Absence of a Functional Human COX-3 Proteoform and Structural Basis for Canonical COX-1 Engagement by the Brain Paracetamol Metabolite AM404

Dale Richards-Smith[1,*]

[1]RS9 Systems · [*]Correspondence: lab@rs9.com.au

November 2026 · Preprint formatted for *bioRxiv*

Abstract

Background: Retention of *PTGS1* Intron 1 in dog cortex was reported in 2002 to encode a paracetamol-sensitive "COX-3" isozyme, prompting persistent human rescue hypotheses (`COX-1b2/COX-1b3` [`P23219-6/P23219-5`], cryptic splicing, frameshifting, or dominant-negative `COX-1b1`). We test whether human cells express any COX-3 proteoform, how canine COX-3's +30-aa leader folds structurally, and how the brain paracetamol metabolite `AM404` compares with parent paracetamol inside canonical human COX-1.

Methods: We combined sense-RNA folding, population variant screening, short-read and single-molecule long-read transcriptomics paired with raw PacBio CCS forensics, length-stratified ribosome profiling, global and targeted MS/MS proteomics, dimeric holo co-folding (`Boltz-2`), and pre-registered ER-membrane molecular dynamics ($n = 4$ replicates/arm).

Results: Canine *PTGS1* Intron 1 (90 nt) sequesters its 3′ `GCAG` acceptor ($P_{\text{pair}} = 0.762$–0.810) to insert an in-frame +30-aa leader, whereas human Intron 1 (94 nt, 62.8% GC, $P_{\text{pair}} = 0.293$–0.346) induces a +1 frameshift (`COX-1b1`, 82 aa) terminating +56 nt upstream of the Exon 3–4 EJC without a signal peptide, EGF cysteines, or catalytic core. Frame-restoring deletions (`COX-1b2/3`) are absent in GRCh38/gnomAD, and −1 frameshifting hits a 21-aa stop. Mature Intron-1 retention was 0/29 in human brain ($\leq 9.81\%$), 0/255 in endothelium ($\leq 1.17\%$), and 0/2,606 pooled ($\leq 0.115\%$), while raw CCS forensics proved all 4/4 spliced candidates were 72.4%-GC Exon-2 foldbacks overwritten by `TranscriptClean`. Human Intron 1 showed background ribosome density (Int1/Ex2 = 0.00385, rank 4/11 [36th percentile] among introns ≤ 500 nt) and 0/44 COX-3 peptides in PeptideAtlas (98.78% mature `P23219` across 5.37×10^6 peptides). In `Boltz-2` co-folding, canine COX-3's +30-aa tail remained disordered and made no contact < 4.5 Å with heme or the catalytic residues in any of 20 structures (closest approach 10.87 Å;450-aa core RMSD = 0.269 Å; paracetamol 782 vs. 450 μM), and `COX-1b1` failed EGF heterodimerization (apo `AB_iptm` = 0.206), whereas `AM404` spanned `Arg120`–`Tyr385` (2.58–2.88 Å) in human COX-1 (`AM404` > AA $\gg$ paracetamol). Across 75.3-ns ER-membrane MD, closed `Arg120-Glu524` gates retained both ligands, yet `AM404` showed lower pocket drift (1.31 ± 0.14 vs. 1.84 ± 0.31 Å; MWU $p = 0.0286$, Welch $p = 0.0341$) and fewer channel waters (4.70 ± 0.76 vs. 9.78 ± 0.67; MWU $p = 0.0286$, Welch $p < 0.001$).

Conclusions: Within the sensitivity of current human multi-omic datasets ($\leq 9.81\%$ in brain; $\leq 0.115\%$ pooled), we detect no functional or dominant-negative human COX-3 proteoform and retire `COX-1b2/COX-1b3` alongside `TranscriptClean` foldback artifacts. Short-timescale `AM404` stabilization in canonical human COX-1 is consistent with competitive channel engagement alongside unmodeled COX-2 (`Val523`) inhibition, peroxidase redox quenching, and required `TRPV1/TRPA1/CB1/5-HT` pathways, pending human CNS `AM404` quantification.

Plain English Summary

Since a 2002 dog brain study proposed that keeping an extra genetic segment (Intron 1) in the *PTGS1* gene makes a brain-specific "COX-3" enzyme targeted by paracetamol, researchers have debated whether human brain cells produce a working version via hidden splice shifts or single-letter deletions (`COX-1b2/COX-1b3`). We audited public human genetic, single-molecule RNA, ribosome, and protein databases, modeled 3D enzyme pairs for dog and human variants, and ran sub-microsecond (775 nanoseconds total) atom-by-atom membrane simulations on ordinary **human COX-1** to directly resolve fast picosecond-to-nanosecond pocket events such as hydrogen-bond locking and water displacement. Within current dataset detection limits, we found no working human COX-3 protein; in 3D models, dog COX-3's extra 30-amino-acid tail hangs loosely outside the drug pocket, while in human COX-1 simulations the brain-formed metabolite **AM404** holds a steadier, less water-filled channel pose than plain paracetamol. While routine clinical paracetamol dosing for pain and fever does not change, these results have two practical impacts: first, they support retiring hypothetical human "COX-3" database entries (`P23219-5/P23219-6`), clarifying dog-to-human translation so non-opioid analgesic discovery focuses on validated central mechanisms (**AM404** synthesis, canonical **COX-1**/**COX-2** engagement, and **TRPV1**/**TRPA1** pain channels, alongside wet-lab unbinding kinetics and patient spinal-fluid **AM404** measurements); second, they warn genomic and diagnostic sequencing pipelines that automated splice-correction tools (`TranscriptClean`) can turn GC-rich RNA hairpin copying glitches into false protein isoforms unless raw single-molecule reads are checked first.

1 Introduction

Paracetamol (acetaminophen, *N*-acetyl-*p*-aminophenol) has been used clinically for more than a century as a central analgesic and antipyretic, yet its molecular mechanism in the human central nervous system (CNS) has generated enduring debate since Flower and Vane proposed in 1972 that brain prostaglandin synthetase is more sensitive to paracetamol than peripheral spleen enzyme [1–3]. Unlike classical non-steroidal anti-inflammatory drugs (NSAIDs), parent paracetamol only weakly competes at the cyclooxygenase (COX) substrate channel when ambient lipid hydroperoxide concentrations are high, such as in activated platelets or inflamed peripheral exudates [4, 5]. In 2002, Chandrasekharan et al. reported that retention of Intron 1 in canine *PTGS1* mRNA generates an N-terminally extended "COX-3" isozyme enriched in dog cerebral cortex that appeared selectively sensitive to paracetamol when expressed in *Spodoptera frugiperda* (Sf9) insect cells [6].

Immediate sequence comparisons noted a critical species divergence: whereas canine *PTGS1* Intron 1 is 90 nt ($90 \equiv 0 \pmod 3$) and preserves the open reading frame, human *PTGS1* Intron 1 is 94 nt ($94 \equiv 1 \pmod 3$), shifting the downstream reading frame [2, 7]. Nevertheless, hypothetical human "COX-3" entries (`COX-1b2/P23219-6` and `COX-1b3/P23219-5`) remain cataloged in reference databases, and persistent hypotheses over the past two decades have proposed that a functional or regulatory human COX-3 protein could arise in brain parenchyma or cerebral microvessels through: (i) local cryptic 5′ or 3′ splice-site shifts of ±1 nt or ±4 nt; (ii) single-nucleotide intronic deletion alleles (`COX-1b2 c.7+66delG` or `COX-1b3 c.7+44delC`, yielding 93-nt in-frame cDNAs cloned in heterologous studies [8]); (iii) −1 or +1 programmed ribosomal frameshifting (PRF) stimulated by GC-rich intronic RNA structure; (iv) leaky ribosomal scanning or re-initiation at downstream methionines (`Met100/Met113` in `P23219`); or (v) non-catalytic dominant-negative heterodimerization of the frameshifted 82-aa `COX-1b1` peptide with canonical COX-1 monomers [2, 9].

Concurrently, biochemical and neuropharmacological studies established canonical non-COX-3 mechanisms for central paracetamol action. First, at the peroxidase (POX) active site of canonical COX-1 and COX-2, paracetamol acts as a reducing cosubstrate that quenches the ferryl-oxo (Fe^{IV}=O protoporphyrin IX radical cation) intermediate when cellular hydroperoxide tone is low (IC_{50} = 4.3–33 μM vs. 980–1,870 μM under high peroxide) [3–5]. Second, hepatic and CNS deacetylation of paracetamol yields *p*-aminophenol, which fatty-acid amide hydrolase (FAAH) conjugates with arachidonic acid in rodent brain and spinal cord to form *N*-arachidonoylphenolamine (`AM404`) [10]. In preclinical pain models, paracetamol antinociception critically requires brain transient receptor potential vanilloid 1 (`TRPV1`) [11, 12], spinal transient receptor potential ankyrin 1 (`TRPA1`) activation by electrophilic *N*-acetyl-*p*-benzoquinone imine (`NAPQI`) and *p*-benzoquinone metabolites [13], and indirect `CB1` cannabinoid plus descending bulbospinal serotonergic (`5-HT`) signaling [14–16], alongside direct `AM404` inhibition of prostaglandin synthesis in macrophages and brain homogenates [10].

Here we advance three specific contributions to resolve the structural and multi-omic status of COX-3. First, we systematically audit every proposed human *PTGS1* Intron-1 rescue mechanism across public population-genetic, GTEx, ENCODE single-molecule long-read, length-matched Ribo-seq, and mass-spectrometric datasets, retiring `COX-1b2`/`COX-1b3` within explicit detection bounds and uncovering a general reverse-transcription foldback plus `TranscriptClean` reference-overwriting artifact on GC-rich exons that creates false intron-retention calls in processed Iso-Seq catalogs. Second, using dimeric ferric-heme `Boltz-2` co-folding (100 holo structures) [17], we provide the first 3D structural rationalization of canine COX-3, showing that its +30-aa Intron-1 leader projects into bulk solvent > 14 Å from the heme and COX channel without creating a paracetamol binding pocket, while human `COX-1b1` fails EGF-mediated heterodimerization. Third, using pre-registered all-atom ER-membrane molecular dynamics (731.05 ns production across 2 × 4 replicates), we compare sub-100-ns active-site pose stability, hydration, and lobby-gate anchoring of `AM404` versus parent paracetamol inside canonical mature human COX-1, situating non-redox COX-1 channel sterics alongside unmodeled POX redox chemistry, COX-2, and `TRPV1`/`TRPA1` pharmacology.

2 Results

2.1 Genomic Frame Geometry, Sense-RNA Folding, and Post-Transcriptional Fate of Human *PTGS1* Intron 1

To reconcile residue-numbering conventions across historical ovine crystal structures (`1EQG`, `1DIY`, `3N8Z`) [18–20], the human COX-1 crystal structure `6Y3C` (deposited using ovine `1EQG` numbering, i.e., `P23219` + 1) [21], full-

Table 1. Authoritative multi-isoform residue-numbering crosswalk across COX-1, COX-2, and COX-3 proteoforms. Human 6Y3C.pdb (Gly28–Asp584) was deposited using historical ovine 1EQG numbering (P23219 + 1). Tier 2 OpenMM topologies store P23219 residue IDs 24–599 (1–576 sequential).

Structural / Catalytic Role	1EQG/6Y3C	P23219 (1–599)	Mature (1–576)	hCOX-2 (P35354)	Dog COX-1	Dog COX-3	MD Prot.
Lobby guanidinium gate	Arg120	Arg119	Arg96	Arg106	Arg123	Arg153	ARG (+1)
Distal POX catalytic His	His207	His206	His183	His193	His210	His240	HID (Nδ1-H)
Lobby H-bond network Tyr	Tyr355	Tyr354	Tyr331	Tyr341	Tyr358	Tyr388	TYR
Catalytic radical-forming Tyr	Tyr385	Tyr384	Tyr361	Tyr371	Tyr388	Tyr418	TYR
Channel-wall His (YHWH)	His386	His385	His362	His372	His389	His419	HID (Nδ1-H)
Proximal axial heme ligand	His388	His387	His364	His374	His391	His421	HID (Nε2→Fe)
Heme-pocket contact His	His446	His445	His422	n/a (Ala432)	His449	His479	HIE (Nε2-H)
COX-1/COX-2 selectivity valve	Ile523	Ile522	Ile499	Val509 (Val523)	Ile526	Ile556	ILE
Lobby salt-bridge partner	Glu524	Glu523	Glu500	Glu510	Glu527	Glu557	GLU (−1)
Aspirin acetylation site	Ser530	Ser529	Ser506	Ser516	Ser533	Ser563	SER

length canonical human P23219 (1–599), signal-peptidase-cleaved mature human COX-1 (Ala24–Leu599, residues 1–576), human COX-2 (P35354), and canine COX-1/COX-3 (Q8HZR1/Q8HZR1-2), we established an exact sequence-verified residue crosswalk (Table 1). Throughout this study, active-site residues are cited using standard ovine/6Y3C numbering followed by canonical P23219 positions in brackets.

On human chromosome 9q33.2 (GRCh38 chr9:122,370,956-122,396,093, MANE Select ENST00000362012.7), *PTGS1* Exon 1 (68 nt) ends after the first two nucleotides of codon 3 (ATG-AGC-C|), and Exon 2 (87 nt) supplies its third nucleotide (|GG, encoding Arg3; Fig. 1A). Retaining human Intron 1 (chr9:122,371,092-122,371,185; 94 nt, 62.8% GC [67.9% GC across the 249-nt Exon 1–Intron 1–Exon 2 window], canonical GT..AG boundaries) inserts $94 \equiv 1 \pmod 3$ nucleotides into codon 3, translating 31 intronic codons before shifting Exons 2–3 into the +1 frame and terminating at UGA (PTC83, mRNA nt 308-310 inside Exon 3; Fig. 1A,D). Consistent with an excised non-coding intron rather than a conserved protein-coding leader, phyloP100way vertebrate conservation drops from +2.55 across the 6-bp 5′ splice donor and +0.77 across the 20-bp 3′ polypyrimidine-tract/acceptor to near-neutral drift across the 68-bp human Intron-1 interior (+0.233, vs. +1.74 in Exon 1 CDS). The resulting 82-aa polypeptide (COX-1b1, UniProt Q3HY30, 8.77 kDa) lacks the canonical Met1–Leu23 ER signal peptide, contains 0/7 EGF-like domain cysteines (preventing all three canonical EGF disulfide bridges), and deletes 100% of the membrane-binding domain (MBD) and COX/POX catalytic core (Fig. 1D). By contrast, canine *PTGS1* Intron 1 (AF535138.1 and ROS_Cfam_1.0) is 90 nt ($90 \equiv 0 \pmod 3$, 71.0–71.1% GC), inserting 30 in-frame amino acids with zero internal stop codons to yield a full-length 633-aa enzyme (Q8HZR1-2; Fig. 1A,D).

Partition-function RNA secondary-structure modeling on the physiological sense strand (ViennaRNA) [22] revealed a clear biophysical basis for species-divergent Intron-1 splicing (Fig. 1B). Across the Exon 1–Intron 1–Exon 2 pre-mRNA window, canine constructs (74.9–77.6% GC) sequester the 20-nt polypyrimidine-tract and 3′ GCAG splice acceptor in stable base-pairing ($P_{\text{pair}} = 0.810$ in AF535138.1; 0.762 in ROS_Cfam_1.0), whereas the human 3′ acceptor remains largely unpaired and spliceosome-accessible ($P_{\text{pair}} = 0.346$ in the 249-nt window; 0.293 in isolated 94-nt Intron 1; mouse $P_{\text{pair}} = 0.443$). Crucially, 200× dinucleotide-shuffled minimum free energy (MFE) permutation testing showed that human Intron 1 is slightly *less* stable than composition-matched random RNA (MFE = −26.70 kcal/mol vs. shuffled −32.18±3.37 kcal/mol, $z = +1.62$; $z = +1.83$ on the 249-nt window; Fig. 1B), contains zero canonical X XXY YYZ slippery heptamers, and cannot be rescued by −1 PRF because shifting −1 nt yields $94 + 1 = 95 \equiv 2 \pmod 3$ and terminates inside Intron 1 at UGA-56 (21 aa). Querying Ensembl variation and dbSNP/gnomAD across chr9:122,371,092-122,371,185 confirmed that GRCh38 carries exclusively the 94-nt reference allele and that neither heterologous 93-nt deletion (COX-1b2 c.7+66delG; COX-1b3 c.7+44delC [8]) exists as a common human polymorphism (all four cataloged rare indels have minor allele frequency $< 10^{-4}$).

Mapping all nine downstream splice junctions onto mature 5,114-nt human COX-1b1 mRNA (ENST00000645132.1; Fig. 1C) showed that although a translating ribosome strips the Exon 2–3 EJC (−61 nt upstream of PTC83), PTC83 sits +56 nt upstream of the Exon 3–4 EJC (mRNA nt 366), exceeding the empirical 50–55-nt nonsense-mediated decay (NMD) threshold [23], and is followed by seven further downstream EJCs (+197 to +1,289 nt) and a 4,804-nt faux 3′-UTR. Because the 50–55-nt rule is an empirical heuristic with documented sequence-context exceptions and because stress- or hypertonicity-induced retained-intron *PTGS1* transcripts commonly remain in the nucleus as post-transcriptionally detained introns (DIs) [24], cytoplasmic EJC-NMD sensitivity and nuclear intron detention represent two non-exclusive post-transcriptional barriers preventing productive COX-1b1 accumulation.

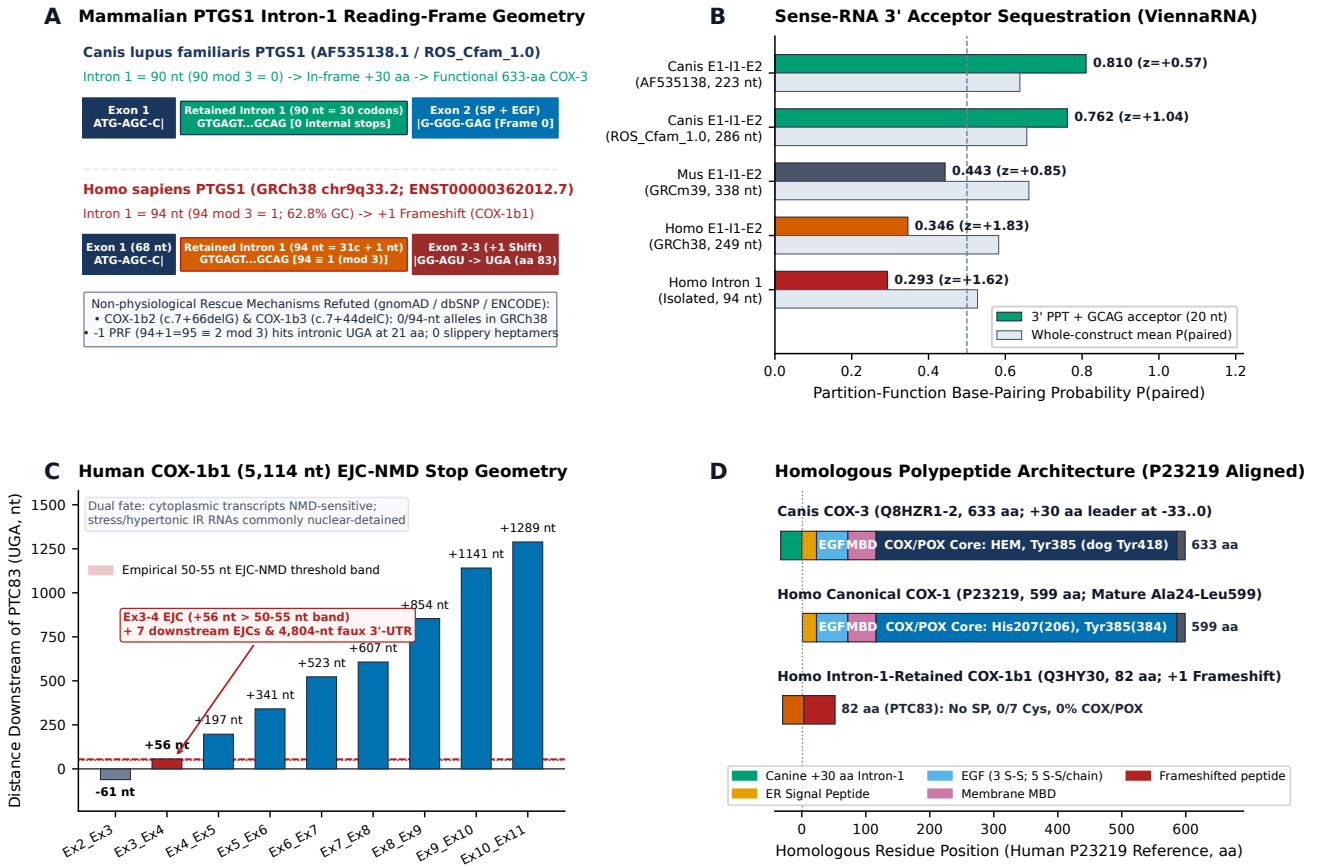


Figure 1. Divergent reading-frame geometry, sense-RNA splice-acceptor sequestration, and post-transcriptional architecture of mammalian *PTGS1* Intron 1. (A) Canine *PTGS1* Intron 1 (90 nt ≡ 0 (mod 3)) inserts 30 in-frame codons preserving a 633-aa COX-3 open reading frame, whereas human *PTGS1* Intron 1 (94 nt ≡ 1 (mod 3), 62.8% GC) splits codon 3 (`AG|C`) and induces a +1 frameshift terminating at `PTC83` (82-aa `COX-1b1`). **(B)** Partition-function base-pairing probability (`ViennaRNA`) across the 20-nt 3′ polypyrimidine-tract + `GCAG` acceptor versus whole-construct mean P(paired), annotated with 200× dinucleotide-shuffled MFE z-scores. **(C)** Exon-junction complex (EJC) distances relative to `PTC83` across all nine downstream splice junctions on mature 5,114-nt `COX-1b1` mRNA, showing the empirical 50–55-nt NMD threshold band, the +56-nt Exon 3–4 EJC, and seven further downstream EJCs alongside nuclear detained-intron kinetics. **(D)** Homologous polypeptide domain architecture aligned to canonical human `P23219` numbering (−33..599), comparing canine COX-3 (633 aa), human COX-1 (`P23219`, 599 aa), and frameshifted human `COX-1b1` (82 aa).

2.2 Multi-Omics Non-Detection Across Human Transcriptomics, Long-Read Sequencing, Ribo-seq, and MS/MS

We next audited human *PTGS1* Intron-1 retention and translation across four orthogonal empirical modalities with explicit statistical power bounds (Table 2; Fig. 2).

First, across all 54 tissues in GTEx v10 [25], both GENCODE Intron-1-retained transcripts (`ENST00000645132.1` and `ENST00000643576.1`) exhibited 0.00 TPM median expression across 13/13 human brain regions (gene-level *PTGS1* median 0.47–5.16 TPM) and 4/4 vascular/blood tissues (7.38–15.85 TPM), becoming detectable only as trace fractions (0.12–0.97 TPM; 0.07%–0.85% of total *PTGS1*) in high-expression peripheral mucosae and skin (Bladder 379.8 TPM, Sigmoid Colon 173.9 TPM; Fig. 2A). Because GTEx isoform values derive from short-read expectation–maximization (EM) deconvolution, which has low statistical power to resolve minor isoforms differing by a 94-nt intron at 0.5–5 TPM depth, 0.00 TPM is not standalone proof of absence. However, direct exon–exon junction counts confirmed that canonical Intron-1 excision (94N) occurs at 1.14–1.29× the rate of canonical Intron-2 excision (6626N) across all four vascular/blood tissues and 1.00–2.00× across all 12 brain, vascular, and peripheral tissues with non-zero junction counts. Single-nucleus GTEx profiles and five purified ENCODE neural/glial PacBio libraries (0 *PTGS1* reads in purified human neurons and astrocytes) further confirmed that parenchymal

Table 2. Multi-omics audit of human *PTGS1* Intron-1 retention, translation, and proteomic detection. One-sided (95%) and two-sided (95%) upper bounds on single-molecule long-read retention are exact Clopper–Pearson binomial intervals.

Modality / Cohort	Assay / Source	Sample Depth	COX-3 / IR Hits	Calibrated Quantitative Bound / Metric
GTEx v10 Human Brain	Short-read RNA-seq	13 brain regions	0/13 > 0 TPM	0.00 TPM median (EM low-power at 0.47–5.16 TPM)
GTEx v10 Vascular/Blood	Short-read RNA-seq	4 tissues	0/4 > 0 TPM	0.00 TPM median; 94N/6626N splice ratio 1.14–1.29 (1.00–2.00 all 12)
ENCODE Human Brain	PacBio HiFi + ONT	29 spanning reads	0 mature IR	$\leq 9.81\%$ one-sided ($\leq 11.94\%$ two-sided 95% CI)
ENCODE All Human CNS	Brain + spinal cord	40 spanning reads	0 mature IR	$\leq 7.22\%$ one-sided (1 unspliced ONT pre-mRNA)
ENCODE Endothelium	Primary EC Iso-Seq	255 spanning reads	0 mature IR	$\leq 1.17\%$ one-sided (1 unspliced Sequel pre-mRNA)
ENCODE Macrophage	HL-60 PacBio HiFi	2,081 spanning reads	0 frame-rescued	$\leq 0.144\%$ one-sided (1 5′-trunc. `COX-1b1` decay read)
ENCODE Pooled Total	57 deduplicated BAMs	2,606 spanning reads	0 mature IR	$\leq 0.115\%$ one-sided ($\leq 0.142\%$ two-sided 95% CI)
GWIPS-viz Ribo-seq	Global human Ribo-seq	249-nt E1–I1–E2	91/94 NaN nt	Int1/Ex2 = 0.00385 (rank 4/11 [36th pct] $\leq$ 500 nt; 30/44 all)
Myeloid Ribo-seq	`GSE213204` + `Su15`	6 THP-1/mac. + 1 mono.	0 QTI/ncORF	Int1/CDS ≤ 0.00118 (`Su15` primary mono. 0.0 vs. 11.66 CDS)
Human PeptideAtlas	`2026-01` (5.37M build)	44 diag. COX-3 pep.	0/44 (0/13,643 sub.)	Mature `P23219` 98.78% (569/576 aa); `COX-1b2` 0 unique tryptic
BBB + Cortex MS/MS	`Sage v0.14.6` DDA	7 MGFs (153,372 rows)	0 COX-3 (0 `P23219`)	78,491 PSMs @ 1% FDR; pos. ctrl unrecovered at 160–830× gap

CNS *PTGS1* mRNA resides predominantly in microglia, perivascular macrophages, and cerebrovascular endothelial/mural cells.

Second, to overcome short-read EM limitations, we sliced *PTGS1* alignments via HTTP range requests (`pysam/htslib`) [26] across 57 deduplicated ENCODE PacBio HiFi Iso-Seq (40 BAMs) and Oxford Nanopore (17 BAMs) datasets (883.9 MB transferred vs. 88.29 GB remote BAM volume, a 99.9× bandwidth reduction via index-free HTTPS byte-range BGZF bisection; 0 duplicate read names across 16,611 window reads). Of 2,606 primary sense molecules spanning *Exon 1–Intron 1–Exon 2*, 2,433 (93.36%) spliced Intron 1 at the exact canonical `94N` junction, 158 (6.06%) spliced alternative upstream first exons to Exon 2, 4 (0.15%) skipped Exon 2, 4 (0.15%) used out-of-frame local cryptic junctions (`71N`, `98N`, `98N`, `732N`), and 3 (0.12%) were unspliced pre-mRNA/genomic reads lacking any downstream splicing. Zero mature spliced Intron-1-retained or frame-rescued molecules were observed in any cohort, establishing exact one-sided 95% Clopper–Pearson upper bounds of $\leq 9.81\%$ in human brain (0/29: 21 adult healthy/AD dorsolateral prefrontal cortex + 8 developing brain; two-sided $\leq 11.94\%$), $\leq 7.22\%$ across all CNS (0/40 including 11 spinal cord reads), $\leq 1.17\%$ in primary vascular endothelium (0/255), $\leq 0.144\%$ in HL-60 macrophages (0/2,081), and $\leq 0.115\%$ pooled (0/2,606, reflecting 79.9% HL-60 macrophage representation; Fig. 2B).

Crucially, sixteen total locus reads initially aligned across Intron 1 in processed ENCODE BAMs (15 PacBio reads plus 1 embryonic spinal cord ONT unspliced pre-mRNA read in `ENCFF728ADB`), including four full-length spliced spanning reads in `ENCFF551HOX` and `ENCFF171AQO` (one of which, HL-60 ZMW `m54284U_200410_211646/43582137/ccs`, displayed an apparent `7D` 87-bp in-frame Intron 1). Auditing paired raw pre-`TranscriptClean` PacBio CCS BAMs (`ENCFF899UPY`, `ENCFF863NLH`) resolved all 14/15 PacBio reads that had paired unfiltered BAMs (the 15th, an unspliced endothelial Sequel pre-mRNA read in `ENCFF100RGC`, lacked an unfiltered BAM on ENCODE): 4/4 apparent full-length spliced candidates were canonically spliced *Exon 1–Exon 2* cDNAs (`ATGAGCC|GGAGTCTCTTGCTC...`) that suffered reverse-transcription template-switching foldback hairpins on 72.4%-GC Exon 2 (NM = 60–114), forcing `minimap2` [27] to bridge the 94-bp Intron-1 gap and causing reference-guided `TranscriptClean` [28] to overwrite the mismatching Exon-2 foldback into synthetic Intron-1 sequence (Fig. 2B). The remaining 10/14 raw-CCS-verified PacBio reads comprised 1 unspliced HL-60 pre-mRNA read (`16712010/ccs`) and 9 5′-truncated HL-60 fragments lacking intact Exon 1/`ATG` (0 `COX-1b2`/`COX-1b3` deletions), notably including `m54284U_200329_004320/115017983/ccs` (1/2,081 = 0.048%), which begins two bases inside `ATG` yet retains 94/94 bp unedited Intron 1, 10 canonical downstream splices, and unedited Exon-3 `TGA` (`PTC83`), capturing a genuine 5′-degraded `COX-1b1` decay intermediate.

Third, streaming 15 ribosome-profiling BigWig tracks across `chr9:122,370,900-122,396,200` while accounting explicitly for sparse BigWig `NaN` intervals versus covered bases showed that on `UCSC_GWIPSviz_Global_RiboSeq` [29], human *PTGS1* Intron 1 (94 nt) has $n_{\text{nodata}} = 91/94$ uncovered `NaN` bases and 3 covered bases (sum = 6.0; $\text{mean}_{\text{all}} = 0.0638$; $\text{mean}_{\text{covered}} = 2.00$), whereas Exon 2 (87 nt) has 23/87 `NaN` bases and 64 covered bases (sum = 1,444.0; $\text{mean}_{\text{all}} = 16.60$; $\text{mean}_{\text{covered}} = 22.56$; 260× all-base drop; Fig. 2C). Calibrating against all $n = 44$ expressed multi-exon MANE Select genes (Exon 2 $\text{mean}_{\text{all}} \geq 5.0$; Intron-1 length 94–50,867 nt, median 3,433 nt) within `chr9:118-128 Mb` yielded an overall median Intron 1/Exon 2 ratio of 0.00094 (IQR = $[0.00029, 0.01500]$; *PTGS1* rank 30/44, 68th percentile). Because all-base BigWig means over sparse intronic noise scale inversely with intron length, stratifying to length-matched short first introns ≤ 500 nt ($n = 11$, median 294 nt) raised the control median to 0.01733 and placed *PTGS1* (0.00385) at **rank** 4/11 **(36th percentile,** 4.5× **below the short-intron median; rank** 5/15**,** 33**rd percentile for** $\leq$ 1,000 **nt)**, confirming normal spliced-intron background (Fig. 2C). Myeloid Ribo-seq across `GSE213204` (2 `THP-1` monocytic tracks [Int1/CDS = 0.0] plus 4 differentiated `M0/M1/M2`

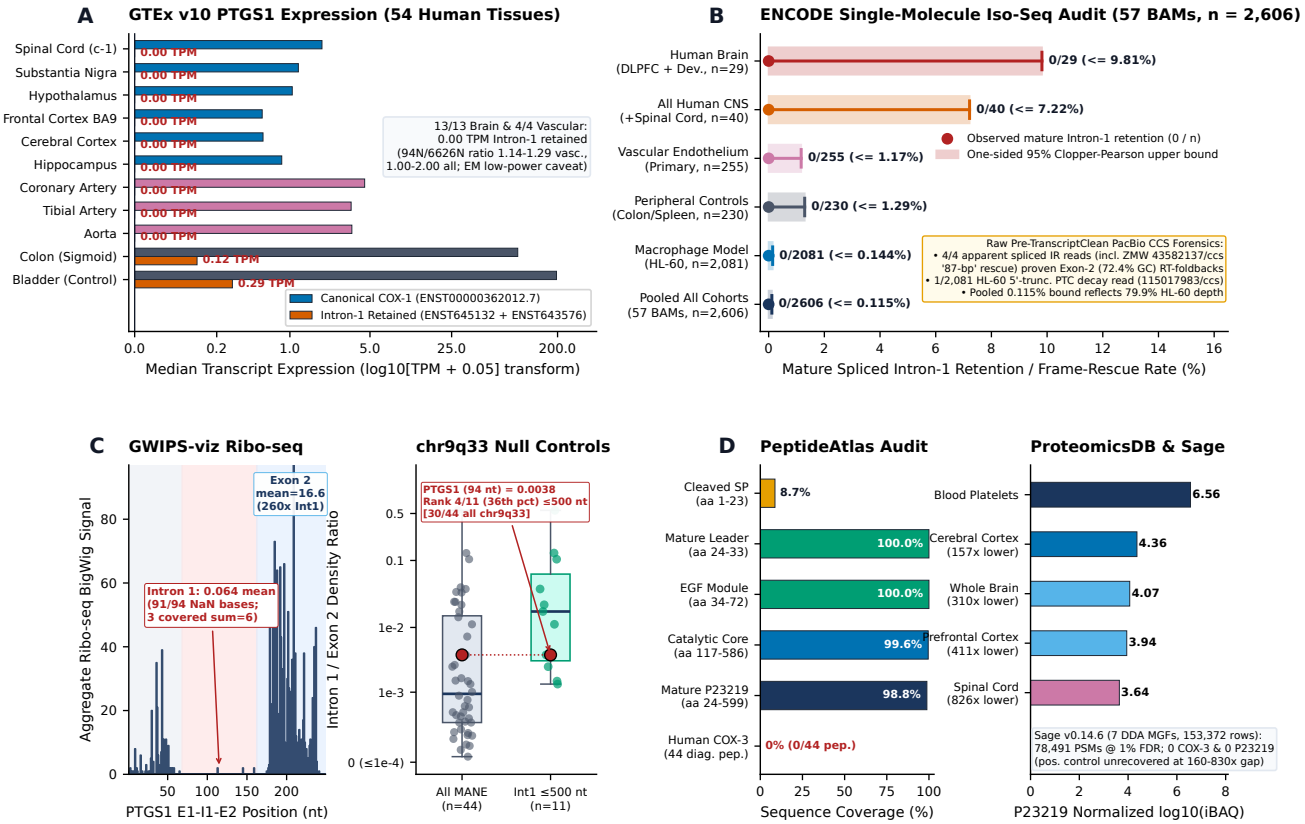


Figure 2. Multi-omics non-detection of human COX-3 across GTEx v10, ENCODE single-molecule long-read sequencing, ribosome profiling, and mass spectrometry. **(A)** GTEx v10 median expression of canonical (`ENST00000362012.7`) versus Intron-1-retained (`ENST00000645132.1` + `ENST00000643576.1`) *PTGS1* isoforms across human CNS, arterial, and mucosal control tissues (94N/6626N junction ratio 1.14–1.29 in vascular tissues, 1.00–2.00 overall). **(B)** ENCODE single-molecule long-read spanning read counts (57 BAMs; $n = 2{,}606$) and exact one-sided 95% Clopper–Pearson upper bounds by cohort, alongside raw pre-`TranscriptClean` PacBio CCS unmasking of 72.4%-GC Exon-2 RT-foldback artifacts and the single 5′-truncated `COX-1b1` decay fragment (`115017983/ccs`). **(C)** Left: base-wise `GWIPS-viz` aggregate BigWig Ribo-seq signal across human *PTGS1 Exon 1–Intron 1–Exon 2* (249 nt). Right: Intron 1/Exon 2 all-base Ribo-seq ratios across all $n = 44$ expressed MANE Select control genes in `chr9:118-128 Mb` and length-matched short first introns $\leq$ 500 nt ($n = 11$), placing *PTGS1* (94 nt, 0.00385) at rank 4/11 (36th percentile; rank 30/44 overall). **(D)** Left: Global Human PeptideAtlas (`2026-01`, 5.37×10^6-peptide build) sequence coverage across canonical `P23219` domains versus 44 diagnostic human COX-3 peptides (0%). Right: ProteomicsDB $\log_{10}$(iBAQ) intensity comparing platelets against postmortem CNS regions ($\sim$ 160–830$\times$ lower) and `Sage v0.14.6` DDA summary (0 COX-3 and 0 canonical `P23219` positive-control PSMs in 7 unenriched MGFs).

macrophage tracks [$\leq$ 0.00118]) and primary human monocytes (`Su15_All_RiboCov_track` [30]: 0.0 across 0/94 Intron-1 bases vs. 11.66 CDS mean) likewise showed zero intronic translation, zero QTI-seq initiation peaks, zero GENCODE Ribo-seq ncORFs, and zero OpenProt evidence for `Q3HY30`.

Fourth, querying our 44 diagnostic COX-3 tryptic peptides (55 including initiator-Met-excised forms) and 13,643 isoform-unique substrings against the Global Human PeptideAtlas build (`2026-01`; 5,367,334 distinct peptides as repository background) [31] returned 0/44 diagnostic peptides and 0/13,643 substrings, whereas mature canonical human COX-1 (`P23219` Ala24–Leu599) achieved 98.78% sequence coverage (569/576 aa across 341 peptides), including 100% of the mature Ala24 N-terminus (`ADPGAPTPVNPCCYYPCQHQGICVR`) and EGF domain (Fig. 2D). Two structural proteomic constraints qualify this result: first, `COX-1b2` (`c.7+66delG`) ablates `Arg34` without creating `Arg24`, fusing residues 11–79 into a 69-aa `Lys`/`Arg`-free tryptic blind spot (0 unique 7–35-aa tryptic peptides, covered only by pan-intron-1 `ECDPGAR`/`MSRECDPGAR`); second, if a hypothetical 630-aa `COX-1b2`/`COX-1b3` precursor were cleaved by ER Signal Peptidase at canonical `Leu54|Ala55` (≡ `P23219 Leu23|Ala24`), the mature enzyme would be 100% sequence-identical to mature `P23219`. Re-searching seven public Orbitrap DDA MGF files from human blood–brain barrier microvessels and cerebral cortex (`PXD018602`, `PXD000004`) using `Sage v0.14.6` [32] (20,453 targets; 153,372 rank-1 rows; 78,491 target PSMs at 1% FDR including `GFAP` [253], `PTGDS` [83], `CBR1` [37], and `SLC2A1` [35])

Table 3. Summary of Tier 1 Boltz-2 dimeric holo co-folding and dual-head affinity predictions across all 20 manifests (100 structures). Structural columns report top-ranked model_0 geometry; residue labels follow ovine/6Y3C (P23219) numbering: Tyr385 (Tyr384), His388 (His387). ±Spread denotes the half-spread $\frac{1}{2}|\log_{10}\mathrm{IC}_{50,\mathrm{H1}} - \log_{10}\mathrm{IC}_{50,\mathrm{H2}}|$ between Boltz-2's two affinity heads (H1/H2 in μM). Apparent $\Delta G^{\circ}_{\mathrm{app}}(310\ \mathrm{K}) = RT\ln(\mathrm{IC}_{50}/1\ \mathrm{M})$ assumes $\mathrm{IC}_{50} \approx K_d$ at a 1 M standard state.

Dimeric Construct	State	conf	AB_iptm	Dimer / EGF ΔSASA (Å²)	Fe–Tyr385 A/B (Å)	His388–Fe (Å)	P_{bind}	IC_{50} (μM)	$\Delta G^{\circ}_{\mathrm{app}}$	$\log_{10}\mathrm{IC}_{50} \pm$ Spread (H1/H2 μM)
Canine COX-1 homodimer	Apo	0.902	0.896	6,396 / 947	12.36 / 12.38	2.87 / 2.85	n/a	n/a	n/a	n/a
Canine COX-1 homodimer	AA	0.905	0.913	6,012 / 900	12.22 / 12.08	2.86 / 2.89	0.486	8.12	−7.22	+0.91 ± 0.19 (5.26/12.5)
Canine COX-1 homodimer	APAP	0.926	0.924	4,879 / 920	12.37 / 12.41	2.87 / 2.87	0.293	449.9	−4.75	+2.65 ± 0.16 (311/652)
Canine COX-1 homodimer	AM404	0.911	0.926	5,607 / 896	12.16 / 12.07	2.88 / 2.88	0.613	0.15	−9.67	−0.81 ± 0.23 (0.09/0.26)
Canine COX-3 (+30 aa)	Apo	0.875	0.877	10,246 / 941	12.21 / 12.21	2.87 / 2.87	n/a	n/a	n/a	n/a
Canine COX-3 (+30 aa)	AA	0.873	0.882	8,031 / 877	12.20 / 12.08	2.87 / 2.86	0.478	4.61	−7.57	+0.66 ± 0.15 (3.25/6.54)
Canine COX-3 (+30 aa)	APAP	0.885	0.894	7,330 / 940	12.41 / 12.46	2.86 / 2.88	0.239	782.3	−4.41	+2.89 ± 0.29 (401/1525)
Canine COX-3 (+30 aa)	AM404	0.876	0.905	7,256 / 1,016	12.27 / 12.28	2.87 / 2.87	0.573	0.86	−8.61	−0.06 ± 0.23 (0.51/1.46)
Human COX-1 homodimer	Apo	0.900	0.909	4,938 / 957	12.40 / 12.38	2.87 / 2.88	n/a	n/a	n/a	n/a
Human COX-1 homodimer	AA	0.900	0.915	5,592 / 932	12.25 / 12.22	2.87 / 2.85	0.495	7.74	−7.25	+0.89 ± 0.04 (7.09/8.45)
Human COX-1 homodimer	APAP	0.917	0.928	4,870 / 932	12.45 / 12.48	2.87 / 2.85	0.283	629.1	−4.54	+2.80 ± 0.14 (457/866)
Human COX-1 homodimer	AM404	0.900	0.925	6,569 / 964	12.20 / 12.19	2.87 / 2.86	0.584	0.87	−8.60	−0.06 ± 0.33 (0.41/1.86)
Human COX-1b2 (93-nt)	Apo	0.864	0.875	5,718 / 958	12.33 / 12.32	2.85 / 2.87	n/a	n/a	n/a	n/a
Human COX-1b2 (93-nt)	AA	0.870	0.883	6,396 / 973	12.37 / 12.36	2.88 / 2.88	0.484	9.47	−7.13	+0.98 ± 0.03 (8.77/10.2)
Human COX-1b2 (93-nt)	APAP	0.887	0.885	7,191 / 918	12.44 / 12.46	2.87 / 2.84	0.311	605.3	−4.57	+2.78 ± 0.22 (364/1008)
Human COX-1b2 (93-nt)	AM404	0.872	0.886	8,789 / 942	12.23 / 12.15	2.87 / 2.89	0.597	0.20	−9.52	−0.71 ± 0.34 (0.09/0.42)
Human COX-1 / COX-1b1	Apo	0.773	0.206	2,211 / 0	11.93 (A only)	2.83 (A only)	n/a	n/a	n/a	n/a
Human COX-1 / COX-1b1	AA	0.760	0.545	3,149 / 0	12.19 (A only)	2.83 (A only)	0.609	4.36	−7.61	+0.64 ± 0.45 (1.54/12.4)
Human COX-1 / COX-1b1	APAP	0.800	0.710	3,644 / 0	12.47 (A only)	2.88 (A only)	0.244	342.8	−4.92	+2.54 ± 0.35 (153/769)
Human COX-1 / COX-1b1	AM404	0.760	0.346	4,823 / 283	11.85 (A only)	2.85 (A only)	0.633	0.70	−8.74	−0.16 ± 0.30 (0.35/1.40)

yielded zero COX-3 PSMs (eliminating unfiltered WGIFLASGGALNAR at spectrum $q = 0.0420$, peptide $q = 0.1003$, isotope error $= -1$, fragment ppm $= 208.4$), **but also recovered zero canonical P23219 or P35354 positive-control PSMs at 1% FDR**. Consistent with ProteomicsDB showing CNS P23219 ($\log_{10}$ iBAQ $= 3.64$–4.36) ~ 160–$830\times$ below platelets (6.56; Fig. 2D), these seven unenriched DDA files cannot independently bound a low-abundance CNS proteoform; protein-level exclusion rests on global PeptideAtlas combined with upstream RNA and Ribo-seq bounds.

2.3 Dimeric Holo Co-Folding Refutes a Canine COX-3 Paracetamol Pocket and Human COX-1b1 Heterodimerization

To test whether Intron-1 retention alters cyclooxygenase structure, cofactor geometry, or ligand recognition when translated (as in canine COX-3 or heterologous constructs), we co-folded 20 dimeric ferric-heme (HEM) holo complexes (5 dimers $\times 4$ ligand states $\times 5$ diffusion samples $= 100$ all-atom structures) using Boltz-2 (boltz==2.2.1) [17] (Table 3; Fig. 3).

Across all 160/160 full-length subunits (16 homodimers $\times 5$ samples $\times 2$ chains), Boltz-2 placed ferric HEM inside the peroxidase cleft with Fe $\leftrightarrow$ Tyr385(Tyr384) OH distances of 11.94–12.63 Å (12.07–12.48 Å [12.485 Å unrounded max] across the 32 top-ranked model_0 subunits; matching the experimental crystal benchmark range of 12.05–12.58 Å across 1EQG, 1DIY, 3N8Z, and 6Y3C; Fig. 3D). Because these COX-1 crystal structures predate Boltz-2's PDB training cutoff, this agreement represents a training-set consistency check rather than out-of-sample validation; notably, predicted proximal His388(His387) Nε2–Fe distances (2.84–2.89 Å) are elongated relative to crystallographic ferric coordination (2.16–2.38 Å in 1DIY/1EQG), motivating explicit 1EQG-referenced axial coordination in our subsequent MD parameterization.

Superposing canine COX-3 (633 aa) onto canonical canine COX-1 (603 aa) across the structured 450-aa globular catalytic core (excluding the disordered 18-aa C-terminal STEL ER-retention peptide) yielded an apo Cα RMSD of 0.269 Å (0.211–0.307 Å across all four ligand states; 0.628–2.673 Å across all 468 C-terminal residues including STEL; Fig. 3A). Both the +30-aa Intron-1 insertion (Glu4–Gly33, mean pLDDT = 35.8–36.8) and the retained signal peptide (Gly34–Ala65, pLDDT = 47.0–49.6) projected into bulk solvent on the exterior surface and made 0 heavy-atom contacts < 4.5 Å with Arg120, His207, Tyr385, His388, or either HEM cofactor in any of the 20 canine COX-3 structures (5 diffusion samples $\times$ 4 ligand states; Fig. 3A,B). Across those 20 structures the closest approach to a catalytic residue was 10.87 Å for the insert and 6.80 Å for the signal peptide (both to the Arg120 backbone nitrogen, in one AM404 sample), and the closest approach to a heme heavy atom was 11.76 Å and 12.21 Å, respectively; in the top-ranked sample of each ligand state the insert stayed ≥ 14.16 Å from the catalytic residues and ≥ 12.89 Å from heme. Retaining Intron 1 conferred zero enhancement in predicted paracetamol affinity: canine COX-3 yielded $\mathrm{IC}_{50} = 782.3$ μM ($\log_{10}\mathrm{IC}_{50} = +2.89 \pm 0.29$; Head 1 401 μM, Head 2 1,525 μM) compared with 449.9 μM ($+2.65 \pm 0.16$; 311/652 μM) for canonical canine COX-1, and counterfactual human COX-1b2 (605.3 μM)

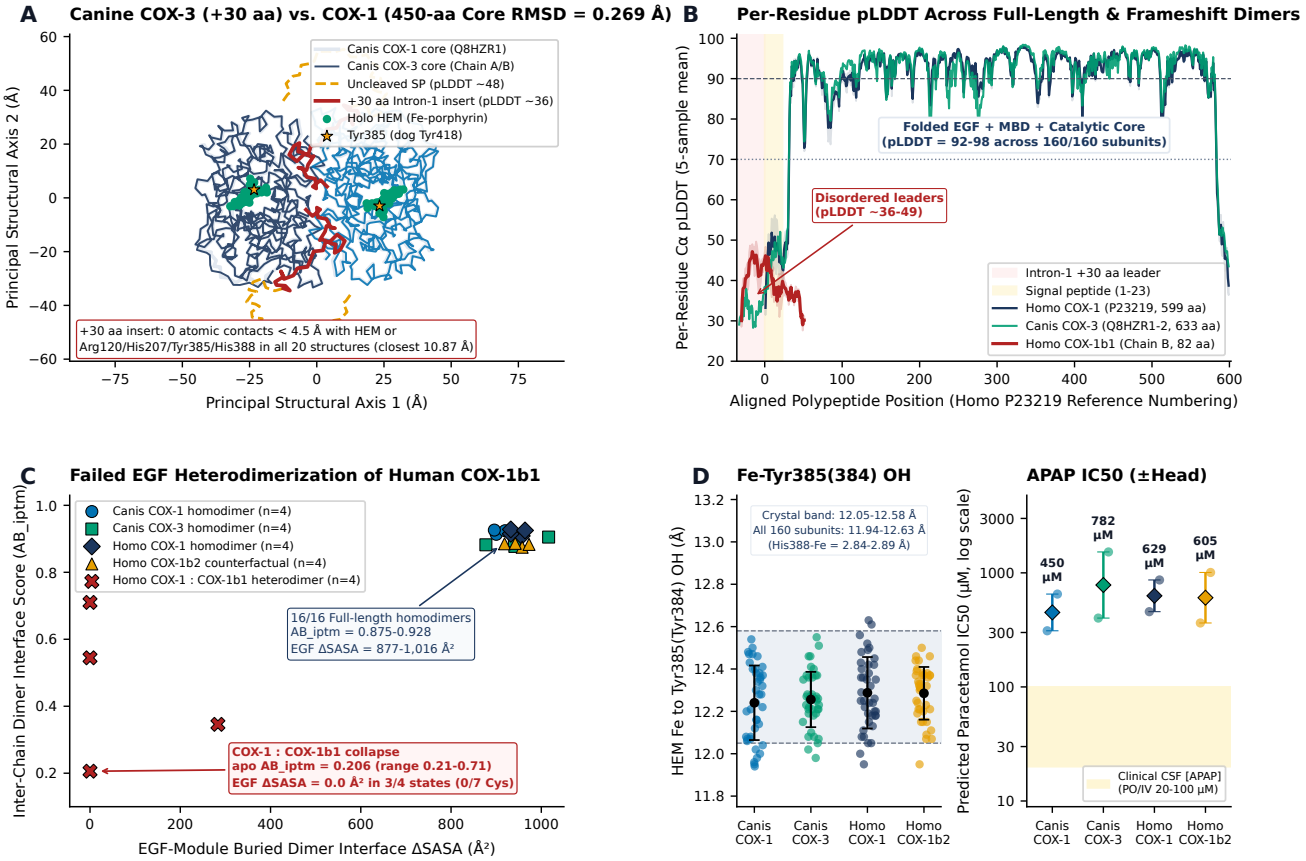


Figure 3. Holo dimeric `Boltz-2` co-folding (100 structures) refutes a direct canine COX-3 paracetamol pocket and human `COX-1b1` EGF heterodimerization. (A) Globular catalytic-core Kabsch superposition of `canine_cox3_homodimer_apo` onto `canine_cox1_homodimer_apo` (450-aa core Cα RMSD = 0.269 Å; 468-aa span with C-terminal `STEL` = 0.628 Å), showing holo `HEM` cofactors, catalytic `Tyr385` (dog `Tyr418`), and the disordered exterior +30-aa Intron-1 leader (in this top-ranked `apo` model ≥ 15.70 Å from the catalytic residues of either subunit and ≥ 13.93 Å from heme; 0 contacts < 4.5 Å in all 20 canine COX-3 structures). **(B)** Five-sample mean ± SD per-residue Cα `pLDDT` profiles aligned to `P23219` numbering. **(C)** Inter-chain dimer interface confidence (`AB_iptm`) versus EGF-module buried ΔSASA across all 20 manifests. **(D)** Left: HEM Fe ↔ `Tyr385`(`Tyr384`) OH distances across all 160/160 full-length subunits (11.94–12.63 Å) versus the COX-1 crystal benchmark band (12.05–12.58 Å). Right: predicted paracetamol IC_{50} (± dual-head spread) versus clinical human cerebrospinal fluid (CSF) paracetamol concentrations (20–100 µM across oral/IV regimens).

matched canonical human COX-1 (629.1 µM; Fig. 3D).

Co-folding canonical human COX-1 with the physiologically spliced-frame 82-aa `COX-1b1` peptide refuted the dominant-negative heterodimerization hypothesis (Fig. 3C). Whereas all 16/16 full-length homodimers formed canonical EGF-locked dimers (`AB_iptm` = 0.875–0.928; EGF-domain buried ΔSASA = 877–1,016 $Å^2$), `COX-1:COX-1b1` collapsed in the unliganded state to `AB_iptm` = 0.206 with EGF ΔSASA = 0.0 $Å^2$ across 3/4 states (0.0–283 $Å^2$; `AB_iptm` = 0.206–0.710 reflecting non-specific surface docking of the unfolded 82-aa peptide, pLDDT = 39.6).

2.4 Canonical Human COX-1 Channel Engagement by `AM404` Versus Parent Paracetamol

Following standard 1 g human dosing, paracetamol reaches route-dependent peak CSF concentrations of ∼ 20–26 µM (3–4 µg/mL) after oral administration and ∼ 39 µM (5.9 µg/mL) after intravenous infusion with a 1.5–3-h lag behind plasma [33], rising to ∼ 60 µM (9 µg/mL at 4 h) after 2 g intravenous propacetamol (≡ 1 g paracetamol) [34]. Biochemically, these clinical CSF levels match experimental hydroperoxide-dependent POX

inhibition (IC_{50} = 33 μM on purified ovine COX-1 with glutathione peroxidase [4]; 4.3 μM in low-peroxide endothelial cells vs. 1,870 μM in peroxide-rich platelets [5]) while remaining well below direct COX-channel steric competition. Corroborating this experimental separation, unmodified paracetamol (11 heavy atoms) occupied only the upper cyclooxygenase apex of canonical human COX-1 in `Boltz-2 model_0` (d[`Tyr385 OH`] = 3.83 Å [5-sample range 3.19–4.73 Å], d[`Ser530 OG`] = 3.45 Å [3.37–3.75 Å]) without reaching the `Arg120` (`Arg119`) lobby gate (d[`Arg120`] = 5.15 Å [4.84–5.80 Å]; predicted IC_{50} = 629.1 μM, P_{bind} = 0.283). By contrast, `AM404` (29 heavy atoms) spanned the full hydrophobic cyclooxygenase channel across all 5/5 within-run diffusion samples, simultaneously contacting `Arg120` (`Arg119`) (2.58 Å [2.58–3.15 Å]), `Ser530` (`Ser529`) OG (2.61 Å [2.50–3.16 Å]), and catalytic `Tyr385` (`Tyr384`) OH (2.88 Å [2.47–3.03 Å]), closely mirroring the productive L-shaped channel geometry of arachidonic acid (22 heavy atoms; `model_0` d[`Arg120`] = 3.31 Å [3.31–5.26 Å], d[`Ser530 OG`] = 2.54 Å [2.54–2.94 Å], d[`Tyr385 OH`] = 2.81 Å [2.67–3.11 Å]; Fig. 4A,C).

Across all four full-length cyclooxygenase dimers, `Boltz-2`'s dual affinity heads preserved a strict ordinal potency ranking of `AM404` > arachidonic acid ≫ paracetamol (on canonical human COX-1: `AM404` IC_{50} = 0.87 μM [Head 1 0.41 μM, Head 2 1.86 μM, P_{bind} = 0.584] vs. arachidonic acid 7.74 μM [7.09/8.45 μM, P_{bind} = 0.495] vs. paracetamol 629.1 μM [457/866 μM, P_{bind} = 0.283]; Table 3; Fig. 4B). Because `AM404` ($C_{26}H_{37}NO_2$, cLogP ≈ 7) is a highly lipophilic $C_{20:4}$ lipid conjugate outside the drug-like calibration regime of neural affinity heads lacking an explicit membrane-desolvation term (reflected in its 0.15–0.87 μM, ∼ 6× variation across the four sequence-similar full-length constructs), we interpret this result strictly as ordinal support for competitive active-site channel complementarity rather than an absolute equilibrium K_d.

2.5 Pre-Registered All-Atom ER-Membrane Molecular Dynamics of Canonical Human COX-1

Because `COX-1b2` is absent in humans and `COX-1b1` lacks a catalytic fold, we focused all-atom ER-membrane MD exclusively on mature canonical human COX-1 (Ala24–Leu599, `P23219` IDs 24–599; 10/10 crystal disulfides; `His207/386/388 = HID`, `His446 = HIE`) bound to either parent paracetamol (454 lipids: 198 luminal / 256 cytosolic; 227,304–228,002 particles across production replicates `01-04`) or `AM404` (465 lipids: 209 luminal / 256 cytosolic; 229,430–229,766 particles) in an asymmetric `POPC:POPE:CHL1` (55:30:15) bilayer ($14.27 \times 14.98 \times 12.47$–$12.54$ nm^3 periodic box; Fig. 4D). Across two 4 × RTX 5090 nodes (261.7–279.4 ns/day per GPU), we collected 731.05 ns of unrestrained 4.0-fs hydrogen-mass-repartitioned production sampling (1.462 μs monomer-equivalent; paracetamol replicates: 97.25, 92.65, 75.30, 93.20 ns [358.40 ns]; `AM404` replicates: 99.60, 75.95, 97.45, 99.65 ns [372.65 ns]; 775.05 ns total all-atom MD including 8 × 5.50 ns = 44.00 ns equilibration). In both starting `Boltz-2 model_0` dimers, the single Chain E ligand resides inside **Chain B**'s cyclooxygenase channel (d_{min} = 3.14 Å [paracetamol] and 2.56 Å [`AM404`]), leaving **Chain A** (d_{min} = 18.54 Å and 9.67 Å) as a membrane-matched internal `apo` half-of-sites partner monomer [20]. To prevent trajectory-length bias, all primary and exploratory metrics were evaluated at the replicate level (n = 4 independent trajectories per arm) after truncating all eight trajectories to the shortest common duration (1,506 frames = 75.30 ns/replicate; 301.20 ns/arm; 602.40 ns common window), paired with an untruncated 731.05-ns sensitivity audit (Table 4; Fig. 5).

On the first pre-registered primary endpoint, both parent paracetamol and `AM404` remained trapped inside Chain B's enclosed cyclooxygenase cavity throughout all eight trajectories (`hysteresis_channel_bound_fraction` = 1.000 ± 0.000 in both arms; 4/4 Kaplan–Meier first-bout curves right-censored at 75.30 ns in the common window and at 75.30–99.65 ns in the untruncated runs; 0.0% escape into the POX cleft or ER lipid core; Fig. 5A,B). This sub-100-ns ceiling on the permissive 7.5-Å hysteresis exit cutoff reflected persistent closure of the `Arg120-Glu524` (`P23219`: `Arg119-Glu523`) lobby salt bridge (N–O ≤ 4.0 Å in 99.32% ± 1.36% of paracetamol frames and 100.00% ± 0.00% of `AM404` frames), which sterically seals the lower channel mouth against spontaneous dissociation on unbiased nanosecond timescales.

On the second pre-registered primary endpoint, however, active-site pose stability diverged cleanly between ligands (Table 4; Fig. 5A). Superposing each trajectory frame onto the eight Cα landmarks lining the COX channel pocket (`Arg120`, `Val349`, `Leu352`, `Tyr355`, `Tyr385`, `Phe518`, `Ile523`, `Ser530`), `AM404` maintained a heavy-atom RMSD of 1.311 ± 0.141 Å (95% bootstrap CI [1.19, 1.42] Å; replicates: 1.405, 1.118, 1.425, 1.295 Å) compared with 1.838 ± 0.310 Å for paracetamol ([1.63, 2.13] Å; replicates: 2.285, 1.610, 1.650, 1.808 Å), a difference Δ(`AM404` − APAP) = −0.53 Å with zero replicate overlap ($\max$ `AM404` = 1.425 Å < $\min$ paracetamol = 1.610 Å; exact midrank-permutation Mann–Whitney U p = 0.0286, the lowest rank p-value attainable at n = 4 vs. n = 4; two-sided Welch p = 0.0341; untruncated 731.05-ns sensitivity: 1.349 ± 0.161 Å vs. 1.855 ± 0.372 Å, zero overlap, MWU p = 0.0286). Whole-catalytic-domain superposition (450-aa Cα core) shifted in the same −0.43-Å direction

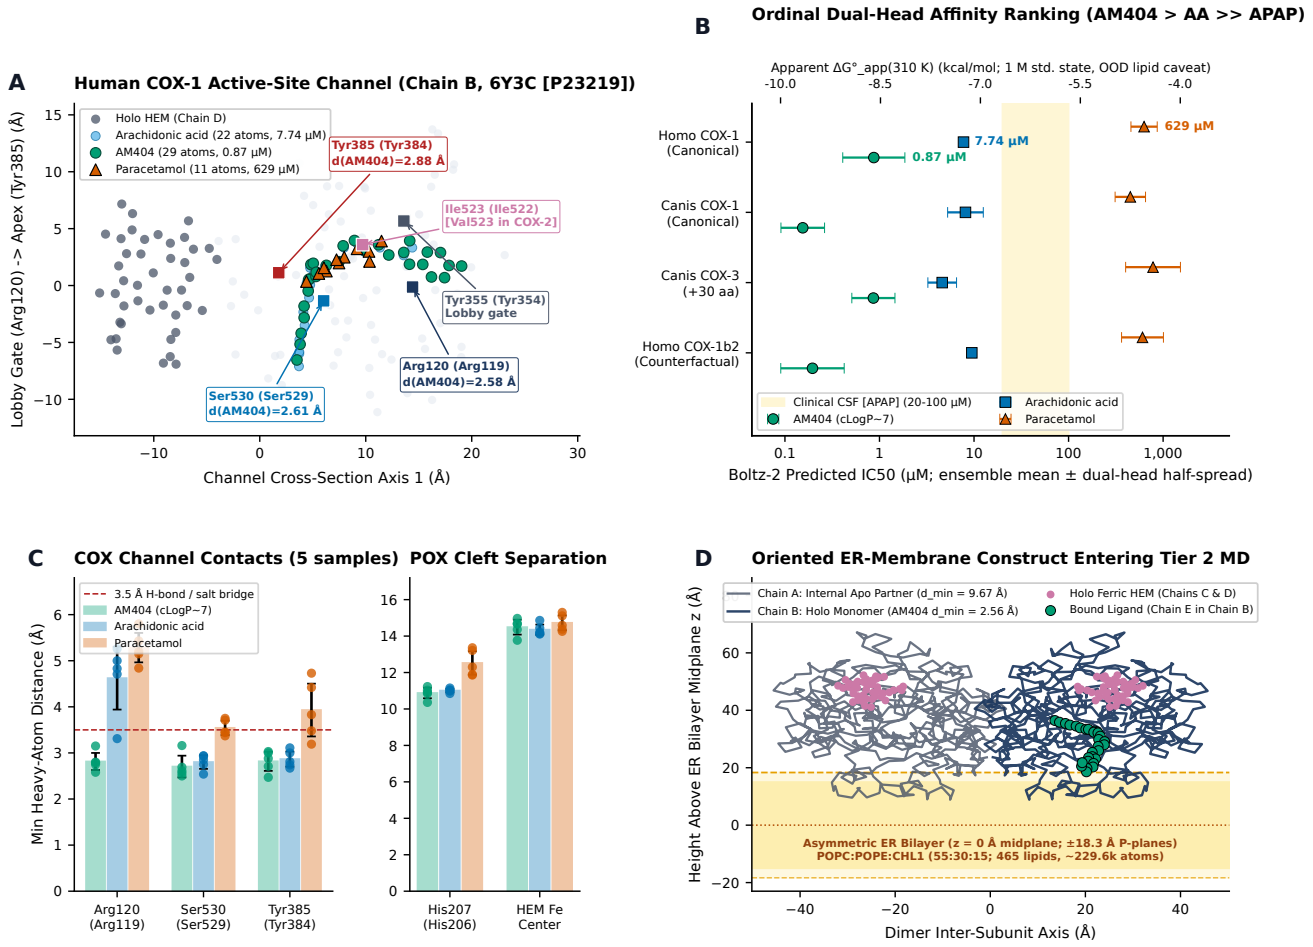


Figure 4. Active-site channel engagement of canonical human COX-1 by the brain paracetamol metabolite `AM404` versus arachidonic acid and parent paracetamol. (A) Catalytic-core Kabsch overlay of Chain B active-site coordinates from `human_cox1_homodimer_{am404,arachidonic_acid,paracetamol}_model_0.pdb`, labeled with dual `6Y3C` (`P23219`) residue numbers. **(B)** Dual-head `Boltz-2` predicted IC_{50} (± inter-head half-spread) and apparent ΔG°_{app}(310 K) across full-length dimers, shaded with clinical human CSF paracetamol levels (20–100 μM across oral/IV regimens). **(C)** Split distance comparison across 5 within-run samples showing active-site COX channel contacts (`Arg120`, `Ser530`, `Tyr385`; left, 0–6.8 Å) and POX cleft separation (`His207`, `HEM Fe`; right, 0–17.2 Å). **(D)** Oriented ER-membrane dimer architecture of mature `human_cox1_homodimer_am404` (Ala24–Leu599, 465 lipids, 229,430–229,766 particles) relative to the $z = 0$ Å ER bilayer midplane, showing amphipathic MBD helices (aa 73–116) inserted into the $+Z$ luminal leaflet ($z \approx +11.1$ Å) and holo Chain B paired with internal `apo` partner Chain A.

(1.419 ± 0.181 Å vs. 1.845 ± 0.442 Å; MWU $p = 0.1143$, Welch $p = 0.1496$), consistent with peripheral helix breathing diluting local active-site motion.

Exploratory active-site hydration and contact metrics across 13 secondary observables (with Benjamini–Hochberg FDR q-values reported in Table 4) were consistent with a steric and hydration divergence inside Chain B (Fig. 5B–D). Whereas `AM404`-bound Chain B averaged 4.70 ± 0.76 channel waters (vs. 10.65 ± 3.30 in its internal `apo` Chain A), paracetamol-bound Chain B retained 9.78 ± 0.67 channel waters ($\Delta = -5.08$ waters; exact MWU $p = 0.0286$ [$q_{\mathrm{MWU}} = 0.0531$], Welch $p < 0.001$ [$q_{\mathrm{Welch}} = 0.0007$], zero overlap), virtually matching unliganded `apo` Chain A (9.73 ± 1.22 waters; Fig. 5C). Because channel waters are counted inside a 5.0-Å cylinder around the eight channel-lining residues with a ≥ 8 heavy-atom burial filter, this -5.08-water reduction reflects both physical solvent expulsion and steric volume displacement by 29-heavy-atom `AM404` relative to 11-atom paracetamol. Geometrically, `AM404` anchored its 4-hydroxyphenylamide headgroup at the `Arg120` (`Arg119`) lobby entrance ($z = +27.70 \pm 1.08$ Å above the $z = 0$ Å ER midplane) while extending its arachidonoyl tail upward (COM $z = +34.18 \pm 1.14$ Å) to within 3.73 ± 0.30 Å of catalytic `Tyr385 OH`, remaining at a hard ceiling on strict ≤ 4.5-Å contact (0.9995 ± 0.0006, replicates $1.000, 0.999, 0.999, 1.000$; max bout $= 66.29 \pm 11.62$ ns). Conversely, paracetamol floated 9.09 Å

Table 4. Pre-registered primary endpoints, exploratory active-site mechanics, and cross-system internal `apo` Chain A falsification controls across all-atom ER-membrane MD ($n = 4$ independent replicates per arm; 1,506 common frames = 75.30 ns/replicate). Values report replicate mean ± SD (10,000-resample percentile bootstrap 95% CI in brackets; replicates `01-04` in parentheses). Exact two-sided Mann–Whitney U (MWU) p-values use full $\binom{8}{4} = 70$ midrank permutation ($p_{\min} = 2/70 = 0.0286$); exploratory rows include Benjamini–Hochberg FDR q-values across all 13 exploratory tests.

Observable	Paracetamol ($n = 4$)	`AM404` ($n = 4$)	Welch p (q_{BH})	MWU p (q_{BH})	Physical / Control Verdict
Pre-registered primary endpoints (75.30 ns common window + untruncated 731.05 ns sensitivity)					
Hysteresis bound (5.0/7.5 Å)	1.000 ± 0.000 [1.00, 1.00] (1, 1, 1, 1)	1.000 ± 0.000 [1.00, 1.00] (1, 1, 1, 1)	1.0000	1.0000	Ceiling (4/4 KM censored at 75.3 ns)
Channel-pocket RMSD (Å, 8 Cα)	1.838 ± 0.310 [1.63, 2.13] (2.29, 1.61, 1.65, 1.81)	1.311 ± 0.141 [1.19, 1.42] (1.41, 1.12, 1.43, 1.30)	**0.0341**	**0.0286**	Zero overlap (Δ = −0.53 Å)
Catalytic-core RMSD (Å, 450 Cα)	1.845 ± 0.442 [1.54, 2.26] (2.47, 1.45, 1.64, 1.82)	1.419 ± 0.181 [1.28, 1.59] (1.66, 1.24, 1.35, 1.43)	0.1496	0.1143	Diluted by 450-aa helix breathing
Untruncated pocket RMSD (Å)	1.855 ± 0.372 [1.62, 2.21] (2.40, 1.59, 1.65, 1.78)	1.349 ± 0.161 [1.20, 1.46] (1.43, 1.12, 1.49, 1.36)	0.0650	**0.0286**	Zero overlap across 75.3–99.7 ns
Exploratory active-site hydration, contact persistence, and bilayer z-geometry (75.30 ns common window; 13-metric BH-FDR)					
Holo Chain B COX waters	9.78 ± 0.67 [9.20, 10.35] (10.3, 9.26, 9.13, 10.4)	4.70 ± 0.76 [3.94, 5.16] (4.96, 5.03, 3.58, 5.23)	< **0.001** (0.0007)	**0.0286** (0.0531)	Δ = −5.08 waters (0 overlap; +steric excl.)
Strict contact (≤ 4.5 Å)	0.827 ± 0.191 [0.66, 0.97] (0.94, 0.81, 1.00, 0.57)	0.9995 ± 0.0006 [1.00, 1.00] (1.0, 1.0, 1.0, 1.0)	n/a (ceiling)	n/a (ceiling)	Δ = +0.173 (`AM404` at 1.0 ceiling)
Max ≤ 4.5 Å bout (ns)	16.69 ± 13.41 [5.68, 27.7] (23.6, 8.90, 31.8, 2.45)	66.29 ± 11.62 [57.0, 75.3] (75.3, 51.0, 63.6, 75.3)	**0.0015** (0.0065)	**0.0286** (0.0531)	4.0× longer contiguous engagement
Headgroup z vs. midplane (Å)	+36.79 ± 0.69 [36.3, 37.4] (37.8, 36.5, 36.7, 36.2)	+27.70 ± 1.08 [26.8, 28.6] (28.5, 28.6, 27.4, 26.3)	< **0.001** (< 0.0001)	**0.0286** (0.0531)	Anchors 9.09 Å lower at `Arg120` gate
Ligand COM z vs. midplane (Å)	+36.71 ± 0.67 [36.2, 37.4] (37.7, 36.4, 36.6, 36.1)	+34.18 ± 1.14 [33.3, 35.1] (34.8, 35.3, 33.8, 32.8)	**0.0134** (0.0421)	**0.0286** (0.0531)	Spans 6.5 Å vertically up channel
Min d(`Tyr385 OH`) (Å)	5.09 ± 0.72 [4.45, 5.64] (4.99, 5.35, 4.15, 5.85)	3.73 ± 0.30 [3.49, 3.97] (3.35, 4.04, 3.89, 3.62)	**0.0245** (0.0531)	**0.0286** (0.0531)	ω-tail 1.36 Å closer to radical Tyr
Holo `Arg120-Glu524` (≤ 4.0 Å)	0.993 ± 0.014 [0.98, 1.00] (1.00, 1.00, 0.97, 1.00)	1.000 ± 0.000 [1.00, 1.00] (1.00, 1.00, 1.00, 1.00)	0.3910 (0.4307)	1.0000 (1.0000)	Lobby salt bridge remains locked
EGF buried ΔSASA (Å^2)	976 ± 66 [914, 1,033] (891, 1051, 983, 979)	983 ± 44 [947, 1,019] (933, 1012, 961, 1027)	0.8581 (0.8581)	1.0000 (1.0000)	Matches crystal 932–1,017 Å^2
Pre-specified cross-system internal `apo` Chain A falsification check (4/4 show no detectable bias at $n = 4$)					
`Apo` Chain A Fe ↔ `Tyr385` (Å)	12.68 ± 0.34 [12.40, 12.95] (12.96, 12.28, 12.52, 12.95)	12.57 ± 0.11 [12.47, 12.66] (12.58, 12.42, 12.61, 12.69)	0.5905	0.8857	No detectable heme–apex bias
`Apo` Chain A COX waters	9.73 ± 1.22 [8.62, 10.66] (8.11, 9.64, 10.15, 11.01)	10.65 ± 3.30 [8.50, 13.73] (8.42, 10.13, 15.45, 8.58)	0.6300	1.0000	No detectable hydration bias
`Apo` Chain A corridor waters	0.36 ± 0.61 [0.01, 0.96] (0.02, 1.27, 0.01, 0.15)	0.23 ± 0.37 [0.04, 0.61] (0.04, 0.03, 0.80, 0.07)	0.7406	0.8857	No detectable corridor bias
`Apo` Chain A `Arg120-Glu524`	0.939 ± 0.092 [0.85, 1.00] (1.00, 1.00, 0.95, 0.80)	0.973 ± 0.023 [0.95, 0.99] (1.00, 0.95, 0.96, 0.99)	0.5243	0.8000	No detectable gate bias

higher inside the hydrated upper cavity (headgroup $z = +36.79 \pm 0.69$ Å, MWU $p = 0.0286$, Welch $p < 0.001$; $d_{\min}$[`Tyr385 OH`] $= 5.09 \pm 0.72$ Å, MWU $p = 0.0286$, Welch $p = 0.0245$), detaching beyond 4.5 Å in 17.4% of frames (contact fraction $= 0.827 \pm 0.191$, $\Delta = +0.173$, p-values omitted in Table 4 due to `AM404` ceiling saturation; max bout = 16.69±13.41 ns, MWU $p = 0.0286$, Welch $p = 0.0015$; mean uncensored bout = 3.50±5.81 ns). Across unliganded partner Chain A (which also absorbs the minor luminal leaflet asymmetry of 198 vs. 209 lipids between builds), all four pre-specified falsification checks yielded Welch $p = 0.5243$–0.7406 and MWU $p = 0.8000$–1.0000 (Table 4; Fig. 5C), showing no detectable cross-system bias at $n = 4$.

3 Discussion

3.1 Resolution of the Human COX-3 Hypothesis and Long-Read Foldback Artifacts

Twenty-four years after Intron-1 retention in canine *PTGS1* was proposed to explain central paracetamol pharmacology [6], multi-omic and structural evidence across human tissues supports retiring hypothetical human COX-3 entries (`COX-1b2/P23219-6`, `COX-1b3/P23219-5`) within explicit empirical bounds (≤ 9.81% in human brain; ≤ 0.115% pooled). At the genomic and RNA levels, human Intron 1 (94 nt ≡ 1 (mod 3), 62.8% GC, near-neutral interior `phyloP100way` = +0.233) lacks both the $3n$ reading-frame modularity and the 3′ `GCAG` splice-acceptor secondary-structure sequestration ($P_{\mathrm{pair}} = 0.293$–$0.346$ vs. canine 0.762–0.810) that enable constitutive canine Intron-1 retention. Frame-restoring 93-nt deletions (`COX-1b2`, `COX-1b3`) [8] are absent from `GRCh38` and population variation; −1 PRF lacks slippery heptamers or excess stability ($z = +1.62$) and hits an intronic stop at 21 aa; *PTGS1* Intron-1 ribosome density sits at rank 4/11 (36th percentile) of length-matched short first introns (≤ 500 nt); and PeptideAtlas detects 98.78% of mature `P23219` alongside 0/44 diagnostic COX-3 peptides. Beyond *PTGS1*, our pre-`TranscriptClean` ZMW forensics expose a general methodological pitfall for long-read transcriptomics: reverse-transcription template-switching foldbacks on GC-rich exons (72.4% GC on Exon 2) can force splice-aware aligners across short adjacent introns, where reference-guided error correctors [28] overwrite mismatching foldback sequence into synthetic intron retention or pseudo-rescued ORFs unless verified against raw CCS reads.

Equally informative is our second pillar: the structural outcome when Intron 1 *is* translated in canine COX-3. Holo dimeric `Boltz-2` co-folding shows that the +30-aa leader folds as a solvent-exposed exterior loop that makes no contact with the heme prosthetic group or the catalytic residues of the cyclooxygenase channel in any of 20 structures (closest approach 10.87 Å), leaving the 450-aa globular catalytic core unchanged (RMSD = 0.269 Å in `apo`; 0.211–0.307 Å across ligand states) and conferring no direct COX-channel paracetamol affinity advantage (782 μM vs. 450 μM). Why, then, did canine COX-3 exhibit apparent paracetamol sensitivity ($IC_{50} \approx 64$ μM) in Sf9 insect cell lysates [6]? Based on prostaglandin synthase ferryl-oxo redox kinetics [4, 5], we hypothesize that prepending +30 aa ahead of the ER signal peptide impairs co-translational ER translocation and maturation, lowering functional holoenzyme yield (consistent with the ∼ 5% relative specific activity reported in Sf9 lysates

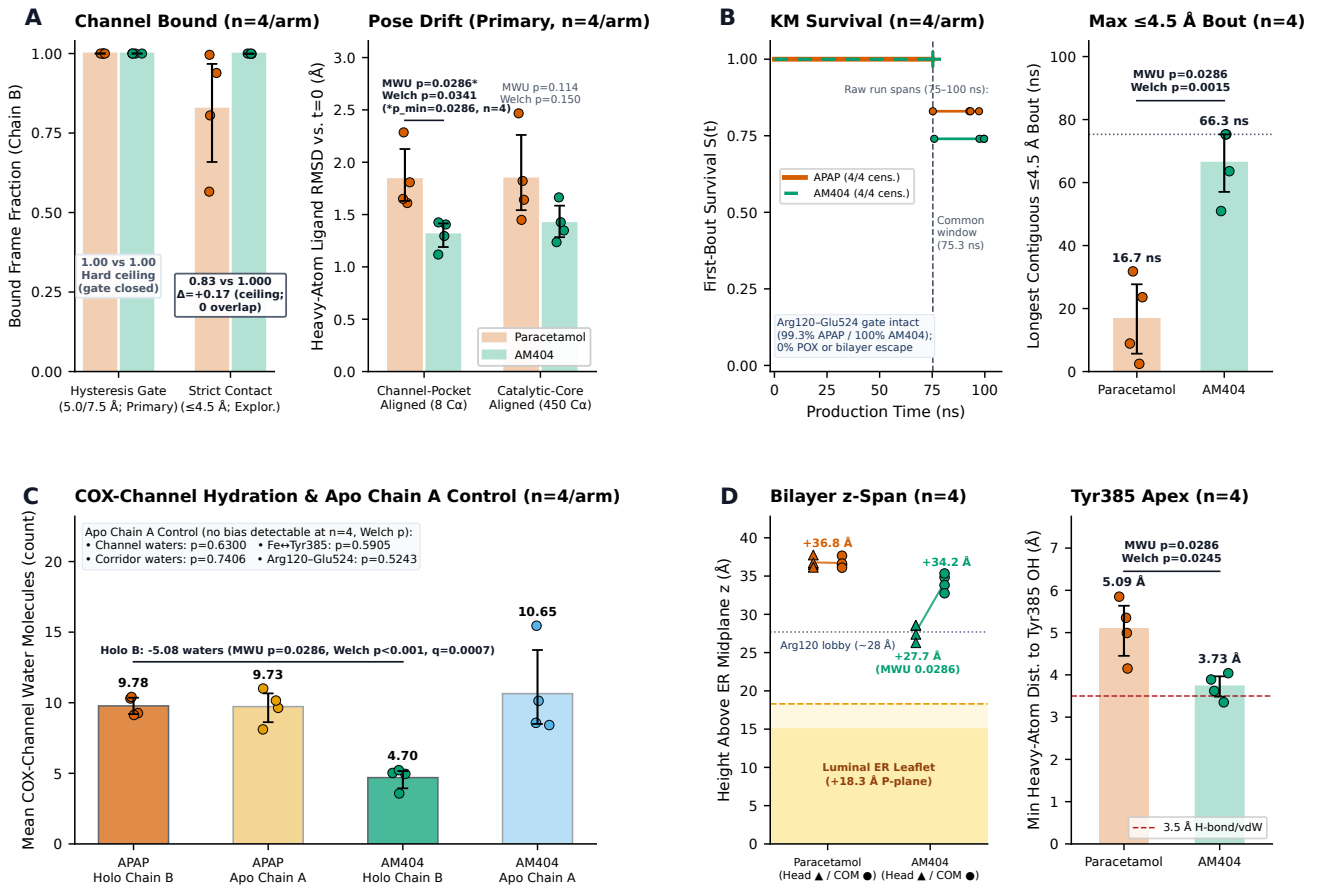


Figure 5. All-atom ER-membrane molecular dynamics (731.05 ns raw production; $n = 4$ replicates/arm truncated to 1,506 common frames = 75.30 ns/replicate) contrast active-site desolvation and channel-plugging by AM404 against upper-cavity tumbling of parent paracetamol. (A) Left: pre-registered hysteresis channel bound fraction (5.0/7.5 Å gate; 1.000 ± 0.000 hard ceiling in both arms) alongside exploratory strict active-site contact occupancy ($\leq$ 4.5 Å, AM404 at 1.000 ceiling, $\Delta = +0.17$). Right: pre-registered channel-pocket-aligned (8 Cα) and catalytic-core-aligned (450 Cα) heavy-atom ligand RMSD versus $t = 0$ (mean $\pm$ 95% bootstrap CI; $*p_{\min} = 0.0286$ at $n = 4$ vs. $n = 4$). **(B)** Left: right-censored Kaplan–Meier first-bout channel survival curves at $S(t) = 1.00$ across the 75.30-ns common window (+ markers indicate right-censoring; separate horizontal bars at $y = 0.83/0.74$ depict raw untruncated run spans of 75.3–99.7 ns). Right: longest contiguous $\leq$ 4.5-Å active-site contact bout per replicate. **(C)** COX-channel water occupancy in holo Chain B versus internal apo partner Chain A ($n = 4$/arm), annotated with Benjamini–Hochberg FDR q-value and all four pre-specified cross-system apo Chain A control p-values. **(D)** Left: vertical position z of ligand polar headgroups (▲) and centers of mass (●) above the $z = 0$ Å ER bilayer midplane. Right: replicate mean minimum heavy-atom distance to catalytic Tyr385 (Tyr384) OH.

[6]) and depressing autocrine PGG_2 hydroperoxide tone into the low-peroxide regime where paracetamol quenches the POX radical intermediate. This translocation–peroxide hypothesis remains a testable biochemical prediction rather than a directly measured result here.

3.2 Situating Canonical COX-1 Channel Dynamics Within Multi-Target Paracetamol Pharmacology

Replacing the COX-3 hypothesis requires cleanly separating two non-overlapping enzymatic sites on canonical cyclooxygenases, the **peroxidase (POX) heme cleft** and the **cyclooxygenase (COX) hydrophobic channel**, within paracetamol's broader central pharmacology [1, 3, 4]. Following a 1 g human dose, CSF paracetamol peaks at $\sim$ 20–26 μM orally and $\sim$ 39 μM intravenously [33] (up to $\sim$ 60 μM after 2 g IV propacetamol [34]). Experimental enzymology establishes that these concentrations are sufficient for reducing-cosubstrate quenching of the POX Fe^{IV}=O radical cation when ambient hydroperoxide tone is low (IC_{50} = 4.3 μM in endothelial cells and 33 μM

with glutathione peroxidase vs. 980–1,870 μM under high peroxide) [4, 5], a redox electron-transfer mechanism that neither static Boltz-2 affinity heads nor classical fixed-charge OpenMM force fields model. What our co-folding and 75.3-ns ER-membrane MD address instead is our third pillar: **non-redox steric occupation of the canonical human COX-1 substrate channel**. Even when kinetically retained behind a closed Arg120–Glu524 salt bridge on sub-100-ns timescales, 11-atom paracetamol leaves the COX channel as hydrated as unliganded apo Chain A (9.78 vs. 9.73 waters), exhibits higher local pose drift (1.84 Å), and detaches beyond 4.5 Å on nanosecond timescales, whereas 29-atom AM404 anchors its phenolamide headgroup at Arg120 ($z = +27.70$ Å), reduces channel water count (4.70 waters), and maintains tighter proximity to catalytic Tyr385 (3.73 Å; 1.31-Å pocket RMSD).

Three translational constraints temper direct extrapolation of AM404 COX-1 channel engagement to human analgesia. First, endogenous AM404 concentrations in human brain parenchyma or CSF after therapeutic 1 g paracetamol dosing have not yet been quantified; rodent brain AM404 measurements (30–300 mg/kg systemic dosing) detected picomolar-to-low-nanomolar bulk tissue levels [10], meaning local membrane enrichment near sites of synthesis would be required to approach effective COX-channel occupancy. Second, hepatic and CNS deacetylation of paracetamol to *p*-aminophenol is a quantitatively minor metabolic route with marked species differences in carboxylesterase/deacetylase activity. Third, whereas FAAH is expressed predominantly in postsynaptic CNS neurons, our GTEx snRNA-seq and ENCODE long-read analyses show that parenchymal human CNS *PTGS1* resides chiefly in microglia, perivascular macrophages, and cerebrovascular endothelial cells, implying that AM404 engagement of microglial or endothelial COX-1 in vivo would require intercellular transcellular lipid transport across neural–glial–vascular compartments.

Furthermore, genetic and pharmacological loss-of-function studies demonstrate that cyclooxygenase inhibition alone does not account for central paracetamol antinociception, which critically requires brain TRPV1 channels [11, 12], spinal dorsal-horn TRPA1 activation by electrophilic NAPQI and *p*-benzoquinone metabolites [13], endocannabinoid/FAAH modulation [14], and indirect CB1 plus bulbospinal 5-HT descending inhibition [15, 16]. Canonical COX-1 channel engagement by AM404 and low-peroxide COX-1/COX-2 POX quenching are therefore consistent with complementary prostanoid-lowering components within a multi-target central network rather than a standalone mechanism.

3.3 Limitations and Scope Calibration

Five methodological boundaries calibrate our conclusions:

1. **Human brain single-molecule depth, tryptic blind spots, and DDA MS/MS sensitivity:** Pooled ENCODE long-read sequencing ($\leq 0.115\%$) is dominated (79.9%) by HL-60 macrophages, leaving post-mortem human brain parenchyma (0/29 reads) bounded at $\leq 9.81\%$ (one-sided 95%). In proteomics, hypothetical COX-1b2 (c.7+66delG) lacks isoform-unique 7–35-aa tryptic peptides (69-aa Lys/Arg-free span, covered only by pan-intron-1 ECDPGAR/MSRECDPGAR), PeptideAtlas sensitivity is bounded by tryptic flyability and CNS representation, and our 7 public cortical/BBB DDA MGF files failed to recover canonical P23219 or P35354 positive controls at 1% FDR ($\sim$ 160–830$\times$ lower CNS vs. platelet abundance), requiring targeted PRM/DIA on enriched human microglia and microvessels for direct CNS protein bounds.

2. **Post-transcriptional fate ambiguity:** Steady-state Ribo-seq and MS/MS non-detection cannot distinguish nuclear intron detention [24] from cytoplasmic pioneer-round NMD [23]; fractionated nuclear/cytoplasmic long-read sequencing following UPF1 depletion would be required to partition their relative kinetic contributions.

3. **Boltz-2 sampling and lipophilic affinity calibration:** Tier 1 co-folding used a single random seed and shared Stage-1 MSA per chain, quantifying within-run diffusion denoising spread rather than multi-seed variance, and AM404's high lipophilicity (cLogP ≈ 7) restricts affinity predictions to ordinal ranking (AM404 > AA $\gg$ paracetamol).

4. **Unbiased MD timescale, water model, axial restraint, and replicate power:** Four unbiased replicates per arm (75.30-ns common window; 731.05 ns raw production) initiated from a single predicted pose test short-timescale pose drift and pocket hydration rather than equilibrium ΔG_{bind} or macroscopic k_{off}; exact two-sided rank tests at $n = 4$ vs. $n = 4$ have a resolution floor of $p_{\min} = 0.0286$, low-viscosity TIP3P-FB water accelerates solvent-exchange kinetics relative to OPC, unliganded Chain A (198 vs. 209 luminal lipids across builds) is an internal half-of-sites partner rather than an independent apo:apo homodimer, and the

stiff `His388-Fe` harmonic restraint (k = 25,000 kJ mol^{-1} nm^{-2}, $\sim$ 0.01-nm thermal excursions) precludes heme/POX breathing claims.

5. **COX-1 versus COX-2 (`Val523`) and non-redox scope:** All structural co-folding and classical MD here modeled non-redox COX-1 (`Ile523`/`Ile522`), because COX-3 was proposed as a *PTGS1* splice variant. Inducible central prostaglandin E_2 synthesis during fever and inflammation depends substantially on COX-2 (`P35354`, carrying the larger `Val523` [`Val509`] side pocket) and POX ferryl-oxo redox state [3, 4, 10].

4 Methods

4.1 Genomic Coordinates, Variant Screening, and Proteogenomic Database Construction

Genomic sequences and exon–intron boundaries for human *PTGS1* (`ENSG00000095303.15`, MANE Select `ENST00000362012.7`, and Intron-1-retained `ENST00000645132.1`, `ENST00000643576.1`), canine *PTGS1* (`ENSCAFG00845009693` on `ROS_Cfam_1.0` and GenBank `AF535138.1`), and murine *Ptgs1* (`ENSMUST00000028329` on `GRCm39`) were retrieved via Ensembl REST v113 and verified for exact splice-donor/acceptor dinucleotides and phase modular arithmetic. Human germline and somatic variants across `chr9:122,371,092-122,371,185` were queried via Ensembl variation endpoints integrating dbSNP and gnomAD. All 27 deterministic translation products (22 human non-canonical candidates spanning +1 native Intron-1 retention, `COX-1b2 c.7+66delG`, `COX-1b3 c.7+44delC`, cryptic splice shifts, $-1/+1$ PRF, and downstream `AUG`/`CUG` initiation, plus canonical/canine controls) and 17 initiator-methionine-excised forms were compiled with UniProtKB/Swiss-Prot human canonical proteome `UP000005640` (20,414 entries) into a non-redundant target FASTA (20,453 targets; 44 unique tryptic diagnostic COX-3 peptides of 7–35 aa).

4.2 GTEx v10 Transcriptomics and ENCODE Single-Molecule Long-Read Forensics

Gene-level, transcript-level (`GENCODE v39`), exon-read, junction-read, and single-nucleus expression profiles across 54 human tissues were retrieved from the GTEx Portal API v2 (`gtex_v10`). For single-molecule auditing, 57 unique released GRCh38 BAM alignments (40 PacBio HiFi Iso-Seq and 17 Oxford Nanopore) spanning adult DLPFC, developing brain, spinal cord, purified neural/glial cells, primary vascular endothelium, HL-60 macrophages, and colon/spleen controls were queried over `chr9:122,370,000-122,397,000` via HTTP byte-range BAM slicing (`pysam`). Primary non-supplementary alignments spanning $\leq$ 122,371,086 to $\geq$ 122,371,191 were classified by exact CIGAR splice operations. Every PacBio read aligning across Intron 1 was cross-examined by read name against paired unfiltered pre-`TranscriptClean` CCS BAMs (`ENCFF899UPY`, `ENCFF863NLH`). Exact binomial upper bounds on mature Intron-1 retention ($0/n$) were computed via the Clopper–Pearson beta-quantile identity $p_{\text{upper,1-sided}} = 1 - \alpha^{1/n}$ ($\alpha = 0.05$) and two-sided $\alpha/2 = 0.025$.

4.3 Sense-RNA Folding, EJC-NMD Mapping, and Matched-Gene Ribosome Profiling

Sense-strand RNA folding (converting genomic `T` to `U` without reverse-complementing) was computed at 37 °C using `ViennaRNA` (`RNA.fold_compound`, Turner 2004 parameters) [22] to obtain MFE structures, McCaskill partition-function base-pairing matrices, and 200× Altschul–Erickson dinucleotide-shuffled MFE z-scores, alongside UCSC `phyloP100way` base-wise conservation scores. Remote BigWig ribosome-profiling tracks (`UCSC_GWIPSviz_Global_RiboSeq` [29], primary human monocytes `Su15_All_RiboCov_track` [30], and `GSE213204` comprising 2 `THP-1` monocytic and 4 differentiated `M0/M1/M2` macrophage tracks) were streamed via `pyBigWig`, recording sparse `NaN` base counts (n_{nodata}), covered-base means, and all-base means (`NaN` $\rightarrow$ 0). Empirical background splicing depletion was calibrated across all $n = 44$ multi-exon MANE Select genes in `chr9:118,000,000-128,000,000` with Exon 2 $\text{mean}_{\text{all}} \geq 5.0$ on `UCSC_GWIPSviz_Global_RiboSeq`, both overall and stratified to short first introns ($\leq$ 500 nt, $n = 11$; $\leq$ 1,000 nt, $n = 15$) to control for inverse length scaling over sparse BigWig intervals.

4.4 PeptideAtlas Audit and Target-Decoy Orbitrap MS/MS Re-Searching

Global Human PeptideAtlas (`2026-01`; 5,367,334 distinct peptides), Human Brain PeptideAtlas, PeptideAtlas Ribo-Seq ncORFs, and ProteomicsDB tissue $\log_{10}$(iBAQ) intensities were audited by exact and Aho–Corasick substring matching against all 44 diagnostic COX-3 peptides and 13,643 non-canonical 6–50-aa substrings. Seven Orbitrap MS/MS MGF files from human blood–brain barrier cerebral microvessels and matched cortex (`PXD018602`; 5 fractions, Thermo Orbitrap Fusion FTMS1 / linear ion-trap ITMS2) and human frontal cortex (`PXD000004`; 2 fractions, Thermo Q Exactive HCD) were searched using `Sage v0.14.6` [32] against the 20,453-entry target database with internal reverse decoys, precursor tolerance ±20 ppm, fragment tolerance ±0.5 Da (`PXD018602` linear ion-trap ITMS2) or ±20 ppm (`PXD000004` Orbitrap HCD), isotope errors −1..2, up to 2 missed tryptic cleavages, fixed carbamidomethyl-`Cys`, variable `Met` oxidation and protein N-terminal acetylation, and picked target-decoy FDR filtering at 1%.

4.5 Tier 1 Dimeric Holo Co-Folding (`Boltz-2`)

Twenty dimeric holo YAML manifests (5 assemblies: `canine_cox1_homodimer`, `canine_cox3_homodimer`, `human_cox1_homodimer`, `human_cox1b2_homodimer`, and `human_cox1_cox1b1_heterodimer`; crossed with `apo`, `arachidonic_acid`, `paracetamol`, and `am404`; 2× ferric `HEM` per full-length dimer and 1× `HEM` on heterodimer) were co-folded on an NVIDIA A100-SXM4-80GB GPU using `boltz==2.2.1` (`cuEquivariance==0.8.1`, `recycling_steps=10`, `diffusion_samples=5`, `diffusion_samples_affinity=5`, `-use_potentials`, without explicit `-use_templates`) [17]. Stage-1 ColabFold MSAs were generated on the five `apo` manifests and reused across liganded states under a single seed. Inter-subunit buried ΔSASA (Shrake–Rupley), catalytic-core Kabsch Cα RMSD against `1EQG`/`6Y3C`, heme coordination distances, and dual-head $\log_{10}(\mathrm{IC}_{50}/\mu\mathrm{M})$ predictions were extracted across all 100 PDBs.

4.6 Tier 2 All-Atom ER-Membrane Molecular Dynamics (`OpenMM 8.4`)

Mature human COX-1 homodimers (Ala24–Leu599 per chain, excising Met1–Leu23 signal peptides) from `human_cox1_homodimer_paracetamol_model_0.pdb` and `human_cox1_homodimer_am404_model_0.pdb` were embedded monotopically via their amphipathic MBD helices (aa 73–116, inserted at $z \approx +11.1$ Å relative to the $z = 0$ Å bilayer midplane) into the $+Z$ luminal leaflet of an asymmetric `POPC:POPE:CHL1` (55:30:15) ER bilayer using a custom tile-and-graft monotopic membrane builder ($14.27 \times 14.98 \times 12.47$–$12.54$ nm^3 periodic box; 454 lipids [198 luminal / 256 cytosolic] for paracetamol and 465 lipids [209 luminal / 256 cytosolic] for `AM404`, reflecting minor luminal lipid-exclusion differences around the starting dimer footprint) and solvated in 0.15 M NaCl `TIP3P-FB` water (227,304–228,002 particles for paracetamol and 229,430–229,766 particles for `AM404` across production replicates `01-04`; 227,244 and 229,406 in Gate-0 validation) [35, 36]. Force fields comprised `Amber ff19SB` (protein; originally benchmarked with `OPC`, whereas `TIP3P-FB` was selected for `Lipid21`/`OpenFF` compatibility and yields faster bulk water self-diffusion that can accelerate pocket hydration exchange) [36, 37], `Lipid21` (phospholipids and cholesterol) [38], `OpenFF 2.2.0` (Sage) with conformation-independent `NAGL openff-gnn-am1bcc-0.1.0-rc.3` charges (paracetamol and `AM404`) [39], and a 73-atom all-atom ferric protoporphyrin IX (`HEM`, net charge $-1.00000000\,e$ [40]) tethered to proximal `His388` (`His387`) Nε2 via an `1EQG`-referenced harmonic axial restraint ($r_0 = 0.240$ nm, $k = 25{,}000$ kJ mol^{-1} nm^{-2}, restricting axial thermal fluctuations to ~ 0.01 nm and precluding heme/POX breathing claims, plus four 90° Nε2–Fe–$\mathrm{N_{pyrrole}}$ harmonic angle restraints at $k_\theta = 209.2$ kJ mol^{-1} rad^{-2} and 1–2/1–3 nonbonded exclusions). Ten crystallographic disulfides (5 per monomer) and explicit histidine tautomers (`His207/386/388 = HID`, `His446 = HIE`; `P23219`: `His206/385/387`, `His445`) were enforced. Four independent production replicates per arm (731.05 ns raw 4.0-fs HMR production across 8 trajectories [75.30–99.65 ns/replicate; 50-ps frame stride], or 775.05 ns total MD including a 7-stage equilibration protocol comprising L-BFGS minimization up to 5,000 iterations and 6 MD stages totaling 5.50 ns/replicate: 0.20-ns 1-fs NVT heating 10 → 310.15 K, four progressively unrestrained semi-isotropic NPT stages of 0.10, 0.40, 0.80, and 1.50 ns, and 2.50 ns of unrestrained 4-fs HMR NPT pre-production; unique Langevin seeds, 310.15 K, 1 bar semi-isotropic `MonteCarloMembraneBarostat`, PME 1.0 nm cutoff) were executed on 2 × (4 × NVIDIA RTX 5090) nodes and analyzed over equal-length common-truncated trajectories (1,506 frames = 75.30 ns/replicate; 602.40 ns combined) using `MDAnalysis` [41]. Between-arm replicate comparisons ($n = 4$ vs. $n = 4$) used two-sided Welch's t-tests and exact two-sided midrank-permutation Mann–Whitney U tests across all $\binom{8}{4} = 70$ group assignments, with Benjamini–Hochberg FDR q-values computed across all 13 exploratory metrics.

Data and Code Availability

All analysis pipelines, `Boltz-2` co-folding manifests, OpenMM ER-membrane simulation scripts, figure-generation scripts, and machine-readable JSON summaries were executed exclusively on publicly archived datasets (Ensembl v113, GTEx v10, ENCODE Portal, UCSC GWIPS-viz, NCBI GEO `GSE213204`, PeptideAtlas `2026-01`, ProteomicsDB, PRIDE `PXD018602/PXD000004`, and RCSB PDB `1EQG/1DIY/3N8Z/6Y3C`) and are included in the project repository.

Competing Interests

The author declares no competing financial or non-financial interests.

Glossary of Key Genetic, Biophysical, and Pharmacological Terms

Term / Shorthand	Plain-Scientific Definition and Context in This Study
Genomics & RNA Splicing	
PTGS1* / *PTGS2	Genes encoding prostaglandin G/H synthase 1 (constitutive COX-1, UniProt `P23219`) and inducible synthase 2 (COX-2, `P35354`), which convert arachidonic acid to prostaglandin H_2.
Intron 1 & Reading-Frame Modularity (mod 3)	First non-coding segment interrupting codon 3 of *PTGS1*. Canine Intron 1 (90 nt $\equiv 0 \pmod 3$) inserts 30 amino acids without shifting downstream codons, whereas human Intron 1 (94 nt $\equiv 1 \pmod 3$) shifts Exons 2–3 into the +1 frame and hits an early stop.
`COX-3`, `COX-1b1`, `COX-1b2/3`	Canine `COX-3` (`Q8HZR1-2`) is the 633-aa dog enzyme retaining 90-nt Intron 1. Human `COX-1b1` (`Q3HY30`) is the native 94-nt-retained, frameshifted 82-aa fragment; `COX-1b2` (`P23219-6`) and `COX-1b3` (`P23219-5`) are hypothetical 93-nt single-base-deleted cDNAs absent in `GRCh38`/gnomAD.
MANE Select	Matched Annotation from NCBI and EMBL-EBI defining one canonical reference transcript per human protein-coding gene (`ENST00000362012.7` for *PTGS1*).
`ViennaRNA` P_{pair} & **Dinucleotide MFE z-Score**	Partition-function base-pairing probability (P_{pair}) at 37 °C and minimum free energy (MFE) stability standardized against 200× dinucleotide-shuffled controls; tests whether pre-mRNA hairpins sequester splice sites or stimulate frameshifting.
Programmed Ribosomal Frameshifting (PRF)	Regulated −1 or +1 slippage of a translating ribosome on a slippery heptamer (`X XXY YYZ`) backed by a stable RNA structure; absent in human *PTGS1* Intron 1 (where a −1 shift terminates at 21 aa).
`phyloP100way`	Base-resolution evolutionary conservation score across 100 vertebrates; positive values mark purifying selection (coding exons/splice donors) and near-zero values mark neutral drift (human Intron-1 interior: +0.233).
PacBio HiFi CCS (ZMW) & Oxford Nanopore (ONT)	Single-molecule long-read RNA/cDNA sequencing platforms reading entire multi-kilobase transcripts end to end without short-read assembly; zero-mode waveguide (ZMW) circular consensus sequencing (CCS) yields > 99.9% raw base accuracy.
RT Foldback Hairpin & `TranscriptClean` Artifact	Reverse-transcriptase template-switching self-priming on GC-rich RNA (72.4% GC Exon 2) that forces splice aligners across adjacent short introns, where reference-guided `TranscriptClean` overwrites mismatching foldbacks into synthetic intron retention.
Translation & Proteomics	
Exon-Junction Complex (EJC) & NMD (`PTC83`)	Protein complex deposited ∼ 20–24 nt upstream of spliced exon–exon junctions; premature termination codons (such as `COX-1b1 PTC83`) located ≥ 50–55 nt upstream of a remaining EJC trigger nonsense-mediated mRNA decay (NMD).
Detained Intron (DI)	Post-transcriptionally unspliced intron causing otherwise spliced polyadenylated transcripts to remain sequestered inside the cell nucleus away from cytoplasmic ribosomes.
GTEx EM Deconvolution	Expectation–maximization statistical allocation of short RNA-seq fragments across overlapping transcript isoforms across 54 human tissues; low-powered for minor splice variants at low total gene expression.
Ribosome Profiling (Ribo-seq)	Deep sequencing of ∼ 28–30-nt mRNA footprints protected inside actively translating ribosomes; benchmarked against length-matched short first introns (≤ 500 nt) to account for sparse BigWig coverage.
PeptideAtlas, DDA MS/MS (`Sage`), & iBAQ	Global repository of tandem mass-spectrometry (MS/MS) peptide identifications (5.37×10^6 peptides), data-dependent acquisition (DDA) spectrum searching (`Sage`) at 1% target-decoy false discovery rate (FDR), and intensity-based absolute quantification (iBAQ).

Continued on next page

Term / Shorthand	Plain-Scientific Definition and Context in This Study (continued)
Structural Modeling & Membrane MD	
Holo vs. Apo & Half-of-Sites Homodimer	*Holo* denotes enzyme bound to its prosthetic heme cofactor (and/or channel ligand), whereas *apo* denotes an unliganded active site; COX-1 operates as an asymmetric homodimer in which one monomer's COX channel binds ligand while its partner monomer remains `apo`.
`Boltz-2`, `pLDDT`, `AB_iptm`, & **ΔSASA**	All-atom biomolecular co-folding and dual-head affinity model; `pLDDT` (0–100) scores local residue confidence, `AB_iptm` (0–1) scores inter-chain dimer interface accuracy, and ΔSASA (Å^2) measures solvent-accessible surface area buried upon dimerization.
Signal Peptide, EGF Module, & MBD	N-terminal ER-targeting leader (Met1–Leu23, cleaved in mature `P23219`), disulfide-knotted epidermal growth factor-like dimerization module, and amphipathic membrane-binding helices (aa 73–116) anchoring COX-1 monotopically into one luminal ER leaflet.
Ferric Heme (`HEM`), POX Cleft, & COX Channel	Spatially distinct catalytic sites on each COX monomer: the solvent-exposed peroxidase (POX) cleft houses ferric protoporphyrin IX (Fe^{III}, coordinated by `His388`) and generates the Fe^{IV}=O ferryl-oxo radical that oxidizes `Tyr385` at the apex of the hydrophobic cyclooxygenase (COX) substrate channel.
`Arg120-Glu524`, `Tyr355`, `Tyr385`, `Ser530`, & `Ile523/Val523`	Key COX-channel landmarks (ovine/`6Y3C` numbering): `Arg120-Glu524` forms the lower lobby salt-bridge gate; `Tyr355` stabilizes the lobby H-bond network; `Tyr385` abstracts the substrate hydrogen radical; `Ser530` is the upper channel aspirin-acetylation residue; and bulky `Ile523` in COX-1 vs. smaller `Val523` (`Val509`) in COX-2 controls side-pocket volume.
`POPC:POPE:CHL1`, `OpenMM` 4-fs **HMR, & Pocket RMSD**	Physiological 55:30:15 phosphatidylcholine : phosphatidylethanolamine : cholesterol ER membrane; GPU molecular dynamics (`OpenMM 8.4`) using hydrogen mass repartitioning (HMR) for 4-fs integration steps; and heavy-atom ligand root-mean-square deviation (RMSD) after aligning the 8 COX-pocket Cα atoms.
Channel Desolvation & Unoccupied Pocket Volume	Expulsion of buried water molecules from the hydrophobic COX channel (via steric displacement and apolar tail packing) and contraction of water-accessible free cavity volume (V_{free}) upon tight lipid-conjugate binding.
Analgesic Pharmacology & Statistics	
Paracetamol (APAP) & Peroxide Tone	*N*-acetyl-*p*-aminophenol (acetaminophen); acts weakly as a direct COX-channel competitor but potently quenches the POX ferryl-oxo radical when cellular lipid hydroperoxide concentrations ("peroxide tone") are low.
`AM404` & `FAAH`	*N*-arachidonoylphenolamine (`AM404`), a lipophilic $C_{20:4}$ brain metabolite synthesized from paracetamol-derived *p*-aminophenol and arachidonic acid by neuronal fatty-acid amide hydrolase (`FAAH`).
`TRPV1`, `TRPA1 (NAPQI)`, `CB1`, & `5-HT`	Validated non-COX central pain-modulating pathways engaged by paracetamol metabolites: brain vanilloid receptor `TRPV1`, spinal ankyrin channel `TRPA1` (activated by electrophilic *N*-acetyl-*p*-benzoquinone imine, `NAPQI`, and *p*-benzoquinone), indirect cannabinoid `CB1` signaling, and bulbospinal serotonergic (`5-HT`) inhibition.
Clopper–Pearson Exact Binomial CI	Conservative, exact beta-distribution confidence bound on a binomial proportion (such as $0/n$ observed intron-retention reads) guaranteed to maintain at least 95% nominal coverage at small sample sizes.
Exact Permutation Mann–Whitney U **& BH-FDR** (q)	Distribution-free rank test evaluating all $\binom{8}{4} = 70$ exact replicate permutations ($p_{\min} = 0.0286$ at $n = 4$ vs. $n = 4$), paired with Benjamini–Hochberg false discovery rate (q) adjustment across exploratory endpoints.

References

[1] R. J. Flower and J. R. Vane. Inhibition of prostaglandin synthetase in brain explains the anti-pyretic activity of paracetamol (4-acetamidophenol). *Nature*, 240(5381):410–411, 1972. doi: 10.1038/240410a0.

[2] Bela Kis, James A. Snipes, and David W. Busija. Acetaminophen and the cyclooxygenase-3 puzzle: sorting out facts, fictions, and uncertainties. *Journal of Pharmacology and Experimental Therapeutics*, 315(1):1–7, 2005. doi: 10.1124/jpet.105.085431.

[3] Burkhard Hinz, Olga Cheremina, and Kay Brune. Acetaminophen (paracetamol) is a selective cyclooxygenase-2 inhibitor in man. *The FASEB Journal*, 22(2):383–390, 2008. doi: 10.1096/fj.07-8506com.

[4] Marc Ouellet and M. David Percival. Mechanism of acetaminophen inhibition of cyclooxygenase isoforms. *Archives of Biochemistry and Biophysics*, 387(2):273–280, 2001. doi: 10.1006/abbi.2000.2232.

[5] Olivier Boutaud, David M. Aronoff, Jacob H. Richardson, Lawrence J. Marnett, and John A. Oates. Determinants of the cellular specificity of acetaminophen as an inhibitor of prostaglandin H2 synthases. *Pro-*

ceedings of the National Academy of Sciences of the United States of America, 99(10):7130–7135, 2002. doi: 10.1073/pnas.102588199.

[6] N. V. Chandrasekharan, Hu Dai, K. Lamar Turepu Roos, Nathan K. Evanson, Joshua Tomsik, Terry S. Elton, and Daniel L. Simmons. COX-3, a cyclooxygenase-1 variant inhibited by acetaminophen and other analgesic/antipyretic drugs: Cloning, structure, and expression. *Proceedings of the National Academy of Sciences of the United States of America*, 99(21):13926–13931, 2002. doi: 10.1073/pnas.162468699.

[7] Joseph E. Dinchuk, Richard Q. Liu, and James M. Trzaskos. COX-3: in the wrong frame in mind. *Immunology Letters*, 86(1):121, 2003. doi: 10.1016/S0165-2478(02)00268-7.

[8] Ning Qin, Sui-Po Zhang, Tasha L. Reitz, Jay M. Mei, and Christopher M. Flores. Cloning, expression, and functional characterization of human cyclooxygenase-1 splicing variants: evidence for intron 1 retention. *Journal of Pharmacology and Experimental Therapeutics*, 315(3):1298–1305, 2005. doi: 10.1124/jpet.105.090944.

[9] B. Kis, J. A. Snipes, T. Gaspar, G. Lenzser, C. D. Tulbert, and D. W. Busija. Cloning of cyclooxygenase-1b (putative COX-3) in mouse. *Inflammation Research*, 55(7):274–278, 2006. doi: 10.1007/s00011-006-0083-z.

[10] Edward D. Högestätt, Bo A. G. Jönsson, Anna Ermund, David A. Andersson, Henrik Björk, Jessica P. Alexander, Benjamin F. Cravatt, Allan I. Basbaum, and Peter M. Zygmunt. Conversion of acetaminophen to the bioactive N-acylphenolamine AM404 via fatty acid amide hydrolase-dependent arachidonic acid conjugation in the nervous system. *Journal of Biological Chemistry*, 280(36):31405–31412, 2005. doi: 10.1074/jbc.M501489200.

[11] Peter M. Zygmunt, Huai-hu Chuang, Pouya Movahed, David Julius, and Edward D. Högestätt. The anandamide transport inhibitor AM404 activates vanilloid receptors. *European Journal of Pharmacology*, 396(1):39–42, 2000. doi: 10.1016/S0014-2999(00)00207-7.

[12] Christophe Mallet, David A. Barrière, Anna Ermund, Bo A. G. Jönsson, Alain Eschalier, Peter M. Zygmunt, and Edward D. Högestätt. TRPV1 in brain is involved in acetaminophen-induced antinociception. *PLoS ONE*, 5(9):e12748, 2010. doi: 10.1371/journal.pone.0012748.

[13] David A. Andersson, Clive Gentry, Lisa Alenmyr, Dan Killander, Simon E. Lewis, Anders S. Andersson, Bernard Bucher, Jean-Luc Galzi, Olov Sterner, Stuart Bevan, Edward D. Högestätt, and Peter M. Zygmunt. TRPA1 mediates spinal antinociception induced by acetaminophen and the cannabinoid Δ^9-tetrahydrocannabiorcol. *Nature Communications*, 2:551, 2011. doi: 10.1038/ncomms1559.

[14] Massimiliano Beltramo, Nephi Stella, Antonio Calignano, S. Y. Lin, Alexandros Makriyannis, and Daniele Piomelli. Functional role of high-affinity anandamide transport, as revealed by selective inhibition. *Science*, 277(5329):1094–1097, 1997. doi: 10.1126/science.277.5329.1094.

[15] Gisèle Pickering, Marie-Anne Loriot, François Libert, Alain Eschalier, Philippe Beaune, and Claude Dubray. Analgesic effect of acetaminophen in humans: first evidence of a central serotonergic mechanism. *Clinical Pharmacology & Therapeutics*, 79(4):371–378, 2006. doi: 10.1016/j.clpt.2005.12.307.

[16] Christophe Mallet, Laurence Daulhac, Jérôme Bonnefont, Catherine Ledent, Monique Etienne, Eric Chapuy, François Libert, and Alain Eschalier. Endocannabinoid and serotonergic systems are needed for acetaminophen-induced analgesia. *Pain*, 139(1):190–200, 2008. doi: 10.1016/j.pain.2008.03.030.

[17] Saro Passaro, Gabriele Corso, Jeremy Wohlwend, Mateo Reveiz, Stephan Thaler, Vignesh Ram Somnath, Noah Getz, Tally Portnoi, Julien Roy, Hannes Stark, David Kwabi-Addo, Dominique Beaini, Tommi Jaakkola, and Regina Barzilay. Boltz-2: Towards accurate and efficient binding affinity prediction. *bioRxiv*, page 2025.06.14.659707, 2025. doi: 10.1101/2025.06.14.659707.

[18] Barry S. Selinsky, Kushol Gupta, Caroline T. Sharkey, and Patrick J. Loll. Structural analysis of NSAID binding by prostaglandin H2 synthase: time-dependent and time-independent inhibitors elicit identical enzyme conformations. *Biochemistry*, 40(17):5172–5180, 2001. doi: 10.1021/bi010045s.

[19] Michael G. Malkowski, Stephan L. Ginell, William L. Smith, and R. Michael Garavito. The productive conformation of arachidonic acid bound to prostaglandin synthase. *Science*, 289(5486):1933–1937, 2000. doi: 10.1126/science.289.5486.1933.

[20] Ranjinder S. Sidhu, Ji Young Lee, Chong Yuan, and William L. Smith. Comparison of cyclooxygenase-1 crystal structures: cross-talk between monomers comprising cyclooxygenase-1 homodimers. *Biochemistry*, 49 (33):7069–7079, 2010. doi: 10.1021/bi1003298.

[21] Morena Miciaccia, Benny Danilo Belviso, Marianna Iaselli, Gino Cingolani, Savina Ferorelli, Maria Cappellari, Paola Loguercio Polosa, Maria Grazia Perrone, Rocco Caliandro, and Antonio Scilimati. Three-dimensional structure of human cyclooxygenase (hCOX)-1. *Scientific Reports*, 11:4312, 2021. doi: 10.1038/s41598-021-83438-z.

[22] Ronny Lorenz, Stephan H. Bernhart, Christian Höner zu Siederdissen, Hakim Tafer, Christoph Flamm, Peter F. Stadler, and Ivo L. Hofacker. ViennaRNA package 2.0. *Algorithms for Molecular Biology*, 6:26, 2011. doi: 10.1186/1748-7188-6-26.

[23] Eszter Nagy and Lynne E. Maquat. A rule for termination-codon position within intron-containing genes: when nonsense affects RNA abundance. *Trends in Biochemical Sciences*, 23(6):198–199, 1998. doi: 10.1016/S0968-0004(98)01208-0.

[24] Paul L. Boutz, Arjun Bhutkar, and Phillip A. Sharp. Detained introns are a novel, widespread class of post-transcriptionally spliced introns. *Genes & Development*, 29(1):63–80, 2015. doi: 10.1101/gad.247361.114.

[25] GTEx Consortium. The GTEx consortium atlas of genetic regulatory effects across human tissues. *Science*, 369(6509):1318–1330, 2020. doi: 10.1126/science.aaz1776.

[26] Petr Danecek, James K. Bonfield, Jennifer Liddle, John Marshall, Valeriu Ohan, Martin O. Pollard, Andrew Whitwham, Thomas Keane, Shane A. McCarthy, Robert M. Davies, and Heng Li. Twelve years of SAMtools and BCFtools. *GigaScience*, 10(2):giab008, 2021. doi: 10.1093/gigascience/giab008.

[27] Heng Li. Minimap2: pairwise alignment for nucleotide sequences. *Bioinformatics*, 34(18):3094–3100, 2018. doi: 10.1093/bioinformatics/bty191.

[28] Dana Wyman and Ali Mortazavi. TranscriptClean: variant-aware correction of indels, mismatches and splice junctions in long-read transcripts. *Bioinformatics*, 35(2):340–342, 2019. doi: 10.1093/bioinformatics/bty483.

[29] Audrey M. Michel, Gearoid Fox, Anmol M. Kiran, Christof De Bo, Patrick B. F. O'Connor, Stephen M. Heaphy, James P. A. Mullan, Claire A. Donohue, Desmond G. Higgins, and Pavel V. Baranov. GWIPS-viz: development of a ribo-seq genome browser. *Nucleic Acids Research*, 42(D1):D859–D864, 2014. doi: 10.1093/nar/gkt1035.

[30] Xiaodi Su, Yingpu Yu, Yi Zhong, Eugenia G. Giannopoulou, Xiaoyu Hu, Hui Liu, Justin R. Cross, Gunnar Rätsch, Charles M. Rice, and Lionel B. Ivashkiv. Interferon-γ regulates cellular metabolism and mRNA translation to potentiate macrophage activation. *Nature Immunology*, 16(8):838–849, 2015. doi: 10.1038/ni.3205.

[31] Eric W. Deutsch, Attila Csordas, Zhi Sun, Andrew Jarnuczak, Yasset Perez-Riverol, Tobias Ternent, David S. Campbell, Manuel Bernal-Llinares, Shujiro Okuda, Shin Kawano, Robert L. Moritz, Julie J. Carver, Mingxun Wang, Yasushi Ishihama, Nuno Bandeira, Henning Hermjakob, and Juan Antonio Vizcaíno. The ProteomeXchange consortium in 2017: supporting the cultural change in proteomics public data deposition. *Nucleic Acids Research*, 45(D1):D1100–D1106, 2017. doi: 10.1093/nar/gkw936.

[32] Michael R. Lazear. Sage: An open-source tool for fast proteomics searching and quantification at scale. *Journal of Proteome Research*, 22(11):3652–3659, 2023. doi: 10.1021/acs.jproteome.3c00486.

[33] Neil K. Singla, Cherri Parulan, Reina Samson, Julie Hutchinson, Rick Bushnell, Evelyn G. Beja, Robert Ang, and Mike A. Royal. Plasma and cerebrospinal fluid pharmacokinetic parameters after single-dose administration of intravenous, oral, or rectal acetaminophen. *Pain Practice*, 12(7):523–532, 2012. doi: 10.1111/j.1533-2500.2012.00556.x.

[34] B. Bannwarth, P. Netter, F. Lapicque, P. Gillet, P. Péré, E. Boccard, R. J. Royer, and A. Gaucher. Plasma and cerebrospinal fluid concentrations of paracetamol after a single intravenous dose of propacetamol. *British Journal of Clinical Pharmacology*, 34(1):79–81, 1992. doi: 10.1111/j.1365-2125.1992.tb04112.x.

[35] Peter Eastman, Raimondas Galvelis, Raúl P. Peláez, Ariana Rustenburg, Stephen E. Craab, John D. Chodera, Thomas E. Markland, and Phineus R. L. Markwick. OpenMM 8: Molecular dynamics simulation with machine learning potentials. *The Journal of Physical Chemistry B*, 128(1):109–116, 2024. doi: 10.1021/acs.jpcb.3c06662.

[36] Lee-Ping Wang, Todd J. Martinez, and Vijay S. Pande. Building force fields: an automatic, systematic, and reproducible approach. *The Journal of Physical Chemistry Letters*, 5(11):1885–1891, 2014. doi: 10.1021/jz500737m.

[37] Chuan Tian, Koushik Kasavajhala, Kellon A. A. Belfon, Lauren Raguette, He Huang, Angela N. Migues, John Bickel, Yuzhang Wang, Jorge Pincay, Qin Wu, and Carlos Simmerling. ff19SB: Amino-acid-specific protein backbone parameters trained against quantum mechanics energy surfaces in solution. *Journal of Chemical Theory and Computation*, 16(1):528–552, 2020. doi: 10.1021/acs.jctc.9b00591.

[38] Callum J. Dickson, Ross C. Walker, and Ian R. Gould. Lipid21: Complex lipid membrane simulations with AMBER. *Journal of Chemical Theory and Computation*, 18(3):1726–1736, 2022. doi: 10.1021/acs.jctc.1c01217.

[39] Simon Boothroyd, Pavan Kumar Behara, Owen C. Madin, David F. Hahn, Hyesu Jang, Vytautas Gapsys, Jeffrey R. Wagner, Joshua T. Horton, David L. Dotson, Matt W. Thompson, Jessica Maat, Trevor Gokey, Lee-Ping Wang, Daniel J. Cole, Michael K. Gilson, John D. Chodera, Christopher I. Bayly, Michael R. Shirts, and David L. Mobley. Development and benchmarking of Open Force Field 2.0.0: the Sage small molecule force field. *Journal of Chemical Theory and Computation*, 19(11):3251–3275, 2023. doi: 10.1021/acs.jctc.3c00039.

[40] Kiumars Shahrokh, Anita Orendt, Garold S. Yost, and Thomas E. Cheatham. Quantum mechanically derived AMBER-compatible heme parameters for various states of the cytochrome P450 catalytic cycle. *Journal of Computational Chemistry*, 33(2):119–133, 2012. doi: 10.1002/jcc.21922.

[41] Richard J. Gowers, Max Linke, Jonathan Barnoud, Tyler J. E. Reddy, Manuel N. Melo, Sean L. Seyler, Jan Domański, David L. Dotson, Sébastien Buchoux, Ian M. Kenney, and Oliver Beckstein. MDAnalysis: A Python package for the rapid analysis of molecular dynamics simulations. In *Proceedings of the 15th Python in Science Conference*, pages 98–105, 2016. doi: 10.25080/Majora-629e541a-00e.