

REMED I: REinforcement learning-driven adaptive MEtabolism modeling of primary sclerosing cholangitis DIsease progression

Chang Hu

University of Illinois at Urbana-Champaign

CHANGHU@ILLINOIS.EDU

Krishnakant V. Saboo

University of Illinois at Urbana-Champaign

KSABOO2@ILLINOIS.EDU

Ahmad H. Ali

University of Missouri School of Medicine

ALIAH@HEALTH.MISSOURI.EDU

Brian D. Juran

Mayo Clinic

JURAN.BRIAN@MAYO.EDU

Konstantinos N. Lazaridis

Mayo Clinic

LAZARIDIS.KONSTANTINOS@MAYO.EDU

Ravishankar K. Iyer

University of Illinois at Urbana-Champaign

RKIYER@ILLINOIS.EDU

Abstract

Primary sclerosing cholangitis (PSC) is a rare disease wherein altered bile acid metabolism contributes to sustained liver injury. This paper introduces REMEDI, a framework that captures bile acid dynamics and the body’s adaptive response during PSC progression that can assist in exploring treatments. REMEDI merges a differential equation (DE)-based mechanistic model that describes bile acid metabolism with reinforcement learning (RL) to emulate the body’s adaptations to PSC continuously. An objective of adaptation is to maintain homeostasis by regulating enzymes involved in bile acid metabolism. These enzymes correspond to the parameters of the DEs. REMEDI leverages RL to approximate adaptations in PSC, treating homeostasis as a reward signal and the adjustment of the DE parameters as the corresponding actions. On real-world data, REMEDI generated bile acid dynamics and parameter adjustments consistent with published findings. Also, our results support discussions in the literature that early administration of drugs that suppress bile acid synthesis may be effective in PSC treatment.

Keywords: Reinforcement learning, Disease progression, Differential equation, Adaptation

1. Introduction

Primary sclerosing cholangitis (PSC) is a rare, complex liver disease in which altered bile acid metabolism contributes to liver injury (Bertolini et al., 2022). There are no effective medications, and liver transplantation is often necessary (Vesterhus and Karlsen, 2020). A critical hurdle in exploring therapeutics is the lack of a model capturing the relevant disease dynamics, the body’s response to the disease, and the effects of treatments. We aim to develop a machine learning (ML) based PSC progression model with a focus on bile acid metabolism dynamics and its bidirectional interactions with the body over time. Such a model could facilitate treatment evaluations and accelerate drug discovery or repurposing. Examples of computational models guiding interventions already exist for prostate cancer (Zhang et al., 2017) and HIV (Xiao et al., 2013).

There were three main challenges in developing the proposed progression model: (1) the absence of a bile acid metabolism model during PSC; (2) limited insight into the body’s adaptive response to the disease; and (3) the lack of data from affected organs and a dearth of longitudinal data. (1) While prior studies have proposed differential equation (DE)-based bile acid metabolism models for healthy individuals (Sips et al., 2018), they do not capture bile duct obstruction, the pathophysiological hallmark of PSC

(Chapman et al., 2010), and its impact on bile acid metabolism. (2) Over the course of the disease, the body responds to changing bile acid levels by continually adapting and altering bile acid metabolism, which plays a central role in keeping PSC patients asymptomatic for many years (Jansen et al., 2017). However, the specific adaptations during PSC progression are not well understood (Milkiewicz et al., 2016), making them difficult to model. (3) Despite the liver and the bile ducts being central to PSC, direct bile acid measurements in these organs are infeasible. We are limited to bile acid data in the blood. Moreover, these data are cross-sectional, i.e., taken only at a single time point, further complicating the modeling of longitudinal disease progression.

We introduce REMEDI to model PSC progression by extending existing bile acid metabolism DEs with PSC pathophysiology and incorporating a reinforcement learning (RL) agent to approximate the body’s adaptations. REMEDI addresses the above challenges with the following key innovations:

(1) We developed a reduced-order bile acid metabolism model to capture dynamics pertinent to PSC, based on an existing DE model for healthy individuals. We extended the reduced model with clinical domain knowledge to capture bile duct obstruction.

(2) We used RL to emulate the body’s adaptive response to the disease. We assume the body is a smart agent that, through evolution, has learned to adapt itself to maintain homeostasis of critical metabolic events (Giordano, 2013). In PSC, the body regulates bile acid metabolism enzymes to maintain homeostasis. These enzymes are naturally represented as parameters in the reduced-order bile acid DE model. During disease progression, the RL agent constantly updates these parameters to maximize a reward function that promotes homeostasis and the generation of close-to-reality bile acid profiles.

(3) Because PSC patients can have a prolonged, partially successful adaptation period (Jansen et al., 2017), we assume the real-world cross-sectional data were collected during this “stable” period, and we encourage the RL agent to generate stabilized trajectories that are close to the data.

The main assumption of REMEDI is that the goal of homeostasis drives biological actions, achieved by smart regulation of body enzymes (Savir et al., 2017; Billman, 2020). Thus, we treat adaptation as a sequential optimization problem with “homeostasis” as the objective function and the sequential regulation of enzymes as the optimization arguments. These en-

zymes are represented as parameters in the DEs. RL offers a framework to solve this sequential optimization problem. In PSC, the bile acid DEs constitute the environment, “homeostasis” is the reward function, and the modulation of DE parameters is the actions. Therefore, RL, in combination with the DEs, approximates the body’s adaptation to disease and enables dynamic modeling of PSC progression.

We validated REMEDI against findings from the literature and real-world clinical data. REMEDI produced biologically realistic results. (1) The reduced-order bile acid model captured the relevant dynamics of a more detailed model from the literature with drastically reduced computational cost. (2) Incorporation of bile duct obstruction mimicked PSC pathophysiology observed in clinical and animal studies (Jansen et al., 2017; Boyd et al., 1966; Popper and Schaffner, 1970). (3) On real-world PSC data, REMEDI generated bile acid dynamics and parameter adaptations consistent with the literature (Jansen et al., 2017). (4) We evaluated *in silico* the effects of two PSC drugs in clinical trials (Vesterhus and Karlsen, 2020) and found REMEDI has the potential to explain the drugs’ biologically observed behaviors (Boyer, 2007; Caballero-Camino et al., 2023).

Our contributions include: (1) developing the first mathematical model of bile acid metabolism in PSC, based on clinical domain knowledge, and (2) providing an *in silico* testbed to evaluate the effects of bile acid modulating therapies. Our approach can be leveraged to determine optimal interventions for PSC in combination with comprehensive clinical data.

In principle, REMEDI’s approach of using RL to estimate time-varying DE parameters can be extended to other diseases where DE-based models with time-varying parameters have been proposed, such as HIV (Liang et al., 2010). Moreover, the innovative strategy of REMEDI that leverages RL to emulate adaptive behaviors holds promise for modeling a variety of homeostatic biological systems.

2. REMEDI Framework

Our approach has three parts (Figure 1): (1) a reduced-order bile acid metabolism model for healthy individuals (Section 2.1), (2) a domain knowledge-based extension to depict the pathophysiology of PSC (Section 2.2), and (3) RL that captures the body’s adaptation to the pathophysiology (Section 2.3).

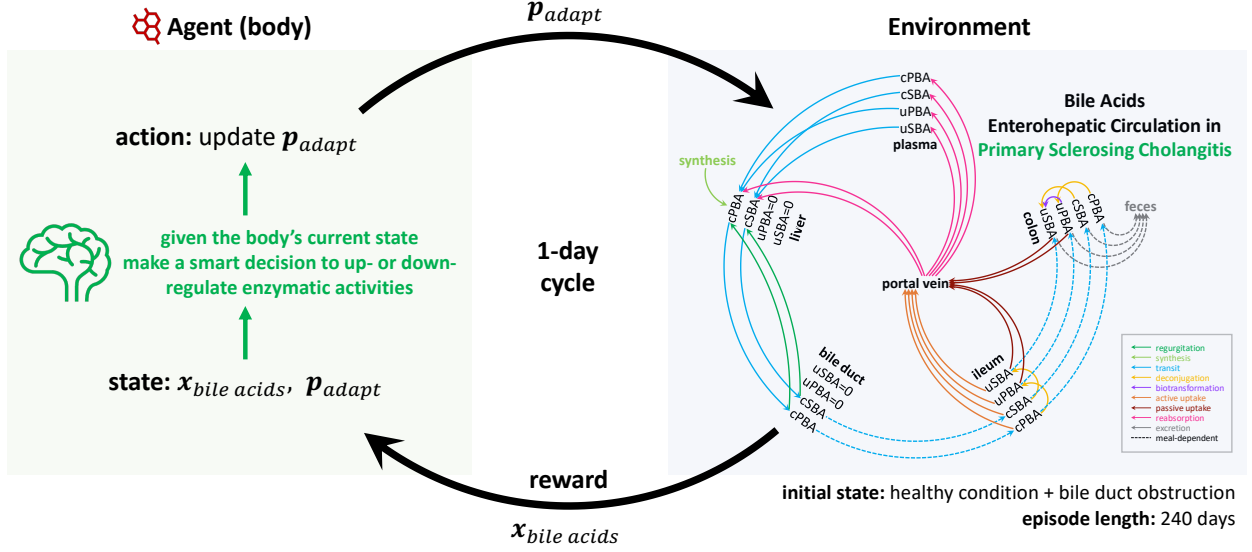


Figure 1: Overview of the reinforcement learning formulation of bile acid metabolism adaptation in PSC.

2.1. Model of healthy bile acid metabolism

Several species of bile acids circulate in multiple organs of the human body in a process called *enterohepatic circulation* (Hofmann, 2009). In the liver (li), cholesterol is metabolized into the unconjugated primary bile acids cholic acid and chenodeoxycholic acid (uCA and uCDCA), which are transformed into conjugated forms (cCA and cCDCA) and secreted into the bile ducts and gallbladder (bd). These bile acids are then released into the intestines, where bacteria in the ileum (il) and colon (co) convert them into secondary bile acids (SBA). Active and passive uptake reabsorb intestinal bile acids back to the liver, with a small portion escaping into plasma (pl). Unreabsorbed bile acids are excreted with feces (fe).

We adopt Sips et al. (2018)’s approach and model bile acid metabolism with a series of DEs. Based on their relevance to PSC, we merge several bile acid species and do not distinguish among certain organ segments (see Appendix A), resulting in a reduced-order model of bile acid $BA \in \{cCA, cCDCA, cSBA, uCA, uCDCA, uSBA\}$ in organ $OG \in \{li, bd, il, co, pl, fe\}$. For each BA in each OG, we (1) model its influxes/outfluxes from relevant biochemical and physical processes and (2) combine the fluxes into one DE to describe how the BA level in OG varies with time. See Appendix A for all processes being modeled and Appendix B for all corresponding DEs.




Figure 2: Fluxes affecting liver cCA level.

As an example, Figure 2 shows the fluxes that affect the liver cCA level: (1) r_{syn}^{cCA} is de novo synthesis in liver cells; (2) r_{li}^{cCA} from gut is active and passive uptake from the gut; (3) r_{li}^{cCA} from pl is influx from systemic blood circulation; and (4) r_{li}^{cCA} to bd is outflux to the gallbladder and bile ducts. Hence, liver cCA level x_{li}^{cCA} varies with time ($\mu\text{mol}/\text{min}$) according to:

$$\frac{dx_{li}^{cCA}}{dt} = r_{syn}^{cCA} + r_{li}^{cCA} \text{ from gut} + r_{li}^{cCA} \text{ from pl} - r_{li}^{cCA} \text{ to bd} \quad (1)$$

$$r_{syn}^{cCA} = p[\text{synthesis}] \cdot p[\text{syn_frac_CA}] \quad (2)$$

$$r_{li}^{cCA} \text{ to bd} = p[\text{li_to_bd_freq}] \cdot x_{li}^{cCA} \quad (3)$$

We model all fluxes with zero- or first-order dynamics. For example, we model the cCA synthesis flux r_{syn}^{cCA} as the product of two constant parameters, making r_{syn}^{cCA} also a constant (zero-order). Here,

$p[\text{synthesis}]$ denotes the total bile acid synthesis rate, and $p[\text{syn_frac_CA}]$ describes the fraction of cholic acid among the newly synthesized bile acids. The liver-to-bile duct transit rate of cCA, $r_{\text{li to bd}}^{\text{cCA}}$, exemplifies first-order dynamics. $x_{\text{li}}^{\text{cCA}}$ represents the current liver cCA level, and the constant parameter $p[\text{li_to_bd_freq}]$ characterizes the first-order dynamics.

2.2. Introducing PSC pathophysiology

Drawing on clinical domain knowledge, we extend the reduced-order bile acid metabolism model with PSC pathophysiology (Figure 3) by (1) implementing an obstruction of bile flow in the bile ducts and (2) introducing bile acid backflow to the liver following excessive bile acid buildup in the bile ducts.

Obstructed bile flow in the bile ducts: In PSC, chronic inflammation causes scarring and narrowing of the bile ducts (Karlsen et al., 2017), impeding the normal bile flow into the small intestine (Figure 3(b)). The extent of the obstruction determines the reduction of bile flow. We introduce a parameter $p[\text{bd_max_flow}]$ to denote the maximum amount of bile acids allowed to flow through, in proportion to the degree of obstruction. Consequently, if $r_{\text{bd to il}}$ calculated from its first-order dynamics exceeds $p[\text{bd_max_flow}]$, we cap it at $p[\text{bd_max_flow}]$:

$$r_{\text{bd to il}} = \min(r_{\text{bd to il}}, p[\text{bd_max_flow}]) \quad (4)$$

Bile acid backflow to the liver: The bile ducts (and gallbladder) have a limited storage capacity for bile acids. In PSC, bile duct obstruction can result in bile acid buildup exceeding the duct’s capacity, leading to regurgitation and backflow of excessive bile acids to the liver (Popper and Schaffner, 1970), as depicted in Figure 3(c). To represent this, we introduce a parameter $p[\text{bd_max_ba}]$ denoting the bile acid holding capacity of the bile duct, and define $r_{\text{bd to li}}$ as the excessive bile acids backflowing from the bile duct to the liver:

$$r_{\text{bd to li}} = \max(x_{\text{bd}} + r_{\text{bd from li}} - r_{\text{bd to il}} - p[\text{bd_max_ba}], 0) \quad (5)$$

x_{bd} is the current bile acid level in the bile duct, $r_{\text{bd from li}}$ is the influx from the liver, and $r_{\text{bd to il}}$ is the outflux to the ileum. See Appendix B for details.

2.3. Model adaptation in PSC

We model adaptation as a series of enzymatic regulations to optimally maintain homeostasis, i.e., preserve

physiological functions without deviating too much from the enzyme levels under healthy conditions. We model this adaptation process using RL (Figure 1).

State: We assume the body self-regulates based on its current status, captured by the state vector $= \{\mathbf{x}_{\text{bile acids}}, \mathbf{p}_{\text{adapt}}\}$ in our context of bile acid metabolism. $\mathbf{x}_{\text{bile acids}}$ denotes bile acid levels across species and organs (30 variables, see Appendix C). $\mathbf{p}_{\text{adapt}}$ denotes parameter values corresponding to regulatable enzymes (five variables, see **Action**).

Action: We assume the body adapts through continual enzyme regulations, which translates to the modulation of the subset of DE parameters representing enzyme levels regulated by the body. Our DEs contain five such regulatable parameters, i.e., $\mathbf{p}_{\text{adapt}} = \{p[\text{synthesis}], p[\text{syn_frac_CA}], p[\text{hep_ratio_conj_tri}], p[\text{hep_ratio_conj_di}], p[\text{max_asbt_rate}]\}$. At every RL step, for each parameter in $\mathbf{p}_{\text{adapt}}$, any of three actions can be taken: up-regulation or down-regulation (with a prespecified fold change or absolute difference), or remaining unchanged. $p[\text{synthesis}]$ and $p[\text{syn_frac_CA}]$ were introduced in Section 2.1. $p[\text{hep_ratio_conj_tri}]$ denotes the fraction of reabsorbed bile acid extracted by the liver (in contrast to going into systemic blood) for cCA, and $p[\text{hep_ratio_conj_di}]$ for cCDCA and cSBA. $p[\text{max_asbt_rate}]$ is the rate parameter in the first-order dynamics of active uptake.

Environment: The RL environment simulates bile acid dynamics following the introduction of PSC pathophysiology to healthy conditions. We use the reduced-order DEs extended with bile duct obstruction as the simulator. At every RL step, the RL agent modifies $\mathbf{p}_{\text{adapt}}$ in the DEs and updates the state vector. The environment takes a step forward via numerical integration of the DEs for a fixed duration. The resultant bile acid levels update $\mathbf{x}_{\text{bile acids}}$ in the state vector. A reward is computed in the environment step and sent to the RL agent to determine the next action. The simulation terminates when a pre-specified time period is reached or when any state variable exceeds physiological ranges.

Reward: We formulate our reward function in consultation with domain experts to guide the RL agent towards meaningful adaptations. The reward comprises several terms that (1) address the dual roles of bile acids, namely their contributions to liver toxicity (*minimizing toxicity*) and their physiological functions (*facilitating digestion* and *main-*




Figure 3: Idealized modeling of PSC pathophysiology.

taining cholesterol elimination), and (2) ensure the generated trajectories are empirically reasonable for PSC (*resembling real-world patient data* and *conforming with values from the literature*).

Minimizing toxicity. One of the main goals of adaptation is to limit liver exposure (LE) to toxic bile acids (Boyer, 2007). We set a negative reward for excessive bile acid exposure in the liver to minimize liver toxicity. We calculate LE as the cumulative liver bile acid level over one day. If current LE exceeds LE under healthy conditions, we set the negative reward to be the normalized excessive exposure.

$$-\max\left(\frac{\text{current } LE - \text{healthy } LE}{\text{maximum possible } LE}, 0\right) \quad (6)$$

Facilitating digestion. Adaptation requires preserving digestive functions under disease conditions (Tappenden, 2014). Because bile acids in the ileum are necessary for fat digestion, we promote ileum access (IA) to bile acids with a reward term defined as the ratio of current IA to IA under healthy conditions. IA is calculated as the cumulative ileum bile acid level over one day. The reward is capped at 1, offering no additional benefit beyond healthy levels.

$$\min\left(\frac{\text{current } IA}{\text{healthy } IA}, 1\right) \quad (7)$$

Maintaining cholesterol elimination. Synthesizing bile acids from cholesterol is one of the main pathways for eliminating cholesterol from the body (Wang et al., 2018). Insufficient elimination of cholesterol increases the risks of multiple diseases. We reward sufficient cholesterol elimination, represented by the ratio of the current bile acid synthesis (BAS) rate to the BAS rate under healthy conditions. The reward is capped at 1, offering no additional benefit from excessive cholesterol elimination.

$$\min\left(\frac{\text{current } BAS \text{ rate}}{\text{healthy } BAS \text{ rate}}, 1\right) \quad (8)$$

Resembling real-world patient data. The way to adapt and sustain physiological functions might not be unique. To obtain an RL agent that mirrors adaptation in humans, we set a reward term to promote RL solutions that resemble real-world plasma data from PSC patients. We select a representative patient from our cohort (see Section 3) and penalize the difference between the patient’s data BA_{data} and the respective RL states BA_{RL} . This difference is divided by the corresponding bile acid’s standard deviation. The negative sum of the squared weighted difference is multiplied by a coefficient $\lambda \in [0, 1]$ to match other reward terms’ range.

$$-\lambda \min\left(\Sigma_{BA} \left(\frac{BA_{RL} - BA_{data}}{\text{std } BA_{data}}\right)^2, \text{CAP}\right) \quad (9)$$

Conforming with values from the literature: We design additional reward terms to ensure the RL agent generates physiologically plausible bile acid levels while keeping the regulatable parameters close to their values under healthy conditions. See Appendix C for further details.

2.4. Implementation of REMEDI

Degree of pathophysiology: To determine the extent of bile duct obstruction, we ran a grid search for the parameter $p[\text{bd_max_flow}]$ across values of 1, 2, 3, 5, 10, 20, and 40 $\mu\text{mol}/\text{min}$, simulating scenarios from near-complete to partial obstruction. Corresponding RL agents were trained independently.

RL timeframes: At each step, the RL agent modified p_{adapt} to simulate the next 24 hours. Considering adaptation unfolds over days to weeks, this daily cycle offered sufficient opportunities for meaningful modulations. We restricted the simulation to a maximum of 240 days to adequately encapsulate the initial adaptation phase (Georgiev et al., 2008).

Adaptation amplitudes: We chose relatively large parameter adaptation amplitudes to match the day-long RL steps. For parameters describing rates, i.e., $p[\text{synthesis}]$ and $p[\text{max_asbt_rate}]$, we simulated up- or down-regulation with a 25% higher or lower fold change; for parameters describing fractions, i.e., $p[\text{syn_frac_CA}]$, $p[\text{hep_extract_ratio_conj_tri}]$, and $p[\text{hep_extract_ratio_conj_di}]$, we applied a 10% addition or subtraction. We also prespecified physiologically plausible ranges for each parameter, and up- or down-regulated values exceeding the ranges were clipped. The RL agent generated stochastic actions during training and evaluation.

Real-world patient data: Our dataset includes 222 PSC patients selected from the PSC Scientific Community Resource (Ali et al., 2021). Plasma bile acid profiling was performed at the Mayo Clinic (Mousa et al., 2021). We selected five representative patients whose measurements minimized the sum of the distances to all patients. We referred to the patient with the smallest sum of the distances to all patients as Patient 1. Patient-specific models were trained using their respective data, resulting in five models for the five patients. Since PSC is chronic, we assume patient data were collected during the stabilized adaptation phase. Hence, the sum of squared errors to encourage resemblance to patient data was only introduced after the initial four weeks.

RL training: We trained the RL agent for 4,000,000 environment steps using the model-free Proximal Policy Optimization (PPO) algorithm implemented in the Python package Stable Baselines3.

See Appendix C for more details on our RL setup.

2.5. Code Availability

Our code, example data, and a comprehensive README file are publicly available at <https://github.com/chang-hu/REMEMEDI>.

3. Results

3.1. Healthy bile acid dynamics

We derived the DE parameters and initial bile acid values from the well-calibrated model by Sips et al. (2018). We simulated our DEs over 60 days, reaching a steady state at which bile acid dynamics repeated every 24 hours. The Runge-Kutta 45 numerical solver was used to integrate the DEs in our experiments.

We validated our reduced-order DEs against the original DEs proposed by (Sips et al., 2018) (see Appendix A for details). The reduced-order DEs completed a 60-day simulation in 2.99s, over ten times faster than the 32.5s of the original DEs, while retaining similar bile acid trends (Appendix A Figure 7). The steady-state (day 60) bile acid levels of the reduced-order DEs were also consistent with real-world plasma data of 302 healthy individuals from Mayo Clinic (Appendix A Figure 8 and Figure 9).

3.2. PSC bile acid dynamics without RL

For this analysis, we employed the reduced-order DEs extended with PSC pathophysiology, excluding RL adaptation. The parameter representing the degree of bile duct obstruction, $p[\text{bd_max_flow}]$, was set to $3 \mu\text{mol}/\text{min}$ (in contrast to $75 \mu\text{mol}/\text{min}$ when unobstructed). Other DE parameters and bile acid levels were initialized with their steady-state values under healthy conditions. We ran the simulation for 60 days, which, as shown by animal models, was enough time to establish adaptation (Georgiev et al., 2008).

Figure 4(a) shows the 60-day cCDCA bile acid dynamics following the introduction of PSC pathophysiology (see Appendix D for dynamics of other bile acids). We observed a 230% surge in bile duct bile acid levels, a direct result of impaired bile flow to the ileum and the subsequent accumulation in the bile ducts. A corresponding decrease in ileum bile acids was also observed. Around day five, bile duct bile acids reached a saturation point, causing excess bile acids to flow back into the liver. These observations align with the biologically expected changes (Boyd et al., 1966; Popper and Schaffner, 1970).

However, once bile backflow starts, it continues at a near-constant rate throughout the remaining simulation, generating unrealistically high liver bile acid levels. The cCDCA level rose to $27,148 \mu\text{mol}$ by day 60, in stark contrast to the $45 \mu\text{mol}$ under healthy conditions. The simulation also indicated decreased conjugated bile acid levels in the plasma, conflicting with data showing elevated levels in PSC patients. These discrepancies arise from the flawed assumption that bile acid metabolism parameters remain unchanged, neglecting the dynamic adaptation occurring in PSC.

3.3. PSC bile acid dynamics with REMEDI

We tested REMEDI upon introducing PSC pathophysiology to healthy conditions. We evaluated a range of $p[\text{bd_max_flow}]$ values and chose the case

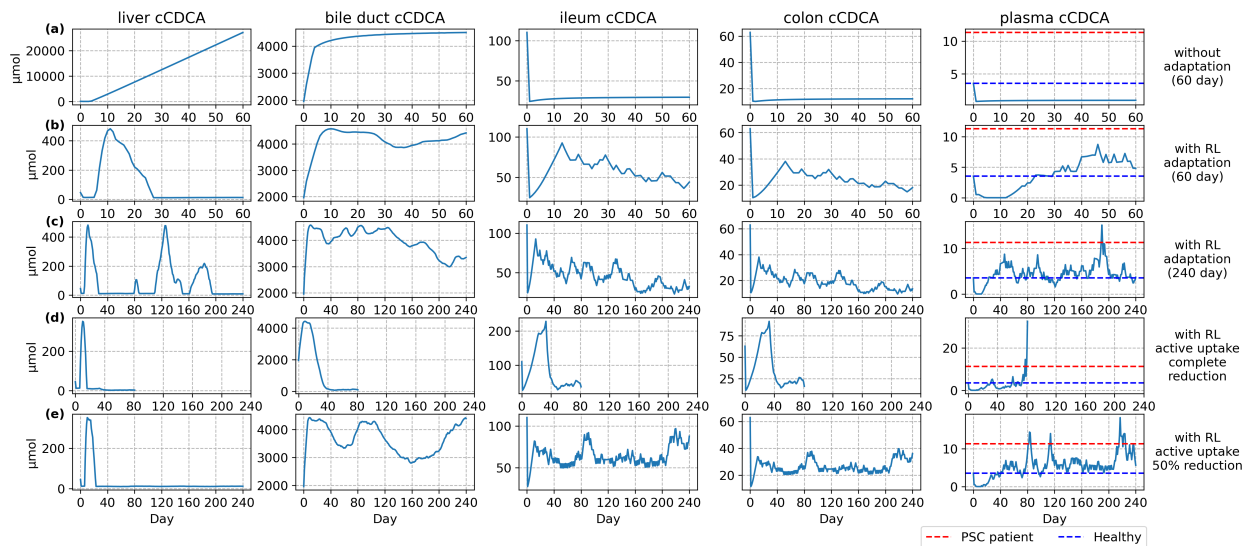


Figure 4: Simulated 60-day (panels (a)–(b)) and 240-day ((c)–(e)) bile acid dynamics after introduction of PSC pathophysiology ($p[\text{bd_max_flow}]=3 \mu\text{mol}/\text{min}$) to healthy conditions at day 0. (a) Without adaptation; (b)–(c) with adaptive DE parameters obtained from a trained RL agent; (d) with complete reduction of active uptake; (e) with 50% reduction of active uptake. Results shown here were derived from the model trained with Patient 1’s data. Only cCDCA values at 8 AM of each day are plotted. The blue dotted line represents the bile acid levels for healthy individuals, specifically, the steady-state (day 60) bile acid levels of the reduced-order DEs (see Section 3.1 and Appendix A). The red dotted line represents the plasma bile acid measurement of Patient 1. See Appendix D and Appendix E for the complete dynamics from all patients.

with $p[\text{bd_max_flow}]=3 \mu\text{mol}/\text{min}$ for further analysis, as it yielded plasma bile acid levels closest to real data and therefore was more likely to reflect the real-world disease conditions (see Appendix C).

Upon introducing PSC pathophysiology in Patient 1’s model, we observed an initial surge of bile acid levels in the bile ducts and a decrease in the downstream intestines (Figure 4(b)), as observed in animal studies of PSC (Boyd et al., 1966). Around day five, liver bile acids started to accumulate following saturation in the bile ducts. Importantly, REMEDI was able to adjust and stabilize the liver bile acid levels by week four, avoiding the unrealistic continuous rise seen without RL adaptation (Figure 4(a)). Furthermore, REMEDI showed an increase in conjugated bile acids in the plasma and a decrease in unconjugated forms (Appendix D Figure 13), better aligning with the real-world PSC patient data than without RL adaptation (Figure 4(c)). Similar results were seen for models trained on Patient 2–5 (Appendix E).

To quantitatively assess the improvement from RL, we compared Patient 1’s measurements with bile acid predictions from models trained on data from four other patients. The average error with RL was smaller than the error without RL (Table 1). We adopted this approach for validation as we only had one cross-sectional measurement for each patient.

Overall, REMEDI with RL-based adaptation captured key disease dynamics and generated a more faithful representation of real-world data.

3.4. Trajectories of adaptive parameters

Analyzing parameter adjustments by REMEDI and linking them to the underlying enzymes can shed light on possible adaptive mechanisms driving bile acid metabolism in PSC. Notably, bile acid synthesis went through a sharp decline following the bile duct obstruction and remained low throughout the simula-

Table 1: Error between the model and PSC Patient 1’s data (μmol , day 50 – 60 averaged).

Model/Plasma	cCA	cCDCA	cSBA	uCA	uCDCA	uSBA
REMEDI without RL	-12.01 (-97.74%)	-10.37 (-91.66%)	-2.21 (-95.91%)	-0.08 (-48.18%)	-0.18 (-65.19%)	-0.35 (-79.23%)
REMEDI	-9.91 (-80.63%)	-6.68 (-59.07%)	0.25 (11.09%)	0.08 (51.69%)	0.10 (34.88%)	0.14 (31.45%)

To make comparison across bile acid species more meaningful, we also calculate the errors as proportions of the observed values (in the parenthesis): $\frac{\text{Model Prediction} - \text{Patient 1's Measurement}}{\text{Patient 1's Measurement}} \times 100\%$.




Figure 5: DE parameter adaptations by REMEDI following the introduction of PSC pathophysiology. *synthesis* (blue): bile acid synthesis rate; *syn_frac_CA* (orange): fraction of CA (vs. CDCA) in newly synthesized bile acid; *hep_ratio_conj_tri* (green): fraction of bile acid uptake entering the liver (vs. plasma) for cCA; *hep_ratio_conj_di* (red): fraction of bile acid uptake entering the liver for cCDCA and cSBA; *max_asbt_rate* (purple): efficiency of bile acid active uptake.

tion period (Figure 5), in line with the known down-regulation of the bile acid synthesis enzyme CYP7A1 in cholestasis (a condition in which the bile flow from the liver stops or slows) (Jansen et al., 2017).

The model also predicted fluctuations in the CA:CDCA ratio among newly synthesized bile acids ($p[\text{syn_frac_CA}]$). CA:CDCA ratio depends on the activities of the classical and alternative bile acid synthesis pathways, which are regulated by enzymes such as CYP8B1 (Li and Chiang, 2012). How the activities of these enzymes change in PSC is unclear, but our real-world PSC patient data also showed an altered

plasma CA:CDCA ratio, warranting further study of potential enzymatic shifts.

We observed reduced liver extraction of reabsorbed intestinal bile acids, with more bile acids entering systemic blood ($p[\text{hep_extract_ratio_conj_tri}]$ and $p[\text{hep_extract_ratio_conj_di}]$), potentially explaining the elevated plasma bile acid levels commonly seen in PSC patients. This trend aligns with cholestasis-related down-regulation of NTCP, a key enzyme in liver bile acid extraction (Donner et al., 2007).

Finally, there was a temporary drop in the ileum bile acid active uptake efficiency in the initial adaptation phase ($p[\text{max_asbt_rate}]$), corroborating the down-regulation of the bile acid uptake mediating enzyme ASBT in animal studies (Hruz et al., 2006).

Trends in Patient 2–5 were similar (Appendix E).

3.5. *In silico* evaluation of bile acid therapies

Using the trained REMEDI model, we assessed two types of experimental therapies targeting bile acid metabolism *in silico* (Boyer, 2007): (1) suppression of bile acid synthesis, which we studied through analyzing bile acid synthesis rate ($p[\text{synthesis}]$) adjustments during adaptation, and (2) reduction of bile acid active uptake, which we simulated by decreasing efficiency of intestinal active uptake ($p[\text{max_asbt_rate}]$). Our analysis suggested partial reduction of bile acid active uptake might protect the liver, highlighting the value of REMEDI in evaluating therapies *in silico*.

Suppressing synthesis: Drugs suppressing bile acid synthesis enzyme CYP7A1 are in clinical trials for PSC and other cholestatic conditions (Chiang and Ferrell, 2020). Interestingly, REMEDI suggested the liver may naturally inhibit bile acid synthesis ($p[\text{synthesis}]$) as an adaptive response (Figure 4(c)), implying the therapeutic window for such drugs may

be limited to the early stages of the disease, before endogenous adaptations establish.

Reducing active uptake: ASBT, a central enzyme in ileum bile acid active uptake, has been the target of multiple drugs (Vesterhus and Karlsen, 2020). We simulated two strategies targeting this mechanism by adjusting $p[\text{max_asbt_rate}]$: (1) complete reduction, down-regulating $p[\text{max_asbt_rate}]$ until it reaches its physiological lower bound, and (2) 50% reduction, down-regulating $p[\text{max_asbt_rate}]$ until it reaches or falls below half of the lower bound.

Complete reduction: REMEDI predicted that full inhibition of active uptake would result in physiologically implausible spikes in plasma and intestine bile acids (Figure 4(d)), which led to premature termination of the simulation, implying complete reduction is likely an unrealistic strategy.

50% reduction: Partial reduction limited the occurrence of liver bile acid spikes (Figure 4(e)), potentially protecting the liver from excessive bile acid accumulation, in line with animal studies showing a liver-protection effect from an ASBT inhibitor (Caballero-Camino et al., 2023). In contrast to complete reduction, 50% reduction yielded plasma and intestine bile acid levels within realistic ranges, making partial reduction a more viable strategy.

4. Related Works

Several studies have proposed DE-based mechanistic models of bile acid metabolism (Hofmann et al., 1983; Molino et al., 1986; Sips et al., 2018; Baier et al., 2019; Voronova et al., 2020), albeit for healthy or non-PSC conditions. Moreover, the adaptive responses of the body in pathological conditions were usually not considered. A notable exception is Voronova et al. (2020), which explicitly modeled the FXR-FGF19 bile acid self-regulation pathway (Eloranta and Kullak-Ublick, 2008). Our approach is unique for using RL to simultaneously consider multiple regulation pathways without explicitly modeling their mechanisms. A previous study combined DEs with RL to model Alzheimer’s disease progression (Saboo et al., 2021). They used RL to estimate DE variable values that maximize cognition and minimize energetic cost, whereas we estimate DE parameters representing enzyme levels that promote homeostasis.

5. Limitations

First, we made several assumptions due to the current limited understanding of PSC, including focusing on bile acid metabolism, modeling bile duct obstruction as a sudden blockage, and in setting adaptation goals. Refining these assumptions as new insights of PSC emerge will lead to a more realistic model. One possible future extension is to include the bile acid-gut microbiota crosstalk (Wahlström et al., 2016). Second, our bile acid trajectory prediction was only compared to a single time point because we only had access to cross-sectional data. However, plasma bile acid data is being collected annually at the Mayo Clinic beginning in 2020. To further validate REMEDI, when the longitudinal data becomes available, we plan to use the first year’s data for training and the subsequent years for validation. This will be achieved by modifying the reward function term that encourages resemblance to real-world patient data.

6. Conclusion

We developed REMEDI, a novel model of PSC progression by combining bile acid metabolism DEs with an RL agent that captures the body’s adaptation. REMEDI captures key bile acid trends in disease progression consistent with the literature and predicted therapy responses *in silico*.

While this work focuses on PSC, our model can be readily adapted to other bile acid-related diseases, such as primary biliary cirrhosis and benign recurrent intrahepatic cholestasis. Additionally, REMEDI’s adaptive approach to disease progression modeling can be extended to a wide range of diseases, including type 2 diabetes (glucose metabolism) and hormone receptor positive cancers like breast, ovarian, or prostate cancers (sex hormone metabolism).

Acknowledgments

We thank August John, Yicheng Wang, Anirudh Choudhary, Yurui Cao, Mosbah Aouad, Jenny Applequist, and Yixin Chen for the insightful discussions. This work was supported by the Mayo Clinic and Illinois Alliance Fellowship for Technology-based Healthcare Research and in part by NSF grant CNS-1624790 and NIH grant RC2DK118619.

References

- Ahmad Hassan Ali, Brian D. Juran, Erik M. Schlicht, Jackie K. Bianchi, Bryan M. McCauley, Elizabeth J. Atkinson, and Konstantinos N. Lazaridis. The PSC scientific community resource: an asset for multi-omics interrogation of primary sclerosing cholangitis. *BMC Gastroenterology*, 21(1):353, 2021. doi: 10.1186/s12876-021-01930-2.
- Yazen Alnouti. Bile acid sulfation: A pathway of bile acid elimination and detoxification. *Toxicological Sciences*, 108(2):225–246, 2009. ISSN 1096-6080. doi: 10.1093/toxsci/kfn268.
- Vanessa Baier, Henrik Cordes, Christoph Thiel, José V. Castell, Ulf P. Neumann, Lars M. Blank, and Lars Kuepfer. A Physiology-Based Model of Human Bile Acid Metabolism for Predicting Bile Acid Tissue Levels After Drug Administration in Healthy Subjects and BRIC Type 2 Patients. *Frontiers in Physiology*, 10:1192, 2019. ISSN 1664-042X. doi: 10.3389/fphys.2019.01192.
- Anna Bertolini, Romina Fiorotto, and Mario Strazzabosco. Bile acids and their receptors: modulators and therapeutic targets in liver inflammation. *Seminars in Immunopathology*, 44(4):547–564, 2022. ISSN 1863-2297. doi: 10.1007/s00281-022-00935-7.
- George E. Billman. Homeostasis: The Underappreciated and Far Too Often Ignored Central Organizing Principle of Physiology. *Frontiers in Physiology*, 11:200, 2020. ISSN 1664-042X. doi: 10.3389/fphys.2020.00200.
- G. S. Boyd, M. A. Eastwood, and N. MacLean. Bile acids in the rat: studies in experimental occlusion of the bile duct. *Journal of Lipid Research*, 7(1):83–94, 1966. ISSN 0022-2275.
- James L. Boyer. New perspectives for the treatment of cholestasis: Lessons from basic science applied clinically. *Journal of Hepatology*, 46(3):365–371, 2007. ISSN 0168-8278. doi: 10.1016/j.jhep.2006.12.001.
- Francisco J. Caballero-Camino, Pedro M. Rodrigues, Fredrik Wängsell, Aloña Agirre-Lizaso, Paula Olaizola, Laura Izquierdo-Sanchez, Maria J. Peruggorria, Luis Bujanda, Bo Angelin, Sara Straniero, Anna Wallebäck, Ingemar Starke, Per-Göran Gillberg, Ellen Strängberg, Britta Bonn, Jan P. Mattsson, Martin R. Madsen, Henrik H. Hansen, Erik Lindström, Peter Åkerblad, and Jesus M. Banales. A3907, a systemic ASBT inhibitor, improves cholestasis in mice by multiorgan activity and shows translational relevance to humans. *Hepatology*, 78(3):709–726, 2023. ISSN 0270-9139. doi: 10.1097/hep.0000000000000376.
- Roger Chapman, Johan Fevery, Anthony Kalloo, David M. Nagorney, Kirsten Muri Boberg, Benjamin Shneider, Gregory J. Gores, and American Association for the Study of Liver Diseases. Diagnosis and management of primary sclerosing cholangitis. *Hepatology*, 51(2):660–678, 2010. ISSN 1527-3350. doi: 10.1002/hep.23294.
- John Y.L. Chiang and Jessica M. Ferrell. Up to date on cholesterol 7 alpha-hydroxylase (CYP7A1) in bile acid synthesis. *Liver Research*, 4(2):47–63, 2020. ISSN 2542-5684. doi: 10.1016/j.livres.2020.05.001.
- Markus G. Donner, Stephanie Schumacher, Ulrich Warskulat, Jane Heinemann, and Dieter Haüssinger. Obstructive cholestasis induces TNF- α - and IL-1 β -mediated periportal downregulation of Bsep and zonal regulation of Ntcp, Oatp1a4, and Oatp1b2. *American Journal of Physiology-Gastrointestinal and Liver Physiology*, 293(6):G1134–G1146, 2007. ISSN 0193-1857. doi: 10.1152/ajpgi.00079.2007.
- Jyrki J. Eloranta and Gerd A. Kullak-Ublick. The Role of FXR in Disorders of Bile Acid Homeostasis. *Physiology*, 23(5):286–295, 2008. ISSN 1548-9213. doi: 10.1152/physiol.00020.2008.
- P. Georgiev, W. Jochum, S. Heinrich, J. H. Jang, A. Nocito, F. Dahm, and P.-A. Clavien. Characterization of time-related changes after experimental bile duct ligation. *British Journal of Surgery*, 95(5):646–656, 2008. ISSN 0007-1323. doi: 10.1002/bjs.6050.
- Mario Giordano. Homeostasis: An underestimated focal point of ecology and evolution. *Plant Science*, 211:92–101, 2013. ISSN 0168-9452. doi: 10.1016/j.plantsci.2013.07.008.
- A F Hofmann, G Molino, M Milanese, and G Belforte. Description and simulation of a physiological pharmacokinetic model for the metabolism and enterohepatic circulation of bile acids in man. Cholic

- acid in healthy man. *Journal of Clinical Investigation*, 71(4):1003–1022, 1983. ISSN 0021-9738. doi: 10.1172/jci110828.
- F Hofmann, Alan. The enterohepatic circulation of bile acids in mammals: form and functions. *Frontiers in Bioscience*, Volume(14):2584, 2009. ISSN 1093-9946. doi: 10.2741/3399.
- P Hruz, C Zimmermann, H Gutmann, L Degen, U Beuers, L Terracciano, J Drewe, and C Beglinger. Adaptive regulation of the ileal apical sodium dependent bile acid transporter (ASBT) in patients with obstructive cholestasis. *Gut*, 55(3):395, 2006. ISSN 0017-5749. doi: 10.1136/gut.2005.067389.
- Peter L.M. Jansen, Ahmed Ghallab, Nachiket Vartak, Raymond Reif, Frank G. Schaap, Jochen Hampe, and Jan G. Hengstler. The ascending pathophysiology of cholestatic liver disease. *Hepatology*, 65(2):722–738, 2017. ISSN 0270-9139. doi: 10.1002/hep.28965.
- Tom H. Karlsen, Trine Folseraas, Douglas Thorburn, and Mette Vesterhus. Primary sclerosing cholangitis – a comprehensive review. *Journal of Hepatology*, 67(6):1298–1323, 2017. ISSN 0168-8278. doi: 10.1016/j.jhep.2017.07.022.
- Tiangang Li and John Y. L. Chiang. Bile Acid Signaling in Liver Metabolism and Diseases. *Journal of Lipids*, 2012:754067, 2012. ISSN 2090-3030. doi: 10.1155/2012/754067.
- Hua Liang, Hongyu Miao, and Hulin Wu. Estimation of constant and time-varying dynamic parameters of HIV infection in a nonlinear differential equation model. *The Annals of Applied Statistics*, 4(1):460–483, 2010. ISSN 1932-6157. doi: 10.1214/09-aos290.
- Malgorzata Milkiewicz, Marta Klak, Agnieszka Kempinska-Podhorodecka, Anna Wiechowska-Kozłowska, Elzbieta Urasinska, Malgorzata Blatkiewicz, Ewa Wunsch, Elwyn Elias, and Piotr Milkiewicz. Impaired hepatic adaptation to chronic cholestasis induced by primary sclerosing cholangitis. *Scientific Reports*, 6(1):39573, 2016. doi: 10.1038/srep39573.
- G. Molino, A. F. Hofmann, C. Cravetto, G. Belforte, and B. Bona. Simulation of the metabolism and enterohepatic circulation of endogenous chenodeoxycholic acid in man using a physiological pharmacokinetic model. *European Journal of Clinical Investigation*, 16(5):397–414, 1986. ISSN 0014-2972. doi: 10.1111/j.1365-2362.1986.tb01015.x.
- Omar Y. Mousa, Brian D. Juran, Bryan M. McCauley, Mette N. Vesterhus, Trine Folseraas, Coleman T. Turgeon, Ahmad H. Ali, Erik M. Schlicht, Elizabeth J. Atkinson, Chang Hu, Denise Harnois, Elizabeth J. Carey, Andrea A. Gossard, Devin Oglesbee, John E. Eaton, Nicholas F. LaRusso, Gregory J. Gores, Tom H. Karlsen, and Konstantinos N. Lazaridis. Bile Acid Profiles in Primary Sclerosing Cholangitis and Their Ability to Predict Hepatic Decompensation. *Hepatology*, 74(1):281–295, 2021. ISSN 0270-9139. doi: 10.1002/hep.31652.
- Hans Popper and Fenton Schaffner. Pathophysiology of cholestasis. *Human Pathology*, 1(1):1–24, 1970. ISSN 0046-8177. doi: 10.1016/s0046-8177(70)80002-8.
- Krishnakant Saboo, Anirudh Choudhary, Yurui Cao, Gregory Worrell, David Jones, and Ravishankar Iyer. Reinforcement learning based disease progression model for alzheimer’s disease. *Advances in Neural Information Processing Systems*, 34:20903–20915, 2021.
- Yonatan Savir, Alexander Martynov, and Michael Springer. Achieving global perfect homeostasis through transporter regulation. *PLoS Computational Biology*, 13(4):e1005458, 2017. ISSN 1553-734X. doi: 10.1371/journal.pcbi.1005458.
- Fianne L. P. Sips, Hannah M. Eggink, Peter A. J. Hilbers, Maarten R. Soeters, Albert K. Groen, and Natal A. W. van Riel. In silico analysis identifies intestinal transit as a key determinant of systemic bile acid metabolism. *Frontiers in Physiology*, 9:631, 2018. ISSN 1664-042X. doi: 10.3389/fphys.2018.00631.
- Kelly A. Tappenden. Intestinal Adaptation Following Resection. *Journal of Parenteral and Enteral Nutrition*, 38(1_suppl):23S–31S, 2014. ISSN 0148-6071. doi: 10.1177/0148607114525210.
- Byron P. Vaughn, Thomas Kaiser, Christopher Staley, Matthew J. Hamilton, Jon Reich, Carolyn Graiziger, Stephanie Singroy, Amanda J. Kabage, Michael J. Sadowsky, and Alexander Khoruts. A

pilot study of fecal bile acid and microbiota profiles in inflammatory bowel disease and primary sclerosing cholangitis. *Clinical and Experimental Gastroenterology*, 12:9–19, 2019. ISSN 1178-7023. doi: 10.2147/ceg.s186097.

Mette Vesterhus and Tom Hemming Karlsen. Emerging therapies in primary sclerosing cholangitis: pathophysiological basis and clinical opportunities. *Journal of Gastroenterology*, 55(6):588–614, 2020. ISSN 0944-1174. doi: 10.1007/s00535-020-01681-z.

Veronika Voronova, Victor Sokolov, Amani Al-Khaifi, Sara Straniero, Chanchal Kumar, Kirill Peskov, Gabriel Helmlinger, Mats Rudling, and Bo Angelin. A Physiology-Based Model of Bile Acid Distribution and Metabolism Under Healthy and Pathologic Conditions in Human Beings. *Cellular and Molecular Gastroenterology and Hepatology*, 10(1):149–170, 2020. ISSN 2352-345X. doi: 10.1016/j.jcmgh.2020.02.005.

Annika Wahlström, Sama I Sayin, Hanns-Ulrich Marschall, and Fredrik Bäckhed. Intestinal crosstalk between bile acids and microbiota and its impact on host metabolism. *Cell metabolism*, 24(1):41–50, 2016.

Yifeng Wang, Wen-Xing Ding, and Tiangang Li. Cholesterol and bile acid-mediated regulation of autophagy in fatty liver diseases and atherosclerosis. *Biochimica et Biophysica Acta (BBA) - Molecular and Cell Biology of Lipids*, 1863(7):726–733, 2018. ISSN 1388-1981. doi: 10.1016/j.bbali.2018.04.005.

Yanni Xiao, Hongyu Miao, Sanyi Tang, and Hulin Wu. Modeling antiretroviral drug responses for hiv-1 infected patients using differential equation models. *Advanced drug delivery reviews*, 65(7):940–953, 2013.

Jingsong Zhang, Jessica J Cunningham, Joel S Brown, and Robert A Gatenby. Integrating evolutionary dynamics into treatment of metastatic castrate-resistant prostate cancer. *Nature communications*, 8(1):1816, 2017.

Appendix A. Reduced-Order Bile Acid Metabolism Model Under Healthy Conditions

A.1. Bile Acids Enterohepatic Circulation




Figure 6: Overview of the biochemical and physical processes modeled in the bile acids enterohepatic circulation. We use primary bile acids (PBA) to represent both cholic acid (CA) and chenodeoxycholic acid (CDCA) because they share the same processes.

Several species of bile acids circulate in multiple organs of the human body in a process called enterohepatic circulation. In the liver (li), cholesterol is metabolized into primary bile acids unconjugated cholic acid (uCA), and unconjugated chenodeoxycholic acid (uCDCA), which are then conjugated with either glycine or taurine into conjugated cholic acid (cCA) and conjugated chenodeoxycholic acid (cCDCA). Conjugated primary bile acids are then secreted into the bile duct and gallbladder (bd), which store and secrete the bile acids into the intestine, mostly after meals. Bacteria in the ileum (il) and the colon (co) deconjugate some of these bile acids and colon bacteria dehydroxylate primary bile acids to form secondary bile acids (SBA). Through active and passive uptake, bile acids are reabsorbed into the portal vein and then back to the liver. Active uptake happens in the ileum, while passive uptake happens in both the ileum and the colon, but only for unconjugated bile acids. Liver cells extract the

reabsorbed bile acids efficiently, with a small portion of various bile acid species escaping into the systemic circulation (plasma, pl). Bile acids that are not reabsorbed in the intestines are excreted with feces (fe).

A.2. Model Reduction from *Sips et al.* (2018)

[Sips et al. \(2018\)](#) presents a detailed bile acid model that is calibrated with a comprehensive dataset compiled from 99 studies. For each bile acid (BA) in each organ (OG), [Sips et al. \(2018\)](#) (1) models its influxes/outfluxes from relevant biochemical and physical processes and (2) joins the fluxes into one differential equation to describe how the level of BA in OG varies with time. We adopt [Sips et al. \(2018\)](#)'s approach and model bile acid circulation with a series of ordinary differential equations (ODEs). Based on relevance to PSC, we reduce [Sips et al. \(2018\)](#)'s model by merging several bile acid species and not distinguishing certain organ segments. Specifically, [Sips et al. \(2018\)](#) divided the intestines into 15 consecutive segments for a detailed description of postprandial dynamics, which we deem irrelevant for PSC and combine the segments the ileum and the colon; [Sips et al. \(2018\)](#) modeled 8 different cSBA species separately, while we combine them together because SBA as a whole is hypothesized to play a role in PSC ([Vaughn et al. \(2019\)](#)). Our reduction results in the modeling of bile acid $BA \in \{cCA, cCDCA, cSBA, uCA, uCDCA, uSBA\}$ in organ $OG \in \{li, bd, il, co, pl, fe\}$. Figure 6 summarizes the biochemical and physical processes in the enterohepatic circulation that constitute the ODEs.

A.3. Meal-related Dynamics

In addition to the zero and first-order dynamics discussed in the main text, our ODEs also contain time-varying parameters that capture the nuances of meal-related dynamics. When we eat, the gallbladder contracts and releases bile acids, and the bowels also move faster. As a result, bile acids move through these compartments at an increased pace immediately after eating, gradually slowing down over time. To model these meal-dependent dynamics, we assume the parameters governing the first-order dynamics of the transit rates along these compartments are functions of the time passed since the last meal, parametrized by normalized Rayleigh functions. Specifically, meal-dependent dynamics affect transit from the gallbladder and bile duct to the

ileum, from the ileum to the colon, and from the colon to feces. As an example,

$$r_{co\ to\ fe}^{cCA} = p[co_to_fe_freq] \cdot p[co_transit_coef] \cdot x_{co}^{cCA} \quad (10)$$

$$p[co_transit_coef] = 1 + p[peak_co] \cdot \left(\frac{t_{meal} \sqrt{e}}{p[loc_co]} \exp\left(-\frac{t_{meal}}{2 \cdot p[loc_co]^2}\right) \right) \quad (11)$$

t_{meal} represents the time that has passed since the subject's most recent meal. As a result of the normalized Rayleigh function, $p[co_transit_coef]$ starts to rise from 1 immediately after a meal, reaching its postprandial peak of $1 + p[peak_co]$ at time $p[loc_co]$. Subsequently, it gradually returns to 1.

A.4. Estimates of Model Parameters and Fasting Bile Acid Levels

We utilize the model fitting and simulation results in [Sips et al. \(2018\)](#) to derive the parameters and fasting bile acid levels for a healthy individual for our model. For example, since $p[synthesis]$ has the same definition in [Sips et al. \(2018\)](#) as in our model, we directly take its value from [Sips et al. \(2018\)](#). However, there is no equivalent variable in [Sips et al. \(2018\)](#) for x_{co}^{cCA} , which represents the level of conjugated cholic acid in the colon. Instead, [Sips et al. \(2018\)](#) separates conjugated cholic acid into taurine-conjugated cholic acid and glycine-conjugated cholic acid and models their levels in 5 consecutive parts of the colon. Therefore, we derive our fasting x_{co}^{cCA} by summing up taurine-conjugated cholic acid and glycine-conjugated cholic acid in the 5 parts of the colon using the fasting time simulation results of [Sips et al. \(2018\)](#). Similarly, we estimate fasting $r_{co\ to\ fe}^{cCA}$ by summing up the fluxes of taurine-conjugated cholic acid and glycine-conjugated cholic acid that transit from the last part of the colon to feces. Because $p[co_transit_coef]$ is assumed to be close to 1 at fasting, we derive $p[co_to_fe_freq]$ by dividing fasting $r_{co\ to\ fe}$ by fasting x_{co} and taking the average across bile acid species. Other parameters and bile acid levels are similarly estimated.

A.5. Specification of Bile Acid Dynamics Simulation under Healthy Conditions

After obtaining the estimations for all parameters and fasting bile acid levels, we run the system of ODEs using the Runge-Kutta 45 (RK45) numerical solver. To mimic the meal-dependent characteristics of bile acid metabolism, we introduce three simulated meals at 0, 6, and 12 hours every 24 hours, corresponding to breakfast, lunch, and dinner (8 AM, 2 PM, and 8 PM). We simulate the ODE system for a duration of 60 days, surpassing the 55-day duration in [Sips et al. \(2018\)](#). We verify that a steady state has been reached by day 60, where the bile acids dynamics approximately repeat themselves every 24 hours. Specifically, the largest relative change in a single bile acid is below 0.01% between two consecutive days and remains below 0.1% after a year (360 days). We then take the steady-state fasting bile acid levels as the fasting bile acid levels representative of healthy individuals.

A.6. 24-hour Bile Acid Dynamics

Figure 7 displays the 24-hour bile acid dynamics with three simulated meals at 8 AM, 2 PM, and 8 PM. The simulated results show that different bile acid species in different compartments all exhibit three-peak patterns that are associated with the three simulated meals. However, as expected, the peaking patterns and peaking amplitudes differ across bile acid species and compartments, validating the importance of modeling them separately. In general, conjugated bile acids in the same compartment follow similar trends, as do unconjugated bile acids, except for uSBA in the colon. Consistent with experimental observations, our results show that the gallbladder and bile duct contains the highest amount of conjugated bile acids among all compartments; uSBA is primarily concentrated in the colon; and plasma bile acid levels are generally low for all bile acid species due to the liver’s effective hepatic clearance under healthy conditions. Our simulation results also show good correspondence with the well-calibrated [Sips et al. \(2018\)](#) while reducing the computation time by more than 10X.

A.7. Validation with A Real-world Dataset

We validate the accuracy of our model by comparing the model simulation results with a real-world dataset from our partnering hospital. The dataset collected

fasting plasma bile acid levels from 302 healthy individuals. Measurements below the detectable threshold were imputed with half of the smallest non-zero value. We take the model plasma bile acid levels at 8 AM, 12 hours after the previous simulated meal, as the corresponding fasting levels for comparison. Both the model simulated results and the real-world data were log 10 transformed. Figure 8 shows that the simulated fasting bile acid levels are located within high likelihood regions of real-world data for all bile acids species, indicating that our model is capable of generating realistic bile acid levels. To further evaluate the representativeness of our simulation results beyond univariate comparison, we estimated the joint distribution of pairs of bile acids from the healthy control data and plotted the simulated data on the same plot in Figure 9. We see the simulated fasting profile is still located within high-likelihood regions when pairs of bile acids are considered simultaneously. These validation results provide confidence in the model’s ability to simulate bile acid dynamics under healthy conditions and encourage its usage as a base model for studying study bile acid-related diseases.

Appendix B. Reduced-Order Bile Acid Metabolism ODEs for PSC (Mathematical Descriptions)

B.1. Abbreviations

- cCA: conjugated cholic acid
- cCDCA: conjugated chenodeoxycholic acid
- cSBA: conjugated secondary bile acid
- uCA: unconjugated cholic acid
- uCDCA: unconjugated chenodeoxycholic acid
- uSBA: unconjugated secondary bile acid
- li: liver
- bd: bild duct
- il: ileum
- co: colon
- pl: plasma
- fe: feces




Figure 7: Simulated 24-hour bile acid dynamics with meals at 8 AM, 2 PM, and 8 PM. Orange curve: dynamics from our reduced-order ODEs; blue curve: dynamics from [Sips et al. \(2018\)](#).




Figure 8: \log_{10} fasting plasma bile acid level (red vertical line) from the reduced-order ODEs compared to real-world data histograms.

- syn: synthesis
- deconj: deconjugation
- biotr: primary to secondary bile acids bio-transformation
- au: active uptake
- pu: passive uptake

B.2. Definition of independent parameters

- $p[\text{synthesis}]$: $\mu\text{mol}/\text{min}$. Total bile acid (BA) synthesis (cCA, cCDCA) rate in li.
- $p[\text{syn_frac_CA}] \in [0, 1]$, unitless. Fraction of CA among the newly synthesized BA.
- $p[\text{li_to_bd_freq}]$: $1/\text{min}$. Parameter for BA secretion from li to bd.
- $p[\text{bd_to_il_freq}]$: $1/\text{min}$. Parameter for BA transit from bd to il.
- $p[\text{il_to_co_freq}]$: $1/\text{min}$. Parameter for BA transit from il to co.
- $p[\text{co_to_fe_freq}]$: $1/\text{min}$. Parameter for BA excretion from co to fe.

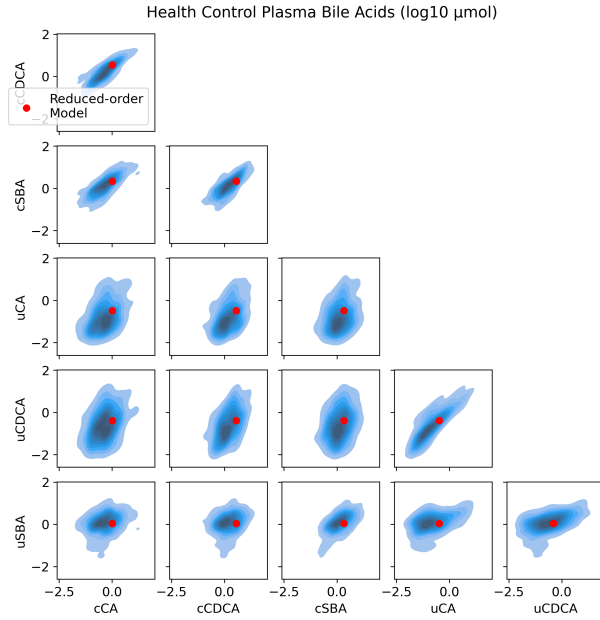


Figure 9: \log_{10} fasting plasma level of bile acid pairs (red dot) from the reduced-order ODEs compared to real-world data bile acid pairs joint distribution.

- $p[\text{pl_volume}]$: L. Total volume of pl in the body.
- $p[\text{pl_flow_thru_li}]$: L/min. Rate of pl flowing through li.
- $p[\text{hep_ratio_conj_tri}] \in [0, 1]$, unitless. For conjugated trihydroxylated BA (cCA) extracted from the gut via passive uptake and active uptake, the proportion that goes to li; for conjugated trihydroxylated BA (cCA) flowing through li from pl, the proportion that goes to li.
- $p[\text{hep_ratio_conj_di}] \in [0, 1]$, unitless. Same as $p[\text{hep_ratio_conj_tri}]$, but for conjugated dihydroxylated BA (primary BA CDCA, secondary BA ursodeoxycholic acid and deoxycholic acid). For simplicity, we assume the other secondary BA, mono-hydroxylated BA lithocholic acid, also belongs to this category. Hence, $p[\text{hep_ratio_conj_di}]$ applies to cCDCA and cSBA.
- $p[\text{hep_ratio_unconj_to_conj}] \in [0, 1]$, unitless. Unconjugated BAs have lower hepatic extraction ratios than conjugated BAs. This parameter quantifies the fraction $\frac{p[\text{hep_ratio_unconj_tri/di}]}{p[\text{hep_ratio_conj_tri/di}]}$.
- $p[\text{gut_deconj_freq_co}]$: 1/min. Parameter for BA deconjugation in co (by bacteria).
- $p[\text{gut_deconj_freq_il_to_co}] \in [0, 1]$, unitless. Deconjugation happens more in co than in il. This parameter quantifies the fraction $\frac{p[\text{gut_deconj_freq_il}]}{p[\text{gut_deconj_freq_co}]}$.
- $p[\text{gut_biotr_freq_CA}]$: 1/min. Parameter for uCA to uSBA biotransformation in co (by bacteria).
- $p[\text{gut_biotr_freq_CDCA_to_CA}] \in [0, 1]$, unitless. Biotransformation to secondary BA happens (slightly) faster for CA than for CDCA. This parameter quantifies the fraction $\frac{p[\text{gut_biotr_freq_CDCA}]}{p[\text{gut_biotr_freq_CA}]}$.
- $p[\text{max_asbt_rate}]$: 1/min. Parameter for BA active uptake in il (via enzyme ASBT).
- $p[\text{gut_pu_freq_co}]$: 1/min. Parameter for uBA passive uptake in co.
- $p[\text{gut_pu_freq_il_to_co}] \in [0, 1]$, unitless. Passive uptake is more efficient in co than in il because of the larger diameter of co. This parameter quantifies the fraction $\frac{p[\text{gut_pu_freq_il}]}{p[\text{gut_pu_freq_co}]}$.

- $p[\text{meal_reflex_loc_bd}]$: min. Parameter characterizing the normalized Rayleigh function that describes the meal-dependent BA transit dynamics from bd to il. BA transit efficiency from bd to il peaks at $p[\text{meal_reflex_loc_bd}]$ min after the most recent meal, while baseline efficiency happens immediately after a meal or after an infinitely extended period.
- $p[\text{meal_reflex_loc_il}]$: min. Same as $p[\text{meal_reflex_loc_bd}]$, but for il to co transit.
- $p[\text{meal_reflex_loc_co}]$: min. Same as $p[\text{meal_reflex_loc_bd}]$, but for co to fe transit.
- $p[\text{meal_reflex_peak_bd}]$: unitless. Parameter characterizing the normalized Rayleigh function that describes the meal-dependent BA transit dynamics from bd to il. At its peak, BA transit efficiency from bd to il is elevated by a factor of $1 + p[\text{meal_reflex_peak_bd}]$ compared to the baseline.
- $p[\text{meal_reflex_peak_il}]$: min. Same as $p[\text{meal_reflex_peak_bd}]$, but for il to co transit.
- $p[\text{meal_reflex_peak_co}]$: min. Same as $p[\text{meal_reflex_peak_bd}]$, but for co to fe transit.
- $p[\text{co_sulfate_freq}]$: 1/min. Parameter for sulfation of secondary BA (LCA). For simplicity, we account for sulfation by modeling its end outcome: the additional excretion of SBA into fe, which is again modeled with first-order dynamics, governed by $p[\text{co_sulfate_freq}]$.
- $p[\text{bd_max_ba}]$: μmol . BA holding capacity of bd. Excessive BAs will backflow to li proportional to their levels in bd.
- $p[\text{bd_max_flow}]$: μmol . Maximum amount of BA allowed to pass through bd to il, in proportion to the degree of bd obstruction.

B.3. ODE joining the bile acid fluxes

How the level of each bile acid $BA \in \{cCA, cCDCA, cSBA, uCA, uCDCA, uSBA\}$ in each organ $OG \in \{\text{li, bd, il, co, pl, fe}\}$ varies with time is described by differential equations summing the corresponding influxes and subtracting the outfluxes. The fluxes take unit $\mu\text{mol}/\text{min}$ unless specified otherwise.

B.3.1. LIVER

$$\begin{aligned} \frac{dx_{\text{li}}^{cCA}}{dt} &= r_{\text{syn}}^{cCA} + r_{\text{li from gut}}^{cCA} + r_{\text{li from pl}}^{cCA} \\ &\quad - r_{\text{li to bd}}^{cCA} + r_{\text{li from bd}}^{cCA} \\ \frac{dx_{\text{li}}^{cCDCA}}{dt} &= r_{\text{syn}}^{cCDCA} + r_{\text{li from gut}}^{cCDCA} + r_{\text{li from pl}}^{cCDCA} \\ &\quad - r_{\text{li to bd}}^{cCDCA} + r_{\text{li from bd}}^{cCDCA} \\ \frac{dx_{\text{li}}^{cSBA}}{dt} &= r_{\text{li from gut}}^{cSBA} + r_{\text{li from pl}}^{cSBA} - r_{\text{li to bd}}^{cSBA} \\ &\quad + r_{\text{li from bd}}^{cSBA} \\ \frac{dx_{\text{li}}^{uCA}}{dt} &= 0 \\ \frac{dx_{\text{li}}^{uCDCA}}{dt} &= 0 \\ \frac{dx_{\text{li}}^{uSBA}}{dt} &= 0 \end{aligned}$$

B.3.2. BILE DUCT

$$\begin{aligned} \frac{dx_{\text{bd}}^{cCA}}{dt} &= r_{\text{bd from li}}^{cCA} - r_{\text{bd to il}}^{cCA} - r_{\text{bd to li}}^{cCA} \\ \frac{dx_{\text{bd}}^{cCDCA}}{dt} &= r_{\text{bd from li}}^{cCDCA} - r_{\text{bd to il}}^{cCDCA} - r_{\text{bd to li}}^{cCDCA} \\ \frac{dx_{\text{bd}}^{cSBA}}{dt} &= r_{\text{bd from li}}^{cSBA} - r_{\text{bd to il}}^{cSBA} - r_{\text{bd to li}}^{cSBA} \\ \frac{dx_{\text{bd}}^{uCA}}{dt} &= 0 \\ \frac{dx_{\text{bd}}^{uCDCA}}{dt} &= 0 \\ \frac{dx_{\text{bd}}^{uSBA}}{dt} &= 0 \end{aligned}$$

B.3.3. ILEUM

$$\begin{aligned} \frac{dx_{\text{il}}^{cCA}}{dt} &= r_{\text{il from bd}}^{cCA} - r_{\text{il deconj}}^{cCA} - r_{\text{il au}}^{cCA} - r_{\text{il to co}}^{cCA} \\ \frac{dx_{\text{il}}^{cCDCA}}{dt} &= r_{\text{il from bd}}^{cCDCA} - r_{\text{il deconj}}^{cCDCA} - r_{\text{il au}}^{cCDCA} - r_{\text{il to co}}^{cCDCA} \\ \frac{dx_{\text{il}}^{cSBA}}{dt} &= r_{\text{il from bd}}^{cSBA} - r_{\text{il deconj}}^{cSBA} - r_{\text{il au}}^{cSBA} - r_{\text{il to co}}^{cSBA} \\ \frac{dx_{\text{il}}^{uCA}}{dt} &= r_{\text{il deconj}}^{uCA} - r_{\text{il au}}^{uCA} - r_{\text{il pu}}^{uCA} - r_{\text{il to co}}^{uCA} \\ \frac{dx_{\text{il}}^{uCDCA}}{dt} &= r_{\text{il deconj}}^{uCDCA} - r_{\text{il au}}^{uCDCA} - r_{\text{il pu}}^{uCDCA} - r_{\text{il to co}}^{uCDCA} \\ \frac{dx_{\text{il}}^{uSBA}}{dt} &= r_{\text{il deconj}}^{uSBA} - r_{\text{il au}}^{uSBA} - r_{\text{il pu}}^{uSBA} - r_{\text{il to co}}^{uSBA} \end{aligned}$$

B.3.4. COLON

$$\begin{aligned}
 \frac{dx_{\text{co}}^{\text{cCA}}}{dt} &= r_{\text{co from il}}^{\text{cCA}} - r_{\text{co deconj}}^{\text{cCA}} - r_{\text{co to fe}}^{\text{cCA}} \\
 \frac{dx_{\text{co}}^{\text{cCDCA}}}{dt} &= r_{\text{co from il}}^{\text{cCDCA}} - r_{\text{co deconj}}^{\text{cCDCA}} - r_{\text{co to fe}}^{\text{cCDCA}} \\
 \frac{dx_{\text{co}}^{\text{cSBA}}}{dt} &= r_{\text{co from il}}^{\text{cSBA}} - r_{\text{co deconj}}^{\text{cSBA}} - r_{\text{co to fe}}^{\text{cSBA}} \\
 &\quad - r_{\text{co sulfate}}^{\text{cSBA}} \\
 \frac{dx_{\text{co}}^{\text{uCA}}}{dt} &= r_{\text{co from il}}^{\text{uCA}} + r_{\text{co deconj}}^{\text{uCA}} - r_{\text{co biotr}}^{\text{uCA}} \\
 &\quad - r_{\text{co pu}}^{\text{uCA}} - r_{\text{co to fe}}^{\text{uCA}} \\
 \frac{dx_{\text{co}}^{\text{uCDCA}}}{dt} &= r_{\text{co from il}}^{\text{uCDCA}} + r_{\text{co deconj}}^{\text{uCDCA}} - r_{\text{co biotr}}^{\text{uCDCA}} \\
 &\quad - r_{\text{co pu}}^{\text{uCDCA}} - r_{\text{co to fe}}^{\text{uCDCA}} \\
 \frac{dx_{\text{co}}^{\text{uSBA}}}{dt} &= r_{\text{co from il}}^{\text{uSBA}} + r_{\text{co deconj}}^{\text{uSBA}} + r_{\text{co biotr}}^{\text{uSBA}} \\
 &\quad - r_{\text{co pu}}^{\text{uSBA}} - r_{\text{co to fe}}^{\text{uSBA}} - r_{\text{co sulfate}}^{\text{uSBA}}
 \end{aligned}$$

B.3.5. PLASMA

$$\begin{aligned}
 \frac{dx_{\text{pl}}^{\text{cCA}}}{dt} &= r_{\text{pl from gut}}^{\text{cCA}} - r_{\text{pl to li}}^{\text{cCA}} \\
 \frac{dx_{\text{pl}}^{\text{cCDCA}}}{dt} &= r_{\text{pl from gut}}^{\text{cCDCA}} - r_{\text{pl to li}}^{\text{cCDCA}} \\
 \frac{dx_{\text{pl}}^{\text{cSBA}}}{dt} &= r_{\text{pl from gut}}^{\text{cSBA}} - r_{\text{pl to li}}^{\text{cSBA}} \\
 \frac{dx_{\text{pl}}^{\text{uCA}}}{dt} &= r_{\text{pl from gut}}^{\text{uCA}} - r_{\text{pl to li}}^{\text{uCA}} \\
 \frac{dx_{\text{pl}}^{\text{uCDCA}}}{dt} &= r_{\text{pl from gut}}^{\text{uCDCA}} - r_{\text{pl to li}}^{\text{uCDCA}} \\
 \frac{dx_{\text{pl}}^{\text{uSBA}}}{dt} &= r_{\text{pl from gut}}^{\text{uSBA}} - r_{\text{pl to li}}^{\text{uSBA}}
 \end{aligned}$$

B.3.6. FECES

$$\begin{aligned}
 \frac{dx_{\text{fe}}^{\text{cCA}}}{dt} &= r_{\text{fe from co}}^{\text{cCA}} \\
 \frac{dx_{\text{fe}}^{\text{cCDCA}}}{dt} &= r_{\text{fe from co}}^{\text{cCDCA}} \\
 \frac{dx_{\text{fe}}^{\text{cSBA}}}{dt} &= r_{\text{fe from co}}^{\text{cSBA}} + r_{\text{fe sulfate}}^{\text{cSBA}} \\
 \frac{dx_{\text{fe}}^{\text{uCA}}}{dt} &= r_{\text{fe from co}}^{\text{uCA}} \\
 \frac{dx_{\text{fe}}^{\text{uCDCA}}}{dt} &= r_{\text{fe from co}}^{\text{uCDCA}} \\
 \frac{dx_{\text{fe}}^{\text{uSBA}}}{dt} &= r_{\text{fe from co}}^{\text{uSBA}} + r_{\text{fe sulfate}}^{\text{uSBA}}
 \end{aligned}$$

B.4. Helper functions and parameters

We define the following helper functions and parameters as functions of the estimated parameters:

B.4.1. HELPERS FOR BILE DUCT OBSTRUCTION AND BILE ACID BACKFLOW

$$\begin{aligned}
 x_{\text{bd}} &= x_{\text{bd}}^{\text{cCA}} + x_{\text{bd}}^{\text{cCDCA}} + x_{\text{bd}}^{\text{cSBA}} \\
 r_{\text{bd from li}} &= r_{\text{bd from li}}^{\text{cCA}} + r_{\text{bd from li}}^{\text{cCDCA}} \\
 &\quad + r_{\text{bd from li}}^{\text{cSBA}} \\
 r_{\text{bd to il}} &= r_{\text{bd to il}}^{\text{cCA}} + r_{\text{bd to il}}^{\text{cCDCA}} + r_{\text{bd to il}}^{\text{cSBA}}
 \end{aligned}$$

$$\begin{aligned}
 p[\text{bd_discount}] &= \\
 \min &\left(\frac{p[\text{bd_max_flow}]}{p[\text{bd_to_il_freq}] \cdot p[\text{bd_transit_coef}] \cdot x_{\text{bd}}}, 1 \right)
 \end{aligned}$$

$$\begin{aligned}
 p[\text{bd_backflow_coef}] &= \\
 \max &\left(\frac{x_{\text{bd}} + r_{\text{bd from li}} - r_{\text{bd to il}} - p[\text{bd_max_ba}]}{x_{\text{bd}}}, 0 \right)
 \end{aligned}$$

B.4.2. HELPERS FOR PARAMETERS DEFINED BY RATIOS

$$\begin{aligned}
 p[\text{hep_ratio_unconj_tri}] &= \\
 p[\text{hep_ratio_conj_tri}] \cdot p[\text{hep_ratio_unconj_to_conj}] & \\
 p[\text{hep_ratio_unconj_di}] &= \\
 p[\text{hep_ratio_conj_di}] \cdot p[\text{hep_ratio_unconj_to_conj}] & \\
 p[\text{gut_deconj_freq_il}] &= \\
 p[\text{gut_deconj_freq_co}] \cdot p[\text{gut_deconj_freq_il_to_co}] & \\
 p[\text{gut_biotr_freq_CDCA}] &= \\
 p[\text{gut_biotr_freq_CA}] \cdot p[\text{gut_biotr_freq_CDCA_to_CA}] & \\
 p[\text{gut_pu_freq_il}] &= \\
 p[\text{gut_pu_freq_co}] \cdot p[\text{gut_pu_freq_il_to_co}] &
 \end{aligned}$$

B.4.3. HELPERS FOR MEAL-RELATED DYNAMICS

$$\begin{aligned}
 \text{GL_reflex}(loc, peak, t) &= 1 + peak \\
 &\cdot \left(\frac{t\sqrt{e}}{loc} \exp\left(-\frac{t^2}{2 \cdot loc^2}\right) \right) \\
 \text{GL_reflex}(loc, peak, 0) &= 1 \\
 \text{GL_reflex}(loc, peak, loc) &= 1 + peak \\
 \text{GL_reflex}(loc, peak, \infty) &= 1
 \end{aligned}$$

$$\begin{aligned}
 p[\text{bd_transit_coef}] &= \text{GL_reflex}(p[\text{meal_reflex_loc_bd}], \\
 p[\text{meal_reflex_peak_bd}], t_{\text{since_meal}}) & \\
 p[\text{il_transit_coef}] &= \text{GL_reflex}(p[\text{meal_reflex_loc_il}], \\
 p[\text{meal_reflex_peak_il}], t_{\text{since_meal}}) & \\
 p[\text{co_transit_coef}] &= \text{GL_reflex}(p[\text{meal_reflex_loc_co}], \\
 p[\text{meal_reflex_peak_co}], t_{\text{since_meal}}) &
 \end{aligned}$$

B.5. Zero- and first-order dynamics defining the fluxes

B.5.1. SYNTHESIS FLUXES

$$\begin{aligned}
 r_{\text{syn}}^{\text{cCA}} &= p[\text{synthesis}] \cdot p[\text{syn_frac_CA}] \\
 r_{\text{syn}}^{\text{cCDCA}} &= p[\text{synthesis}] \cdot (1 - p[\text{syn_frac_CA}])
 \end{aligned}$$

B.5.2. TRANSIT FLUXES

Transit fluxes are further categorized based on the organs involved in the transit process:

Liver vs. Bile Duct

$$\begin{aligned}
 r_{\text{li to bd}}^{\text{cCA}} &= p[\text{li_to_bd_freq}] \cdot x_{\text{li}}^{\text{cCA}} \\
 r_{\text{li to bd}}^{\text{cCDCA}} &= p[\text{li_to_bd_freq}] \cdot x_{\text{li}}^{\text{cCDCA}} \\
 r_{\text{li to bd}}^{\text{cSBA}} &= p[\text{li_to_bd_freq}] \cdot x_{\text{li}}^{\text{cSBA}} \\
 r_{\text{bd from li}}^{\text{cCA}} &= r_{\text{li to bd}}^{\text{cCA}} \\
 r_{\text{bd from li}}^{\text{cCDCA}} &= r_{\text{li to bd}}^{\text{cCDCA}} \\
 r_{\text{bd from li}}^{\text{cSBA}} &= r_{\text{li to bd}}^{\text{cSBA}} \\
 r_{\text{bd to li}}^{\text{cCA}} &= p[\text{bd_backflow_coef}] \cdot x_{\text{bd}}^{\text{cCA}} \\
 r_{\text{bd to li}}^{\text{cCDCA}} &= p[\text{bd_backflow_coef}] \cdot x_{\text{bd}}^{\text{cCDCA}} \\
 r_{\text{bd to li}}^{\text{cSBA}} &= p[\text{bd_backflow_coef}] \cdot x_{\text{bd}}^{\text{cSBA}} \\
 r_{\text{li from bd}}^{\text{cCA}} &= r_{\text{bd to li}}^{\text{cCA}} \\
 r_{\text{li from bd}}^{\text{cCDCA}} &= r_{\text{bd to li}}^{\text{cCDCA}} \\
 r_{\text{li from bd}}^{\text{cSBA}} &= r_{\text{bd to li}}^{\text{cSBA}}
 \end{aligned}$$

Bile Duct vs. Ileum

$$\begin{aligned}
 r_{\text{bd to il}}^{\text{cCA}} &= p[\text{bd_discount}] \cdot p[\text{bd_to_il_freq}] \\
 &\cdot p[\text{bd_transit_coef}] \cdot x_{\text{bd}}^{\text{cCA}} \\
 r_{\text{bd to il}}^{\text{cCDCA}} &= p[\text{bd_discount}] \cdot p[\text{bd_to_il_freq}] \\
 &\cdot p[\text{bd_transit_coef}] \cdot x_{\text{bd}}^{\text{cCDCA}} \\
 r_{\text{bd to il}}^{\text{cSBA}} &= p[\text{bd_discount}] \cdot p[\text{bd_to_il_freq}] \\
 &\cdot p[\text{bd_transit_coef}] \cdot x_{\text{bd}}^{\text{cSBA}} \\
 r_{\text{il from bd}}^{\text{cCA}} &= r_{\text{bd to il}}^{\text{cCA}} \\
 r_{\text{il from bd}}^{\text{cCDCA}} &= r_{\text{bd to il}}^{\text{cCDCA}} \\
 r_{\text{il from bd}}^{\text{cSBA}} &= r_{\text{bd to il}}^{\text{cSBA}}
 \end{aligned}$$

Ileum vs. Colon

$$\begin{aligned}
 r_{\text{il to co}}^{\text{cCA}} &= p[\text{il_to_co_freq}] \cdot p[\text{il_transit_coef}] \cdot x_{\text{il}}^{\text{cCA}} \\
 r_{\text{il to co}}^{\text{cCDCA}} &= p[\text{il_to_co_freq}] \cdot p[\text{il_transit_coef}] \cdot x_{\text{il}}^{\text{cCDCA}} \\
 r_{\text{il to co}}^{\text{cSBA}} &= p[\text{il_to_co_freq}] \cdot p[\text{il_transit_coef}] \cdot x_{\text{il}}^{\text{cSBA}} \\
 r_{\text{il to co}}^{\text{uCA}} &= p[\text{il_to_co_freq}] \cdot p[\text{il_transit_coef}] \cdot x_{\text{il}}^{\text{uCA}} \\
 r_{\text{il to co}}^{\text{uCDCA}} &= p[\text{il_to_co_freq}] \cdot p[\text{il_transit_coef}] \cdot x_{\text{il}}^{\text{uCDCA}} \\
 r_{\text{il to co}}^{\text{uSBA}} &= p[\text{il_to_co_freq}] \cdot p[\text{il_transit_coef}] \cdot x_{\text{il}}^{\text{uSBA}} \\
 r_{\text{co from il}}^{\text{cCA}} &= r_{\text{il to co}}^{\text{cCA}} \\
 r_{\text{co from il}}^{\text{cCDCA}} &= r_{\text{il to co}}^{\text{cCDCA}} \\
 r_{\text{co from il}}^{\text{cSBA}} &= r_{\text{il to co}}^{\text{cSBA}} \\
 r_{\text{co from il}}^{\text{uCA}} &= r_{\text{il to co}}^{\text{uCA}} \\
 r_{\text{co from il}}^{\text{uCDCA}} &= r_{\text{il to co}}^{\text{uCDCA}} \\
 r_{\text{co from il}}^{\text{uSBA}} &= r_{\text{il to co}}^{\text{uSBA}}
 \end{aligned}$$

Plasma vs. Liver

$$\begin{aligned}
 r_{\text{pl to li}}^{\text{cCA}} &= p[\text{hep_ratio_conj_tri}] \\
 &\cdot p[\text{pl_flow_thru_li}] \cdot \frac{x_{\text{pl}}^{\text{cCA}}}{p[\text{pl_volume}]} \\
 r_{\text{pl to li}}^{\text{cCDCA}} &= p[\text{hep_ratio_conj_di}] \\
 &\cdot p[\text{pl_flow_thru_li}] \cdot \frac{x_{\text{pl}}^{\text{cCDCA}}}{p[\text{pl_volume}]} \\
 r_{\text{pl to li}}^{\text{cSBA}} &= p[\text{hep_ratio_conj_di}] \\
 &\cdot p[\text{pl_flow_thru_li}] \cdot \frac{x_{\text{pl}}^{\text{cSBA}}}{p[\text{pl_volume}]} \\
 r_{\text{pl to li}}^{\text{uCA}} &= p[\text{hep_ratio_unconj_tri}] \\
 &\cdot p[\text{pl_flow_thru_li}] \cdot \frac{x_{\text{pl}}^{\text{uCA}}}{p[\text{pl_volume}]} \\
 r_{\text{pl to li}}^{\text{uCDCA}} &= p[\text{hep_ratio_unconj_di}] \\
 &\cdot p[\text{pl_flow_thru_li}] \cdot \frac{x_{\text{pl}}^{\text{uCDCA}}}{p[\text{pl_volume}]} \\
 r_{\text{pl to li}}^{\text{uSBA}} &= p[\text{hep_ratio_unconj_di}] \\
 &\cdot p[\text{pl_flow_thru_li}] \cdot \frac{x_{\text{pl}}^{\text{uSBA}}}{p[\text{pl_volume}]} \\
 r_{\text{li from pl}}^{\text{cCA}} &= r_{\text{pl to li}}^{\text{cCA}} + r_{\text{pl to li}}^{\text{uCA}} \\
 r_{\text{li from pl}}^{\text{cCDCA}} &= r_{\text{pl to li}}^{\text{cCDCA}} + r_{\text{pl to li}}^{\text{uCDCA}} \\
 r_{\text{li from pl}}^{\text{cSBA}} &= r_{\text{pl to li}}^{\text{cSBA}} + r_{\text{pl to li}}^{\text{uSBA}}
 \end{aligned}$$

B.5.3. DECONJUGATION FLUXES

$$\begin{aligned}
 r_{\text{il deconj}}^{\text{cCA}} &= p[\text{gut_deconj_freq_il}] \cdot x_{\text{il}}^{\text{cCA}} \\
 r_{\text{il deconj}}^{\text{cCDCA}} &= p[\text{gut_deconj_freq_il}] \cdot x_{\text{il}}^{\text{cCDCA}} \\
 r_{\text{il deconj}}^{\text{cSBA}} &= p[\text{gut_deconj_freq_il}] \cdot x_{\text{il}}^{\text{cSBA}} \\
 r_{\text{co deconj}}^{\text{cCA}} &= p[\text{gut_deconj_freq_co}] \cdot x_{\text{co}}^{\text{cCA}} \\
 r_{\text{co deconj}}^{\text{cCDCA}} &= p[\text{gut_deconj_freq_co}] \cdot x_{\text{co}}^{\text{cCDCA}} \\
 r_{\text{co deconj}}^{\text{cSBA}} &= p[\text{gut_deconj_freq_co}] \cdot x_{\text{co}}^{\text{cSBA}} \\
 r_{\text{il deconj}}^{\text{uCA}} &= r_{\text{il deconj}}^{\text{cCA}} \\
 r_{\text{il deconj}}^{\text{uCDCA}} &= r_{\text{il deconj}}^{\text{cCDCA}} \\
 r_{\text{il deconj}}^{\text{uSBA}} &= r_{\text{il deconj}}^{\text{cSBA}} \\
 r_{\text{co deconj}}^{\text{uCA}} &= r_{\text{co deconj}}^{\text{cCA}} \\
 r_{\text{co deconj}}^{\text{uCDCA}} &= r_{\text{co deconj}}^{\text{cCDCA}} \\
 r_{\text{co deconj}}^{\text{uSBA}} &= r_{\text{co deconj}}^{\text{cSBA}}
 \end{aligned}$$

B.5.4. BIOTRANSFORMATION FLUXES

$$\begin{aligned}
 r_{\text{co biotr}}^{\text{uCA}} &= p[\text{gut_biotr_freq_CA}] \cdot x_{\text{co}}^{\text{uCA}} \\
 r_{\text{co biotr}}^{\text{uCDCA}} &= p[\text{gut_biotr_freq_CA}] \cdot x_{\text{co}}^{\text{uCDCA}} \\
 r_{\text{co biotr}}^{\text{uSBA}} &= r_{\text{co biotr}}^{\text{uCA}} + r_{\text{co biotr}}^{\text{uCDCA}}
 \end{aligned}$$

B.5.5. ACTIVE UPTAKE FLUXES

$$\begin{aligned}
 r_{\text{il au}}^{\text{cCA}} &= p[\text{max_asbt_rate}] \cdot x_{\text{il}}^{\text{cCA}} \\
 r_{\text{il au}}^{\text{cCDCA}} &= p[\text{max_asbt_rate}] \cdot x_{\text{il}}^{\text{cCDCA}} \\
 r_{\text{il au}}^{\text{cSBA}} &= p[\text{max_asbt_rate}] \cdot x_{\text{il}}^{\text{cSBA}} \\
 r_{\text{il au}}^{\text{uCA}} &= p[\text{max_asbt_rate}] \cdot x_{\text{il}}^{\text{uCA}} \\
 r_{\text{il au}}^{\text{uCDCA}} &= p[\text{max_asbt_rate}] \cdot x_{\text{il}}^{\text{uCDCA}} \\
 r_{\text{il au}}^{\text{uSBA}} &= p[\text{max_asbt_rate}] \cdot x_{\text{il}}^{\text{uSBA}}
 \end{aligned}$$

B.5.6. PASSIVE UPTAKE FLUXES

$$\begin{aligned}
 r_{\text{il pu}}^{\text{uCA}} &= p[\text{gut_pu_freq_il}] \cdot x_{\text{il}}^{\text{uCA}} \\
 r_{\text{il pu}}^{\text{uCDCA}} &= p[\text{gut_pu_freq_il}] \cdot x_{\text{il}}^{\text{uCDCA}} \\
 r_{\text{il pu}}^{\text{uSBA}} &= p[\text{gut_pu_freq_il}] \cdot x_{\text{il}}^{\text{uSBA}} \\
 r_{\text{co pu}}^{\text{uCA}} &= p[\text{gut_pu_freq_co}] \cdot x_{\text{co}}^{\text{uCA}} \\
 r_{\text{co pu}}^{\text{uCDCA}} &= p[\text{gut_pu_freq_co}] \cdot x_{\text{co}}^{\text{uCDCA}} \\
 r_{\text{co pu}}^{\text{uSBA}} &= p[\text{gut_pu_freq_co}] \cdot x_{\text{co}}^{\text{uSBA}}
 \end{aligned}$$

B.5.7. REABSORPTION FLUXES

$$\begin{aligned}
 r_{\text{li from gut}}^{\text{cCA}} &= p[\text{hep_ratio_conj_tri}] \cdot r_{\text{il au}}^{\text{cCA}} \\
 &+ p[\text{hep_ratio_unconj_tri}] \\
 &\cdot (r_{\text{il au}}^{\text{uCA}} + r_{\text{il pu}}^{\text{uCA}} + r_{\text{co pu}}^{\text{uCA}}) \\
 r_{\text{li from gut}}^{\text{cCDCA}} &= p[\text{hep_ratio_conj_di}] \cdot r_{\text{il au}}^{\text{cCDCA}} \\
 &+ p[\text{hep_ratio_unconj_di}] \\
 &\cdot (r_{\text{il au}}^{\text{uCDCA}} + r_{\text{il pu}}^{\text{uCDCA}} + r_{\text{co pu}}^{\text{uCDCA}}) \\
 r_{\text{li from gut}}^{\text{cSBA}} &= p[\text{hep_ratio_conj_di}] \cdot r_{\text{il au}}^{\text{cSBA}} \\
 &+ p[\text{hep_ratio_unconj_di}] \\
 &\cdot (r_{\text{il au}}^{\text{uSBA}} + r_{\text{il pu}}^{\text{uSBA}} + r_{\text{co pu}}^{\text{uSBA}}) \\
 r_{\text{pl from gut}}^{\text{cCA}} &= (1 - p[\text{hep_ratio_conj_tri}]) \cdot r_{\text{il au}}^{\text{cCA}} \\
 r_{\text{pl from gut}}^{\text{cCDCA}} &= (1 - p[\text{hep_ratio_conj_di}]) \cdot r_{\text{il au}}^{\text{cCDCA}} \\
 r_{\text{pl from gut}}^{\text{cSBA}} &= (1 - p[\text{hep_ratio_conj_di}]) \cdot r_{\text{il au}}^{\text{cSBA}} \\
 r_{\text{pl from gut}}^{\text{uCA}} &= (1 - p[\text{hep_ratio_unconj_tri}]) \\
 &\cdot (r_{\text{il au}}^{\text{uCA}} + r_{\text{il pu}}^{\text{uCA}} + r_{\text{co pu}}^{\text{uCA}}) \\
 r_{\text{pl from gut}}^{\text{uCDCA}} &= (1 - p[\text{hep_ratio_unconj_di}]) \\
 &\cdot (r_{\text{il au}}^{\text{uCDCA}} + r_{\text{il pu}}^{\text{uCDCA}} + r_{\text{co pu}}^{\text{uCDCA}}) \\
 r_{\text{pl from gut}}^{\text{uSBA}} &= (1 - p[\text{hep_ratio_unconj_di}]) \\
 &\cdot (r_{\text{il au}}^{\text{uSBA}} + r_{\text{il pu}}^{\text{uSBA}} + r_{\text{co pu}}^{\text{uSBA}})
 \end{aligned}$$

B.5.8. EXCRETION FLUXES

$$\begin{aligned}
 r_{\text{co to fe}}^{\text{cCA}} &= p[\text{co_to_fe_freq}] \cdot p[\text{co_transit_coef}] \cdot x_{\text{co}}^{\text{cCA}} \\
 r_{\text{co to fe}}^{\text{cCDCA}} &= p[\text{co_to_fe_freq}] \cdot p[\text{co_transit_coef}] \cdot x_{\text{co}}^{\text{cCDCA}} \\
 r_{\text{co to fe}}^{\text{cSBA}} &= p[\text{co_to_fe_freq}] \cdot p[\text{co_transit_coef}] \cdot x_{\text{co}}^{\text{cSBA}} \\
 r_{\text{co to fe}}^{\text{uCA}} &= p[\text{co_to_fe_freq}] \cdot p[\text{co_transit_coef}] \cdot x_{\text{co}}^{\text{uCA}} \\
 r_{\text{co to fe}}^{\text{uCDCA}} &= p[\text{co_to_fe_freq}] \cdot p[\text{co_transit_coef}] \cdot x_{\text{co}}^{\text{uCDCA}} \\
 r_{\text{co to fe}}^{\text{uSBA}} &= p[\text{co_to_fe_freq}] \cdot p[\text{co_transit_coef}] \cdot x_{\text{co}}^{\text{uSBA}} \\
 r_{\text{fe from co}}^{\text{cCA}} &= r_{\text{co to fe}}^{\text{cCA}} \\
 r_{\text{fe from co}}^{\text{cCDCA}} &= r_{\text{co to fe}}^{\text{cCDCA}} \\
 r_{\text{fe from co}}^{\text{cSBA}} &= r_{\text{co to fe}}^{\text{cSBA}} \\
 r_{\text{fe from co}}^{\text{uCA}} &= r_{\text{co to fe}}^{\text{uCA}} \\
 r_{\text{fe from co}}^{\text{uCDCA}} &= r_{\text{co to fe}}^{\text{uCDCA}} \\
 r_{\text{fe from co}}^{\text{uSBA}} &= r_{\text{co to fe}}^{\text{uSBA}}
 \end{aligned}$$

B.5.9. SULFATION FLUXES

In human, a portion of secondary bile acids exists in sulfated form (Alnouti (2009)). However, according

to Sips et al. (2018), the process of desulfation happens at a very slow rate. This means only a small fraction of sulfated bile acids can re-enter the entero-hepatic circulation through desulfation, while the majority is excreted through feces. Hence, for simplicity, we account for sulfation by modeling its end outcome: the additional excretion of secondary bile acids into feces, which is again modeled with first-order dynamics.

$$\begin{aligned}
 r_{\text{co sulfate}}^{\text{cSBA}} &= p[\text{co_sulfate_freq}] \cdot p[\text{co_transit_coef}] \\
 &\cdot x_{\text{co}}^{\text{cSBA}} \\
 r_{\text{co sulfate}}^{\text{uSBA}} &= p[\text{co_sulfate_freq}] \cdot p[\text{co_transit_coef}] \\
 &\cdot x_{\text{co}}^{\text{uSBA}} \\
 r_{\text{fe sulfate}}^{\text{cSBA}} &= r_{\text{co sulfate}}^{\text{cSBA}} \\
 r_{\text{fe sulfate}}^{\text{uSBA}} &= r_{\text{co sulfate}}^{\text{uSBA}}
 \end{aligned}$$

Appendix C. Reinforcement Learning Specifications And Training Statistics

C.1. Rationale for training a separate model for each patient

Our current model essentially serves as a prototypical in-silico model for PSC, answering the question: what does a typical PSC patient’s trajectory look like? To this end, we selected a single patient whose measurements minimized the sum of the distances to all patients and found our model’s predictions to be consistent with existing literature. To ensure these findings were not coincidences from this one patient’s data, we repeated the training for four additional patients and observed similar trends.

We want to emphasize that our approach is rooted in a biological mechanistic model. Ultimately, this mathematical model approximates the real world. Parameterized with real data, the mathematical model serves as a starting point in RL that is subsequently optimized through the reward function. The primary objective of employing RL is to understand the change in enzyme activities in response to the disruptions in bile acid metabolism. This understanding is critical, as it provides valuable biological insights into this dynamic process, as well as lays the groundwork for developing an interpretable, individualized predictive model. Such a model could inform clinicians in making disease management decisions, for example, determining when a liver transplant may

become necessary for a given patient. Our current work is a very important step in this direction.

In future work, we plan to achieve a more individualized approach by incorporating patient-specific characteristics that are pertinent to disease progression. For example, body weight, which is linked to blood volume, could individualize the mechanistic model to generate patient-specific trajectories, which can then be compared to the real-world patient measurements in the reward function.

C.2. Additional Details of RL Setup

State: the state vector comprises bile acid levels $\mathbf{x}_{bile\ acids}$ and parameter values that represent host enzyme activity levels \mathbf{p}_{adapt} . Specifically, $\mathbf{x}_{bile\ acids}$ includes 30 variables across six bile acid species and six organs, denoted by $\{x_{OG}^{BA}, \forall BA \in \{cCA, cCDCA, cSBA\}, OG \in \{li, bd, il, co, pl, fe\}\}$ and $\{x_{OG}^{BA}, \forall BA \in \{uCA, uCDCA, uSBA\}, OG \in \{il, co, pl, fe\}\}$. Unconjugated bile acids in the liver and bile duct are not included as we assume bile acids are always conjugated in these two organs. \mathbf{p}_{adapt} includes $\{p[synthesis], p[syn_frac_CA], p[hep_ratio_conj_tri], p[hep_ratio_conj_di], p[max_asbt_rate]\}$, same as the parameters in the action space.

Reward: Here we introduce the additional reward terms used to ensure that the RL agent generates bile acid levels that are physiologically plausible while keeping the regulatable parameters close to their values in healthy conditions:

Enforcing physiologically plausible ranges. All bile acid levels are nonnegative, and they are upper bound by the maximum bile acid storage capacity in each organ. This reward term penalizes the RL agent for violating physiologically plausible bile acids ranges, pre-specified according to the literature. If the agent takes actions that lead to bile acid levels outside of any of the ranges, a large negative reward (VIOLATE_BOUNDARY_PENALTY) is applied, and the episode ends. We set VIOLATE_BOUNDARY_PENALTY to be -100.

Reward for not violating boundary conditions. To encourage the RL agent to take actions that lead to physiologically plausible states, we design this auxiliary reward term to give a small constant positive reward for every RL step that stays within the physiologically plausible ranges. We set the small positive reward to be 4.

Encouraging parameters resembling healthy conditions. This reward term encourages the RL agent to keep the regulatable parameters close to their values in healthy conditions. A negative reward is applied to penalize deviation of parameters from healthy conditions. We quantify the deviation as the smallest number of steps it takes to move from the healthy condition values to current values. For example, for $p[synthesis]$, a 25% higher or lower fold change is allowed at every step, then a value that is 156.25% of the original value takes at least two steps to reach, resulting in a deviation of 2. We sum up the deviation across all parameters and set the negative value of the sum as the reward. We multiply the negative deviation sum by a coefficient within $[0, 1]$ to match the range of other reward terms. Roughly each parameter is allowed a deviation penalty of at most 10 (and we clip the penalty at 10 if the deviation exceeds 10), and five adaptable parameters result in a total penalty of at most 50. We set the coefficient of this reward term to be 0.02 such that the penalty for parameter deviation is at most 1.

We also specify additional implementation details of the reward terms mentioned in the main text:

Minimizing toxicity. The maximum possible LE is calculated as the upper bound of physiologically plausible bile acid in the liver integrated over one day.

Resembling real-world patient data. Bile duct ligation experiments in mice showed that the adaptation usually stabilized after the first two weeks. Because ligation is an extreme model of bile duct obstruction, we allowed a four-week adaptation period in the RL simulation before introducing the fitting error to encourage RL trajectories that are similar to stabilized patient data. Because bile acid measurements are roughly log-normally distributed, log transform is applied to both the actual measurements and the RL-generated bile acid states. We set the cap for fitting error, CAP, to be a generous value of 20. We set the coefficient λ_{error} to be 0.2 such that the fitting error is bounded at 4, giving more emphasis on the data-guided fitting error than the other biology-inspired reward terms that are at most 1.

C.3. RL algorithm training

The actor and critic in PPO shared a neural network with the architecture of a three-layer multilayer perceptron with 100, 50, and 25 hidden units in each

layer. We set the learning rate to be 0.002. We normalized the states during the RL training for easier optimization of the algorithm. We used 16 vectorized environments on different processes to speed up training. We leave the other training parameters to be the default of the implementation, including a discount factor of 0.99 to mimic the optimization of long-term homeostasis. We trained the RL agent for 4,000,000 environment steps to ensure convergence of the algorithm.

C.4. RL training statistics

Figure 10 shows the RL training statistics. We see that PPO is able to converge across different values of $p[\text{bd_max_flow}]$, with smaller $p[\text{bd_max_flow}]$, i.e., more severe obstruction of the bile ducts, resulting in less reward. This is expected because with more severe diseases, balancing between minimizing liver toxicity and maintaining cholesterol elimination and digestion becomes more difficult, and restoring homeostasis requires a higher level of adaptation that will result in larger deviation from healthy conditions. However, after training completes, as the simulation in Figure 11 shows, after the initial adaptation period right after the introduction of PSC pathophysiology, simulation corresponding to $p[\text{bd_max_flow}]=3 \mu\text{mol}/\text{min}$ generates fasting plasma bile acid profiles closest to a representative PSC patient, shown by the smallest fitting error indicated by the green curve. This scenario represents a relatively severe bile duct obstruction. We chose to focus on this scenario for the later analysis under the assumption that a small fitting error from actual data implies that this is more likely to reflect what actually happens in PSC patients.

Appendix D. Complete Bile Acid Dynamics After Introducing PSC Pathophysiology

For all simulations in this section, PSC pathophysiology was simulated by setting $p[\text{bd_max_flow}]=3 \mu\text{mol}/\text{min}$ at day 0, much lower than the maximum bile duct to the small intestine flux of $75 \mu\text{mol}/\text{min}$ observed in the healthy individual simulation. We initialized the bile acid values with their steady-state values under healthy conditions. This is to mimic one animal model of PSC — where bile duct ligation surgeries are performed on healthy mice and PSC-like

symptoms will develop in the mice. Other ODE parameters were also initialized using their values under healthy conditions. Only fasting values (8 AM) were plotted in the figures.

D.1. Bile acid dynamics without RL

To see the simulated PSC bile acid dynamics without RL, we ran the simulation for 60 days after introducing PSC pathophysiology. We chose a simulation period of 60 days to be longer than the two-week adaptation period in the mouse experiment (after which the mouse liver adaptation stabilizes) and to achieve stable bile acid states (all organs other than the liver) in the simulation. Simulated meals at 8 AM, 2 PM, and 8 PM are included as usual. Figure 12 shows the complete simulated 60-day bile acid dynamics without RL adaptation. We observe unrealistically high liver bile acid levels and decreased plasma conjugated bile acids that contradict clinical data. Changing $p[\text{bd_max_flow}]$ to $10 \mu\text{mol}/\text{min}$ results in similar trends. Moreover, continuing the simulation would eventually run into numerical issues due to the ever-increasing liver bile acid levels.

These discrepancies arise from the flawed assumption that bile acid metabolism parameters are constants. While this assumption might be valid under healthy conditions, where a homeostatic state is maintained, it fails to account for the dynamic nature of PSC. Bile duct ligation experiments in mice have shown bile duct obstruction rapidly disrupts bile acid homeostasis, as reflected by the immediate shift in bile acid profiles upon introducing PSC pathophysiology into our model. When such disruption of homeostasis occurs, it is crucial to recognize that biological systems are inherently dynamic and have the ability to self-adapt in order to restore homeostasis.

For instance, in cholestatic conditions, bile acids accumulate in the liver. The liver responds by suppressing the synthesis of bile acids to mitigate potential toxicity associated with excessive bile acid exposure. However, it is important to note that complete suppression of bile acid synthesis is not a feasible option. Bile acid synthesis is the sole source of incoming bile acids in the enterohepatic circulation. If synthesis were completely halted, it would eventually deplete bile acids in all compartments, thereby leaving no bile acids for the digestion of fat. Therefore, modeling bile acid metabolism under PSC pathophysiology requires accounting for the dynamic nature of the system, its self-adjusting capabilities, and the various




Figure 10: PPO training steps vs. mean reward evaluated on 5 episodes, each curve represents a different level of disease severity reflected by $p[\text{bd.max_flow}]$.




Figure 11: Fitting error between simulated fasting plasma bile acids (with adaptation) and fasting plasma bile acid measurement from a representative PSC patient. Bile acid measurement from the actual patient is cross-sectional, while the fitting error is calculated daily, over a 240-day simulation after introducing PSC pathophysiology. Each curve represents a different level of disease severity reflected by $p[\text{bd.max_flow}]$. The adaptive ODE parameters are obtained from PPO agents trained with the corresponding $p[\text{bd.max_flow}]$ for 4M steps.

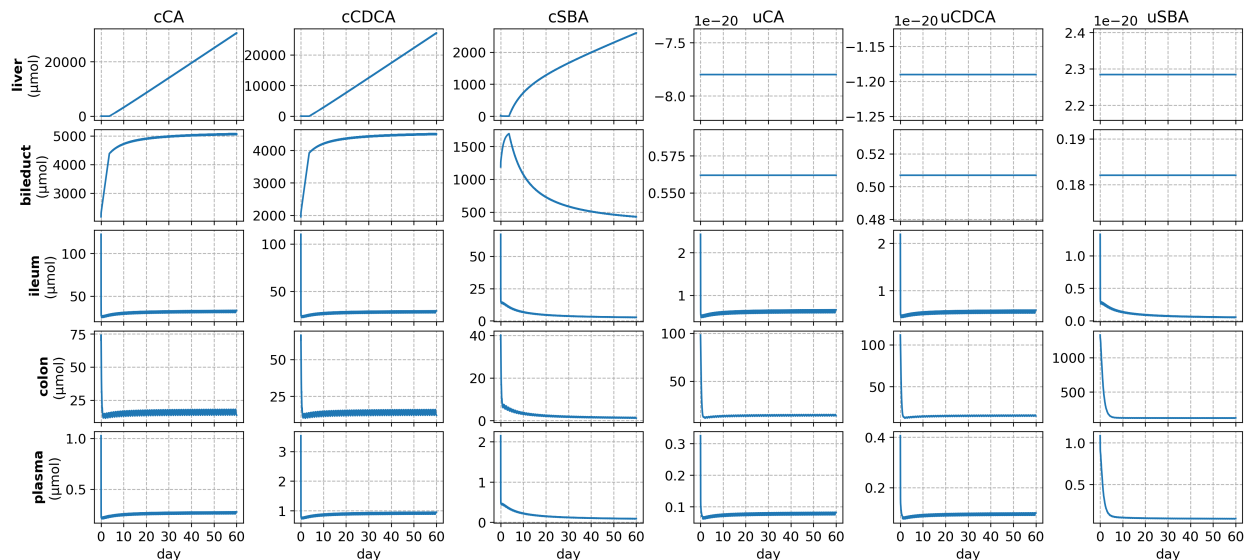


Figure 12: Simulated 60-day bile acid dynamics after introducing PSC pathophysiology, without RL adaptations.

mechanisms aimed at preserving homeostasis. We achieve this in our model by allowing the values of the biomolecular parameters to undergo continuous changes that ensure the preservation of homeostasis.

D.2. Bile acid dynamics with RL (Patient 1)

Figure 13 shows the complete 240-day PSC bile acid dynamics with adaptive ODE parameters obtained from the trained PPO agent. As indicated by the red and blue dashed lines in the trajectories of plasma bile acids, the RL adaptation led to the elevation of plasma conjugated bile acids values and the decrease of plasma unconjugated bile acids. The resulting plasma bile acids profile was closer to the actual PSC patient measurements, correcting the artifacts in the simulation without RL where plasma bile acids decreased across all species.

D.3. Bile Acid Dynamics with RL for ten random seeds (Patient 1)

To understand how the stochasticity in the RL agent’s actions affects the bile acid dynamics. Figure 14 showed the 240-day PSC bile acid dynamics with adaptive ODE parameters obtained from the trained PPO agent for ten random seeds. Trajectories from all ten seeds shared a similar adaptation period

in the beginning, where bile duct bile acids increased, ileum and colon bile acids decreased, and liver bile acids experienced a sharp increase and then decreased following the saturation of bile duct bile acids. After the initial adaptation period, trajectories from different seeds exhibited larger but bounded variability from the stochasticity of the RL actions. Interestingly, all trajectories contained sporadic peaks of liver bile acids, potentially explaining the episodic symptom flare-ups in PSC patients.

D.4. *In silico* evaluation of bile acid therapies (Patient 1)

ASBT inhibitors such as maralixibat have been in clinical trials for PSC. As mentioned before, ASBT mediates the active uptake of bile acid from the ileum. We used our trained RL agent as an in-silico testbed to test ASBT inhibitors for PSC by decreasing the corresponding parameter $p[\text{max_asbt_rate}]$. Figure 15 showed the resulting simulated trajectories. Compared to the case of no inhibition, the complete reduction of active uptake prevented peaks in liver bile acids after the initial adaptation period. However, significant increases in intestine and plasma bile acid levels led to premature termination of the simulation,

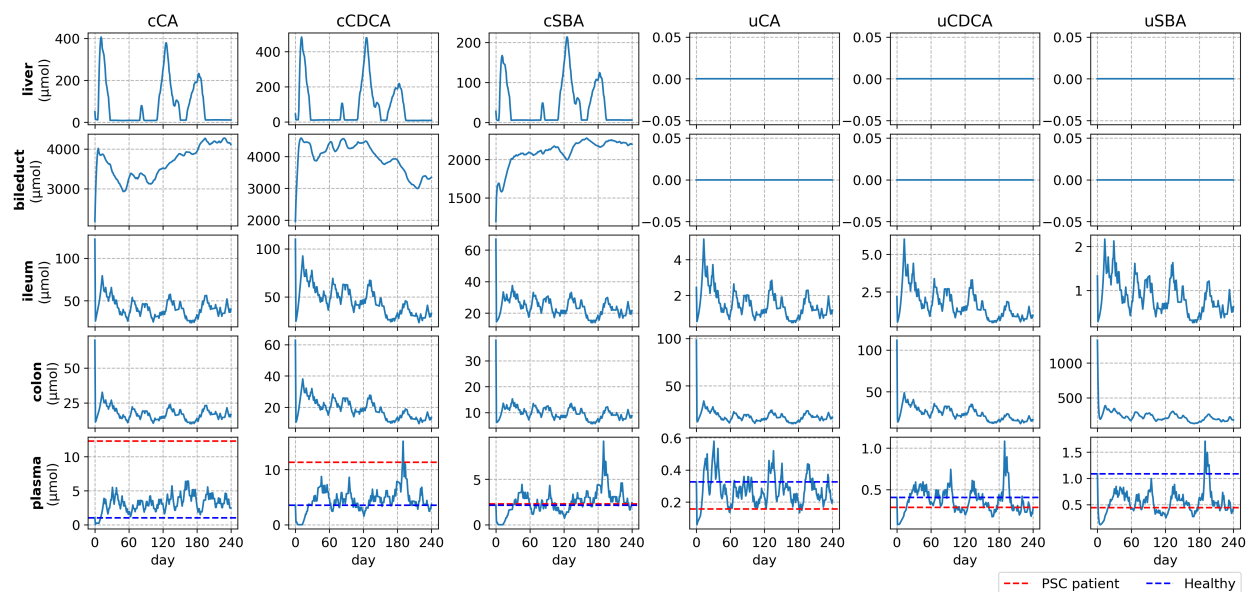


Figure 13: Simulated 240-day bile acid dynamics trajectory after introducing PSC pathophysiology with adaptive ODE parameters obtained from a trained PPO agent.

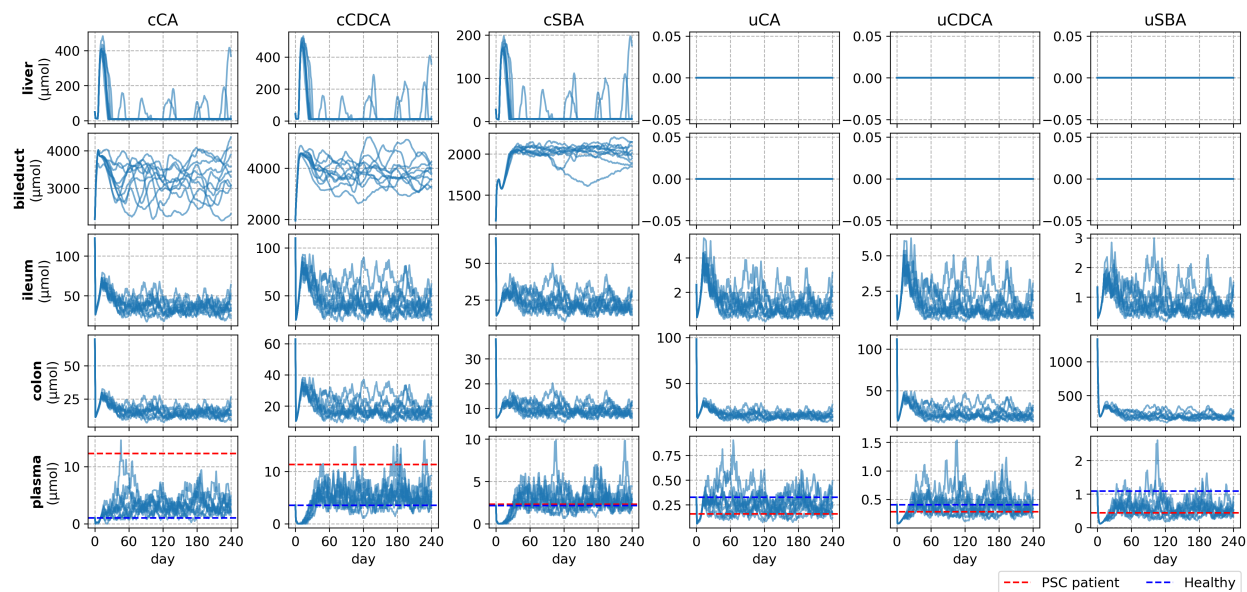


Figure 14: Simulated 240-day bile acid dynamics trajectories after introducing PSC pathophysiology with adaptive ODE parameters obtained from a trained PPO agent and ten random seeds.

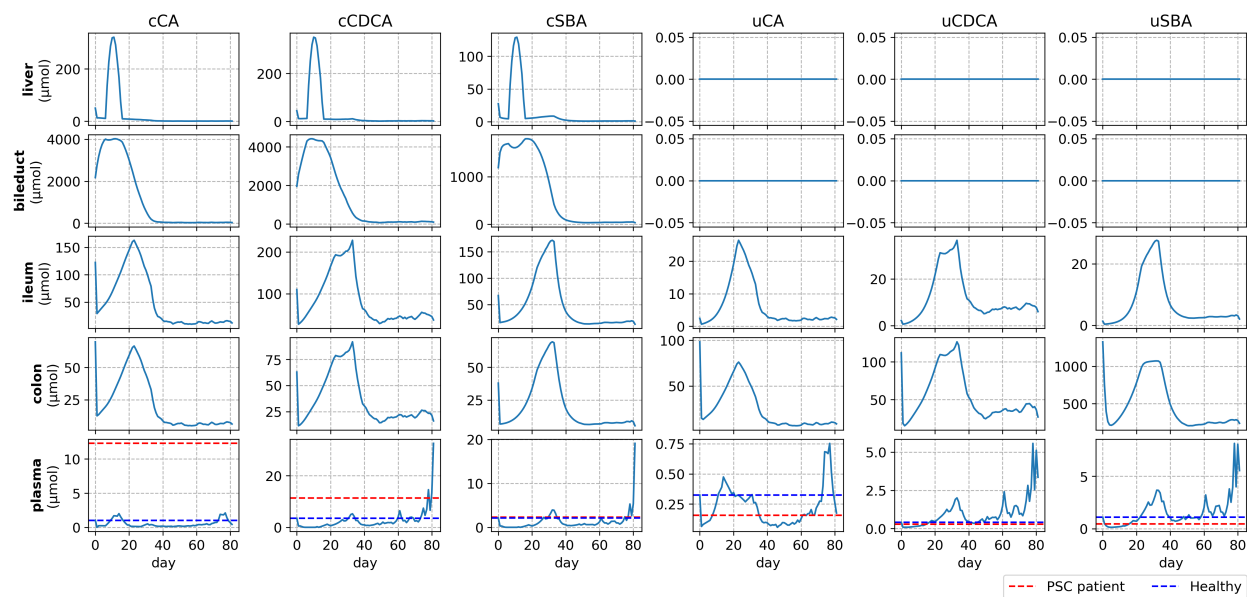


Figure 15: Simulated 240-day bile acid dynamics after introducing PSC pathology with adaptive ODE parameters obtained from a trained PPO agent. Complete reduction of active uptake was simulated by forcing $p[\text{max_asbt_rate}]$ to be down-regulated unless it reaches the lower bound of its predefined physiological ranges.

rendering complete reduction a perhaps unrealistic strategy.

We further tested a more moderate strategy — 50% reduction of active uptake with our trained RL agent. Figure 16 showed a reduction of liver bile acid peaks compared to the case of no intervention, suggesting that partial suppression of active uptake might help alleviate the toxicity burden on the liver. Furthermore, unlike complete reduction of active uptake, all bile acid species were able to stay within plausible ranges across the organs, making partial reduction a more realistic strategy.

Appendix E. Bile Acid Dynamics After Introducing PSC Pathophysiology (Patient 2–4)

To avoid potential bias from patient selection, we replicated our experiments with four other representative patients who had the second to the fourth smallest sum of distances to all patients. We showed their bile acid dynamics as well as parameter adaptation trajectories below. For consistency with results in the main text, we chose the case of $p[\text{bd_max_flow}]=3 \mu\text{mol}/\text{min}$ to train REMEDI across all cases.

In Table 2, we also showed the fitting error between REMEDI prediction trained with all patients and the plasma measurements from the representative PSC patient chosen in the main text.

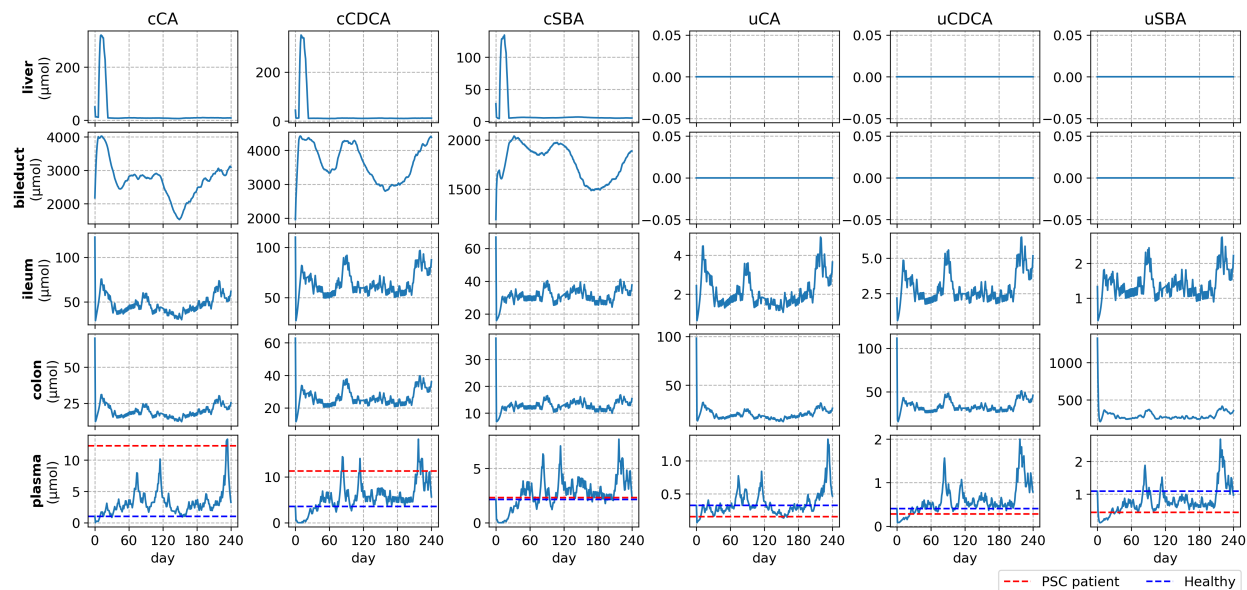


Figure 16: Simulated 240-day bile acid dynamics after introducing PSC pathology with adaptive ODE parameters obtained from a trained PPO agent and ten pre-specified random seeds. 50% reduction of active uptake was simulated by forcing $p[\text{max_asbt_rate}]$ to be down-regulated unless it reaches or is below half of the lower bound of its predefined physiological ranges.

Table 2: Fitting Error (μmol , day 50 – 60 averaged): REMEDI – Data from PSC **Patient 1**.

Scenario/Plasma	cCA	cCDCA	cSBA	uCA	uCDCA	uSBA
REMEDI (without RL)	-12.01	-10.37	-2.21	-0.08	-0.18	-0.35
REMEDI (trained with patient 1)	-9.79	-5.27	0.67	0.05	0.16	0.15
REMEDI (trained with patient 2)	-10.67	-4.75	0.54	0.02	0.33	0.26
REMEDI (trained with patient 3)	-9.56	-7.32	0.19	0.15	0.10	0.22
REMEDI (trained with patient 4)	-10.34	-5.14	1.17	0.01	0.09	0.14
REMEDI (trained with patient 5)	-9.06	-9.51	-0.88	0.14	-0.12	-0.07

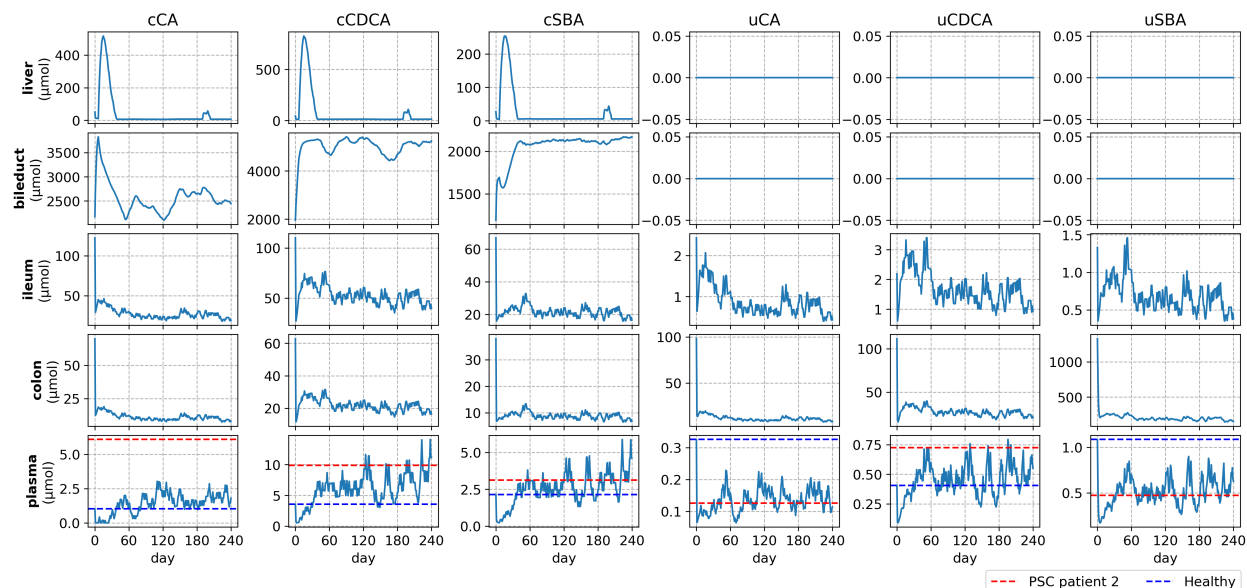


Figure 17: **Patient 2**: Simulated 240-day bile acid dynamics trajectory after introducing PSC pathophysiology with adaptive ODE parameters obtained from a trained PPO agent.




Figure 18: Adaptation of ODE parameters corresponding to the simulation scenario illustrated in Figure 17. *synthesis* (blue): bile acid synthesis rate; *syn_frac_CA* (orange): fraction of CA (vs. CDCA) in newly synthesized bile acid; *hep_ratio_conj_tri* (green): fraction of bile acid uptake entering the liver (vs. plasma) for cCA; *hep_ratio_conj_di* (red): fraction of bile acid uptake entering the liver for cCDCA and cSBA; *max_asbt_rate* (purple): efficiency of bile acid active uptake.

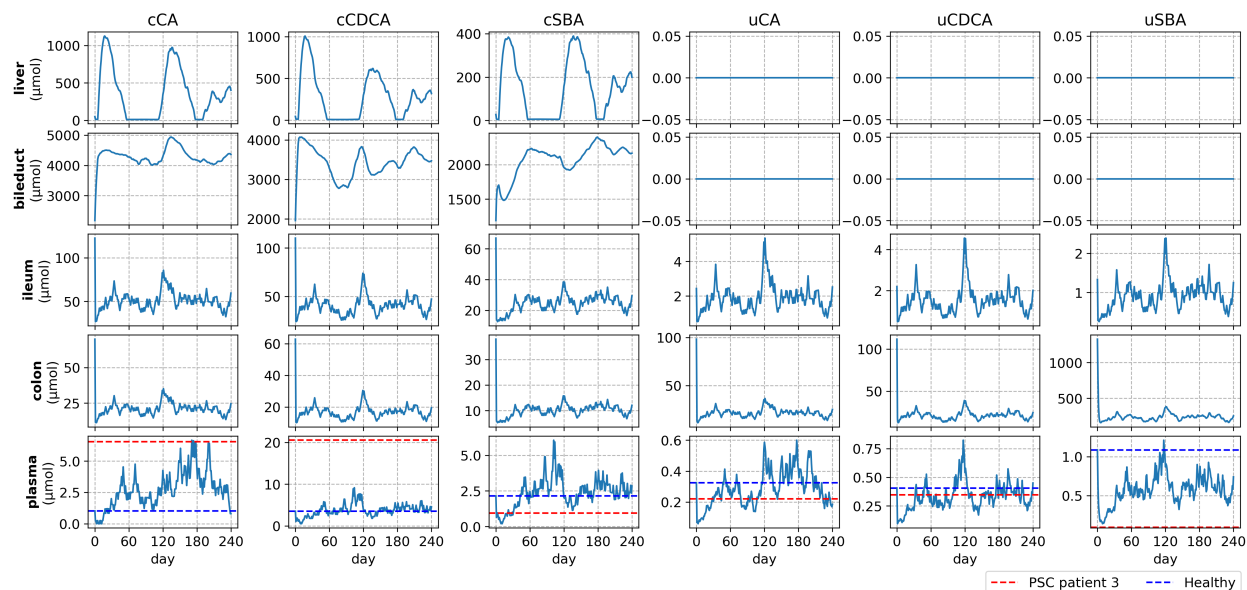


Figure 19: **Patient 3**: Simulated 240-day bile acid dynamics trajectory after introducing PSC pathophysiology with adaptive ODE parameters obtained from a trained PPO agent.




Figure 20: Adaptation of ODE parameters corresponding to the simulation scenario illustrated in Figure 19. *synthesis* (blue): bile acid synthesis rate; *syn_frac_CA* (orange): fraction of CA (vs. CDCA) in newly synthesized bile acid; *hep_ratio_conj_tri* (green): fraction of bile acid uptake entering the liver (vs. plasma) for cCA; *hep_ratio_conj_di* (red): fraction of bile acid uptake entering the liver for cCDCA and cSBA; *max_asbt_rate* (purple): efficiency of bile acid active uptake.

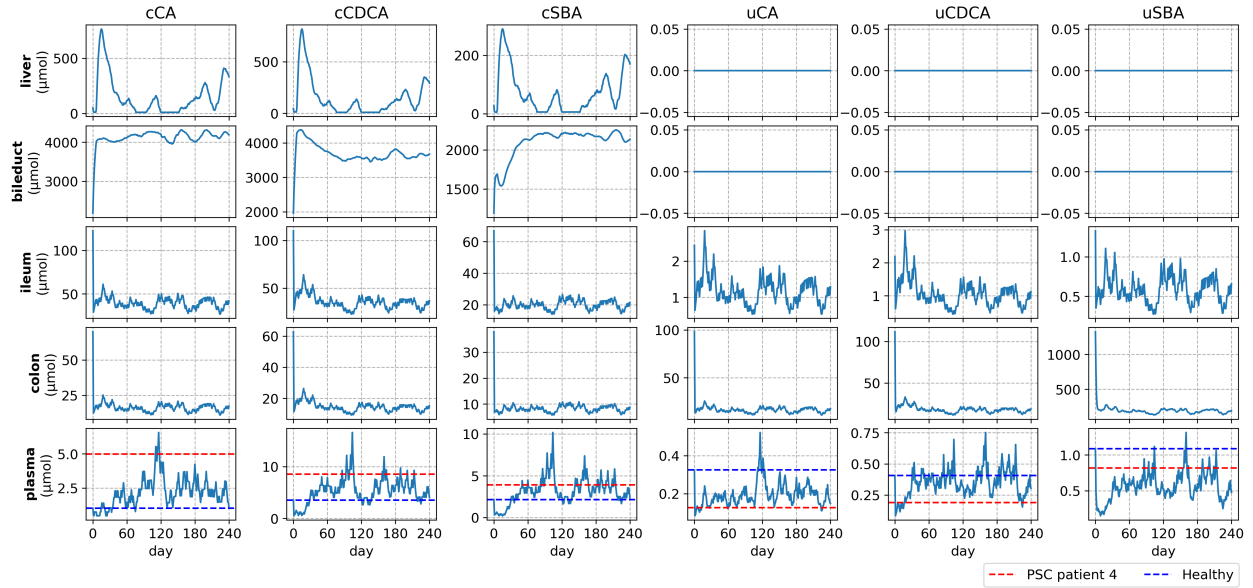


Figure 21: **Patient 4**: Simulated 240-day bile acid dynamics trajectory after introducing PSC pathophysiology with adaptive ODE parameters obtained from a trained PPO agent.




Figure 22: Adaptation of ODE parameters corresponding to the simulation scenario illustrated in Figure 21. *synthesis* (blue): bile acid synthesis rate; *syn_frac_CA* (orange): fraction of CA (vs. CDCA) in newly synthesized bile acid; *hep_ratio_conj_tri* (green): fraction of bile acid uptake entering the liver (vs. plasma) for cCA; *hep_ratio_conj_di* (red): fraction of bile acid uptake entering the liver for cCDCA and cSBA; *max_asbt_rate* (purple): efficiency of bile acid active uptake.

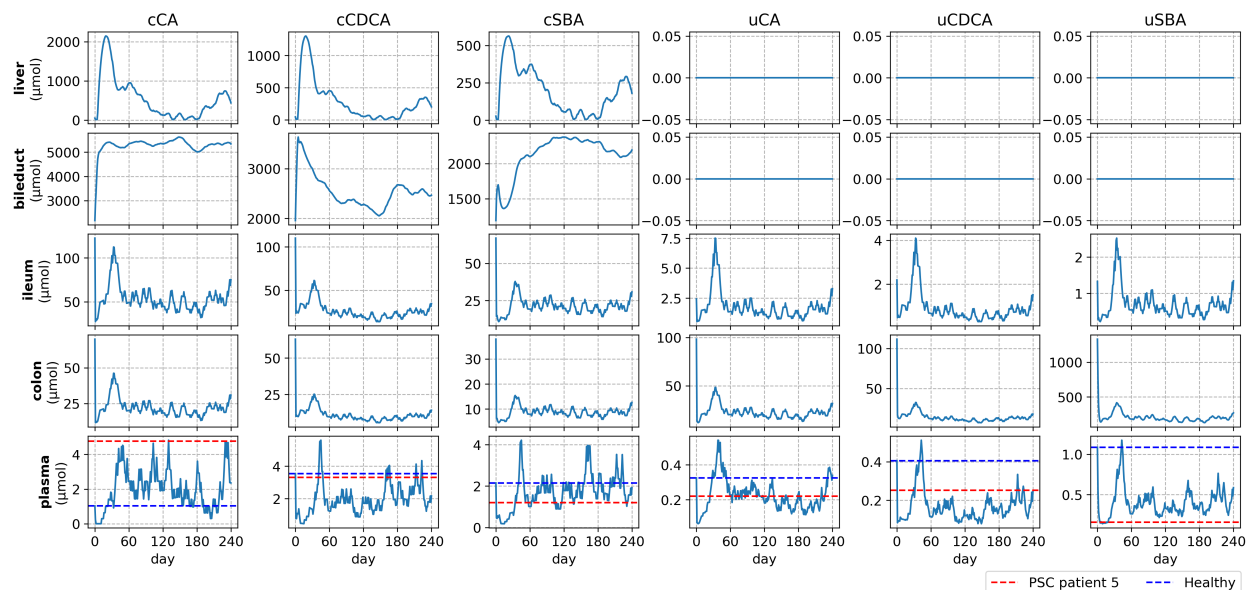


Figure 23: **Patient 5**: Simulated 240-day bile acid dynamics trajectory after introducing PSC pathophysiology with adaptive ODE parameters obtained from a trained PPO agent.




Figure 24: Adaptation of ODE parameters corresponding to the simulation scenario illustrated in Figure 23. *synthesis* (blue): bile acid synthesis rate; *syn_frac_CA* (orange): fraction of CA (vs. CDCA) in newly synthesized bile acid; *hep_ratio_conj_tri* (green): fraction of bile acid uptake entering the liver (vs. plasma) for cCA; *hep_ratio_conj_di* (red): fraction of bile acid uptake entering the liver for cCDCA and cSBA; *max_asbt_rate* (purple): efficiency of bile acid active uptake.