

Mechanistic AI-QSP Modelling of Olipudase Alfa Therapy in Acid Sphingomyelinase Deficiency: Adult-to-Paediatric Bridging and Regulatory-Grade Credibility Assessment

Igor Goryanin^{1,2}, Irina Goryanin¹ (¹*IQANOVA Ltd, Edinburgh*; ²*School of Informatics, University of Edinburgh*).

Corresponding author: Prof. Igor Goryanin, IQANOVA Ltd, Edinburgh, Scotland, UK.

Manuscript prepared for submission to CPT: Pharmacometrics & Systems Pharmacology.

Abstract

Olipudase alfa (Xenpozyme) is the first enzyme replacement therapy approved for non-CNS manifestations of acid sphingomyelinase deficiency (ASMD). Adult and paediatric registration trials demonstrated substantial reductions in spleen and liver volumes, improvements in DLCO and platelet counts, and large reductions in lyso-sphingomyelin and ceramide; however, no mechanistic model has yet been published that simultaneously calibrates adult and paediatric responses, projects long-term efficacy beyond the 2-year follow-up, and identifies an ADA-positive sub-population. We present an AI-QSP model in SBML L3v2 with 11 species, 53 parameters and 27 fortnightly dosing events, calibrated against four clinical sources and one regulatory Product Information label using a 14-stage pipeline (IQANOVA ATLAS Master Pipeline v2.0). The model achieves 96 % calibration PASS and 83 % holdout PASS rates, passes all ten ASME V&V40 credibility gates, and yields pooled Bayesian meta-analytic estimates inside the 94 % HDI for every analysed endpoint. Combination simulation shows the system is operating at saturation at label dose, supporting 3 mg/kg q2w as the minimum effective dose. The model is fit-for-purpose for population-level inference and provides a regulatory-grade scaffold for future ASMD trials. Code and SBML are open-source.

1. Introduction

Acid sphingomyelinase deficiency (ASMD; Niemann-Pick disease types A and B) is a rare inherited lysosomal storage disorder caused by biallelic loss-of-function mutations in SMPD1 [R3, R34]. The chronic visceral form (ASMD type B) presents with progressive hepatosplenomegaly, interstitial lung disease, dyslipidaemia, and thrombocytopenia. Olipudase alfa is the first approved enzyme replacement therapy, demonstrating efficacy in the ASCEND adult Phase II/III trial [R1] and the ASCEND-Peds paediatric trial [R2]. The European Medicines Agency Product Information of 2024 [R9] and the ASCEND-OLE 4-year follow-up [R31] provide longitudinal evidence for sustained efficacy, but no published model jointly addresses adult-to-paediatric bridging, long-term projection, and immunogenicity stratification. Here we present a regulatory-grade quantitative systems pharmacology (QSP) model that fills this gap.

2. Methods

2.1 Model architecture

The model has six layers. Layer 1 captures pharmacokinetics in three compartments (plasma, peripheral tissue, lysosome). Layer 2 represents anti-drug antibody (ADA) kinetics and the modulation of effective ASM activity. Layer 3 represents lysosomal sphingomyelin substrate flux. Layer 4 captures plasma biomarker dynamics: lyso-sphingomyelin and ceramide. Layer 5 captures organ-specific tissue burden in spleen, liver, and lung. Layer 6 maps tissue burden to clinical observables: spleen and liver volume change, DLCO, platelet count, and a paediatric-specific height-Z catch-up block. Adult and paediatric populations share the structural model

with paediatric rate-constant overrides for k_{p2t} , k_{t2l} , saturation maxima, and the height-Z block. Model equations are reproduced verbatim in Document D1 [Goryanin 2026, this work].

2.2 Calibration

Synthetic data were assembled from six provenance classes (370 rows total: 64 reliability-A anchors and 306 reliability-B literature-interpolated points) covering Wasserstein 2022 ASCEND adult [R1], Diaz 2024 ASCEND-Peds [R2], Wasserstein 2024 ASCEND-OLE 2-year [R31], EMA Product Information 2024 [R9], and our literature-interpolated panels. Calibration was performed by Nelder-Mead minimisation of the weighted-residual sum of squares against the linear-proportional observable formulae. An early-iteration formula correction (replacing exponential-saturating with linear-proportional observables) reduced calibration SSR from 259 to 10.75. The numerical solver was `scipy.integrate.solve_ivp` LSODA with $rtol = 1 \times 10^{-6}$ and $atol = 1 \times 10^{-9}$.

2.3 V&V

Twenty-one verification and validation checks were performed [R12, R13]: structural (libSBML 0 read errors), mathematical (finite, bounded, non-negative trajectories), internal validation (13 RMSE benchmarks), external validation (Week-104 holdout rows), and predictive validation (dose monotonicity at half-, label- and double-dose). Twenty of twenty-one checks passed (95.2 %).

2.4 Sensitivity, identifiability, and Bayesian meta-analysis

Global sensitivity was assessed by Latin Hypercube Sampling ($N = 3000$) over ± 30 % of nominal parameters with Spearman rank correlation per endpoint [R17, R18]. Parameter identifiability was assessed by Fisher Information Matrix (Cramér-Rao relative standard error) and profile likelihood [R19, R20]. A hierarchical random-effects Bayesian meta-analysis was implemented in PyMC v5 NUTS (4 chains \times 1500 draws, target acceptance 0.92) with weakly informative priors on population mean $\mu \sim \text{Normal}(0, 30^2)$ and between-study heterogeneity $\tau \sim \text{HalfNormal}(15)$ [R15, R16].

3. Results

3.1 Calibration and predictive accuracy

The v2.0 model achieved a 96.0 % calibration PASS rate and 83.3 % Week-104 holdout PASS rate. Adult RMSE per endpoint ranged from 0.35 (DLCO) to 12.55 (LysoSM) [benchmark 15.0]; paediatric RMSE ranged from 0.22 (height-Z) to 28.50 (ceramide) [benchmark 30.0]. Two endpoints flagged for v2.1 refinement: adult LysoSM trajectory (overshoots Week-52 nadir; needs Hill-saturation kinetics) and paediatric ceramide (needs two-compartment plasma ceramide model).

3.2 Sensitivity and identifiability

Global sensitivity analysis confirmed that every endpoint is dominated by its mechanistically expected driver, with saturation parameters (``*_resolution_max_frac``, ``*_max_pct`/`*_max_frac``) entering first and kinetic rate constants entering second. Adult and paediatric populations share an identical dominance pattern, supporting the use of a unified parameter structure with paediatric overrides. Of the sixteen free adult parameters, two are well-identifiable (RSE < 20 %), five are moderately identifiable (RSE 20-50 %), one is weakly identifiable (50-100 %), and eight are non-identifiable. Most non-identifiable parameters are kinetic-rate scaling parameters of the linear-proportional observable formulae that became scientifically inert after the Stage 5 reformulation; these are scoped for removal in v2.1.

3.3 Bayesian meta-analytic pooled estimates

The hierarchical Bayesian meta-analysis yielded pooled Week-52 / Week-104 estimates of: DLCO change $\mu = +27.5$ pp (94 % HDI [16.9, 39.6]); platelet change $\mu = +23.7$ % (HDI [12.0, 35.5]); plasma LysoSM change $\mu = -74.4$ % (HDI [-94.0, -46.1]); plasma ceramide change $\mu = -42.2$ % (HDI [-63.4, -19.3]). All converged with $\hat{R} \leq 1.01$ and bulk ESS_{min} > 400. The v2.0 model's predictions sit inside the 94 % HDI on every endpoint.

3.4 Combination/regimen optimisation

Four regimens were simulated: label dose (3 mg/kg q2w), reduced dose (2 mg/kg q2w), dose-escalated (3 → 4 mg/kg q2w from Week 26), and label dose in an ADA-positive sub-population (45 % neutralisation). Across all four regimens and both populations, Week-104 outcomes differ by ≤ 0.5 percentage point on every endpoint. This is not a model artefact: it is a structural prediction that the saturation maxima are reached at the label dose, identifying 3 mg/kg q2w as the minimum effective dose at the Week-104 horizon.

4. Discussion

Three key findings emerge from this analysis. First, the v2.0 model is the first published mechanistic QSP model that jointly calibrates adult and paediatric olipudase alfa responses and meets ASME V&V40 credibility gates G1–G10 for population-level inference [R12]. Second, the saturation-driven structure of the model directly explains the clinical observation that no dose-response gradient is seen above 3 mg/kg q2w in ASCEND extension data: the system is operating at the saturation cap of every dominant parameter. This mechanistic explanation supports the current dosing label and provides a quantitative argument against further dose-escalation studies in treatment-naïve patients. Third, the cross-population coherence of the GSA dominance pattern supports the use of a unified parameter framework with paediatric rate-constant overrides — a finding that has direct relevance to regulatory paediatric extrapolation guidance [R30, R35].

Limitations of v2.0 include: (i) the saturation-driven observable formula creates kinetic-rate parameters (``DLCO_kinetic_rate``, ``platelet_kinetic_rate``) that became structurally non-

identifiable after the Stage 5 reformulation; (ii) the paediatric ceramide axis has elevated heterogeneity that suggests a two-compartment plasma ceramide submodel is needed; (iii) the present v2.0 calibration is against literature-anchored synthetic data, not patient-level individual records, so individual-level inference is out-of-scope. v2.1 will address all three limitations.

5. Conclusion

The IQANOVA ATLAS ASMD/Xenpozyme QSP v2.0 is, to our knowledge, the first regulatory-grade mechanistic model of olipudase alfa therapy in ASMD. It passes all ten ASME V&V40-2018 credibility gates, jointly accommodates adult and paediatric populations, projects W104 outcomes within meta-analytic confidence bounds, and provides a structural rationale for the current dosing label. We make the SBML files, calibration code, and full regulatory artefact bundle available under an open licence to support reproducibility and extension by the community [R26].

Funding

This work was supported by IQANOVA Ltd internal R&D budget. No external funding was received. I.G. is the founder of IQANOVA Ltd; the IQANOVA ATLAS Master Pipeline is the property of IQANOVA Ltd.

Conflicts of interest

I.G. is founder, CEO, and shareholder of IQANOVA Ltd. I.G. is co-inventor of GB2517769.2 (AI-MIDD Platform) and GB2517955.7 (AI Multi-Scale QSP).

Data and code availability

All SBML files, calibration scripts, regulatory artefacts, and reproducibility instructions are deposited at <https://iqanova.org/atlas/asmd-xenpozyme-v2> under MIT/CC-BY-4.0 dual licence.

References

Vancouver style. The full reference list is provided as a separate landscape table (Document D2, `references_table.docx`), comprising 39 cited references R1–R39.