \qquad c := K((v,2), A).$$

The output relation is symmetric and contains the two adjacent pairs, so admissibility in both directions yields

$$c \le tb + \gamma \le t^2 a + (t+1)\gamma, \qquad a \le tb + \gamma \le t^2 c + (t+1)\gamma. \tag{13}$$

The probabilities of A under the two contextual laws are

$$\mu := (P_{x_0} K)(A) = \frac{a + 2c}{3}, \qquad \nu := (P_{x_1} K)(A) = \frac{2a + c}{3}.$$

Direct calculation gives

$$3(\mu - t\nu) = (1 - 2t)a + (2 - t)c. \tag{14}$$

If $t \ge 2$, both coefficients on the right-hand side are nonpositive. If $1 \le t < 2$, the coefficient of c is positive, so the first bound in Equation (13) gives

$$\begin{aligned} 3(\mu - t\nu) &\le \big(1 - 2t + (2-t)t^2\big)a + (2-t)(t+1)\gamma \\ &= -(t-1)(t^2 - t + 1)a + (2 + t - t^2)\gamma \\ &\le (2 + t - t^2)\gamma. \end{aligned}$$

Combining the two cases gives

$$\mu \le t\nu + \frac{\max\{0, 2 + t - t^2\}}{3}\gamma.$$

Exchanging a and c and using the second bound in Equation (13) gives the same inequality with μ and ν interchanged. The inequalities for the two reflexive input pairs hold with zero additive term because $t \ge 1$. For $t \ge 1$, the displayed coefficient is at most $2/3$; since $2\gamma/3 \le \gamma$, this proves $(0,0)$-universal contextual composability, as well as the stated sharper comparison.

Finally, let $S = \{(v, 2)\}$. Its relational neighborhood is

$$R(S) = \{(v, 1), (v, 2)\}.$$

Equation (12) gives

$$P_{x_0}(S) = \frac{2}{3} > \frac{1}{3} = P_{x_1}(R(S)),$$

which contradicts Equation (4) with $\varepsilon = \delta = 0$. Thus $M_\star$ is not $(0,0)$-NPDO. □

Lemma 4.2 proves the sufficiency assertion of Theorem 1.1, and Lemma 4.3 supplies the finite constant-time counterexample. Hence universal contextual composability is strictly weaker than NPDO in the stated finite-or-countable statistical execution model.

References

[ABG+21] Alejandro Aguirre, Gilles Barthe, Marco Gaboardi, Deepak Garg, Shin ya Katsumata, and Tetsuya Sato. Higher-order probabilistic adversarial computations: Categorical semantics and program logics. *Proceedings of the ACM on Programming Languages*, 5(ICFP), 2021.

[BEH+19] Gilles Barthe, Thomas Espitau, Justin Hsu, Tetsuya Sato, and Pierre-Yves Strub. Relational $\star$-liftings for differential privacy. *Logical Methods in Computer Science*, 15(4):18:1–18:32, 2019.

[CCMS22] T.-H. Hubert Chan, Kai-Min Chung, Bruce M. Maggs, and Elaine Shi. Foundations of differentially oblivious algorithms. *Journal of the ACM*, 69(4):27:1–27:49, 2022.

[CZSC21] Shumo Chu, Danyang Zhuo, Elaine Shi, and T-H. Hubert Chan. Differentially oblivious database joins: Overcoming the worst-case curse of fully oblivious algorithms. In Stefano Tessaro, editor, *2nd Conference on Information-Theoretic Cryptography (ITC 2021)*, volume 199 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 19:1–19:24. Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2021.

[DMNS06] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam D. Smith. Calibrating noise to sensitivity in private data analysis. In Shai Halevi and Tal Rabin, editors, *Theory of Cryptography*, volume 3876 of *Lecture Notes in Computer Science*, pages 265–284. Springer, 2006.

[GKLX22] S. Dov Gordon, Jonathan Katz, Mingyu Liang, and Jiayu Xu. Spreading the privacy blanket: Differentially oblivious shuffling for differential privacy. In Giuseppe Ateniese and Daniele Venturi, editors, *Applied Cryptography and Network Security*, volume 13269 of *Lecture Notes in Computer Science*, pages 501–520. Springer, 2022.

[GO96] Oded Goldreich and Rafail Ostrovsky. Software protection and simulation on oblivious RAMs. *Journal of the ACM*, 43(3):431–473, 1996.

[QJS+22] Lianke Qin, Rajesh Jayaram, Elaine Shi, Zhao Song, Danyang Zhuo, and Shumo Chu. Adore: Differentially oblivious relational database operators. *Proceedings of the VLDB Endowment*, 16(4):842–855, 2022.

[Sat16] Tetsuya Sato. Approximate relational hoare logic for continuous random samplings. *Electronic Notes in Theoretical Computer Science*, 325:277–298, 2016.

[yKSU18] Shin ya Katsumata, Tetsuya Sato, and Tarmo Uustalu. Codensity lifting of monads and its dual. *Logical Methods in Computer Science*, 14(4):6:1–6:31, 2018.

[ZS22] Mingxun Zhou and Elaine Shi. The power of the differentially oblivious shuffle in distributed privacy mechanisms. Cryptology ePrint Archive, Paper 2022/177, 2022.

[ZSCM23] Mingxun Zhou, Elaine Shi, T.-H. Hubert Chan, and Shir Maimon. A theory of composition for differential obliviousness. In Carmit Hazay and Martijn Stam, editors, *Advances in Cryptology – EUROCRYPT 2023*, volume 14006 of *Lecture Notes in Computer Science*, pages 3–34. Springer, 2023.

[ZZCS24] Mingxun Zhou, Mengshi Zhao, T-H. Hubert Chan, and Elaine Shi. Advanced composition theorems for differential obliviousness. In Venkatesan Guruswami, editor, *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*, volume 287 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 103:1–103:24. Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2024.