Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Leveraging mechanistic modelling for the identification and control of CAR-T cell therapy

Adriana González Vázquez
CITMAga
Alejandro Fernández Villaverde
afvillaverde@uvigo.gal
CITMAga

Article

Keywords:

Posted Date: November 25th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11729000/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Leveraging mechanistic modelling for the identification and control of CAR-T cell therapy

Adriana González Vázquez[1,2] and Alejandro F. Villaverde[1,2,*]

[1]CITMAga, 15782 Santiago de Compostela, Galicia, Spain
[2]Universidade de Vigo, Department of Systems Engineering and Control, 36310 Vigo, Galicia, Spain
[*]Corresponding author: Alejandro F. Villaverde, afvillaverde@uvigo.gal

November 16, 2026

Abstract

CAR-T cell therapy has revolutionized the treatment of certain blood cancers, and mechanistic models are increasingly employed to study it and support treatment optimisation. The predictive capability of a model depends on its proper identification, that is, the possibility of inferring its states and parameters, both theoretically and in practice. In parallel, the ability to influence a system's dynamics is determined by properties such as accessibility and controllability. Sensitivity analysis complements these assessments by quantifying how changes in model parameters affect relevant outputs. In this study, we consider two representative CAR-T therapy models: one describing tumour growth under CAR-T therapy (HCART), and another capturing the immune response leading to Cytokine Release Syndrome (CRS). We use these models to investigate the inverse problem of parameter and state estimation, and to determine whether they can describe the natural evolution of the underlying processes and our ability to drive it. To this end, we analyse their structural and practical identifiability, observability, accessibility, controllability, and parameter sensitivities. We use a differential geometry approach to study structural properties, joint Monte Carlo parameter estimation for practical identifiability, and a custom complex-step derivative approximation for computing sensitivities. For the controllability and accessibility analyses, we reformulate the models to represent CAR-T administration as an external input. The results highlight a distinction between structural identifiability under the adopted parameter restrictions and practical estimation performance: CRS exhibited substantial parameter uncertainty, whereas HCART showed comparatively more consistent estimation behaviour. In regard to controllability, HCART is locally controllable at the tumour-bearing equilibrium, but not at the tumour-free equilibrium; structural computations of the CRS were inconclusive, while its numerical analyses indicated accessibility limitations. The input-independent dynamics of the total macrophage population provided a mechanistic explanation for part of these limitations. Lastly, sensitivity analyses revealed which parameters have a significant impact on the models' behaviour. Together, these results provide a framework for assessing model capabilities and limitations and informing parameter estimation, experimental design, and future investigations of treatment control.

1 Introduction

Tumour growth models based on ordinary differential equations (ODEs) have become essential tools for understanding the behaviour of cancer progression and evaluating treatment strategies. Mechanistic models are based on first principles (mass balance, reaction kinetics, immune-cell dynamics) and are capable of linking measurable biomarkers to underlying biological mechanisms [1, 2]. Mechanistic models thus allow not only to reproduce tumour growth dynamics, but also to predict treatment outcomes and optimize dosing regimens [3]. They can incorporate therapeutic interventions such as radiotherapy, chemotherapy, and immunotherapy, and they provide a mechanistic representation of biological processes, allowing the systematic analysis of tumour–therapy interactions under controlled assumptions [4]. Consequently, it is essential to ensure their adequacy and reliability.

Immunotherapy has experienced significant advances, with Chimeric Antigen Receptor T-cells (CAR-T) being one of the most promising strategies, which involves genetically modifying a patient's T-cell to express a receptor that specifically targets cancer cells [5, 6]. By enhancing the patient's immune

response against malignant cells, CAR-T therapies have achieved remarkable clinical success in treating haematologic malignancies such as leukaemia [7] and lymphoma, while their application to solid tumours remains more challenging. Despite these clinical successes, CAR-T therapy remains complex due to its nonlinear dynamics, patient-specific variability, and the strong coupling between immune activation and tumour burden [8]. Challenges posed by CAR-T cell therapy include variable patient responses, relapse, and severe adverse effects such as Cytokine Release Syndrome (CRS). These challenges have inspired extensive research focused on understanding the mechanisms underlying CAR-T cell therapy and optimizing its application [9]. CRS, in particular, has been identified as a major dose-limiting toxicity, driven by excessive cytokine production, with interleukin-6 (IL_6) playing a central role. Experimental evidence suggests that macrophages, rather than CAR-T cells themselves, are the main contributors to IL_6 secretion, highlighting the importance of including innate immune components in mathematical models of CAR-T therapy [10].

To address these challenges, computational models have been developed to simulate diverse aspects of the treatment, including T-cell activation, tumour cell interactions and cytokine dynamics, among others. These models range from single predator-prey frameworks, which describe the interplay between CAR-T cells and tumour cells, to more complex systems that incorporate intracellular signalling pathways and patient-specific factors. For example, Liu et al. [7] developed a computational model to predict patient responses to CAR-T cell therapy by analysing the dynamics of CAR-T cells and tumour burden. Their model revealed key determinants of therapeutic outcomes, such as the functionality of CAR-T cells and the role of bystander killing in CD19-negative relapse. Nukala et al. [8] conducted a systematic review of various modelling approaches, highlighting the use of ordinary differential equations (ODEs) to study CAR-T cell activation, expansion, and exhaustion. These models have provided valuable insights into the factors influencing CAR-T cell efficacy, including the impact of tumour burden, lymphodepletion, and CAR-T cell composition. Nevertheless, many of these studies implicitly assume that model parameters are known or can be reliably estimated, an assumption that may not hold in practice. A good agreement with calibration data does not guarantee that the parameter estimates are unique or that model predictions remain reliable outside calibration conditions [11, 12].

A central question in model development is whether the structure of the equations allows reliable inference of parameters and state variables from experimental or clinical data. The two relevant properties are *structural identifiability and observability* (SIO), which determine whether model parameters and internal states, respectively, can be inferred from ideal, noise-free measurements [13]. In their absence, different parameter combinations may produce indistinguishable outputs, compromising model calibration and predictive accuracy. In realistic settings, where data are finite and affected by experimental noise, *practical* identifiability, which determines whether parameters can be estimated with acceptable confidence intervals, becomes relevant [14]. A different set of properties, *accessibility and controllability*, characterize the ability of external inputs to influence the system dynamics. Accessibility determines whether the system can evolve towards a full-dimensional neighbourhood of a given state under external inputs, whereas controllability addresses whether suitable inputs can drive the system towards a target state, e.g. when aiming to control tumour growth or to reach specific treatment goals. Lastly, *sensitivity* analysis quantifies how changes in model parameters affect variables such as tumour volume or cytokine levels, providing complementary information on the parameters that most strongly influence the system dynamics. Sensitivities may also be used to prioritize the most influential factors for model calibration and guide experimental design, indicating which parameters require more accurate measurement and which can be fixed without significantly affecting model predictions. Although the aforementioned properties are often studied independently, together they provide a comprehensive mathematical characterization of a model, revealing not only whether it can reproduce data, but also whether it can be reliably calibrated, interpreted, and ultimately used for predictive or control-oriented applications.

To study the aforementioned properties, in this work we combine our previously presented implementations of techniques, included in the MATLAB toolbox STRIKE-GOLDD [15, 16], with AMIGO2 for practical identifiability analysis [17] and with a new implementation of a numerical method for sensitivity analysis [18]. Additionally, we use StrucID for complementing the accessibility analyses in computationally demanding cases [19]. We use these methodologies to characterize two representative CAR-T models systematically through structural identifiability, observability, accessibility, controllability, sensitivity, and practical identifiability analyses. By integrating these complementary mathematical properties into a unified framework, this work provides a comprehensive assessment of the reliability of CAR-T mechanistic models and identifies their strengths and limitations for parameter estimation, prediction, and future model-based therapeutic optimization.

2 Models

2.1 Mathematical framework

The CAR-T cell therapy models considered in this work are systems of nonlinear ordinary differential equations (ODE) describing the interactions between tumour cells, CAR-T cell subpopulations, immune cells and cytokines. In their general form, these models can be written as

$$M : \begin{cases} \dot{x}(t) = f\left(x(t), u(t), \theta\right), \\ y(t) = h\left(x(t), u(t), \theta\right), \end{cases} \tag{1}$$

where $x(t) \in \mathbb{R}^{n_x}$ denotes the vector of state variables, including the biological populations and molecular species involved in the therapy; $u(t) \in \mathbb{R}^{n_u}$ are external inputs that may represent therapeutic interventions, such as CAR-T cell infusion or treatment administration; $y(t) \in \mathbb{R}^{n_y}$ is the vector of experimentally measurable outputs; and $\theta \in \mathbb{R}^{n_\theta}$ contains the model parameters involved in proliferation, death, activation, differentiation, and interaction processes. The vector of nonlinear functions $f(\cdot)$ describes the biological dynamics, whereas $h(\cdot)$ specifies the measured quantities available for model calibration. Throughout the paper, the explicit dependence on time will be omitted whenever it does not compromise clarity.

2.2 CAR-T immunotherapy model of haematological cancers (HCART)

This model, originally presented by Barros et al. [20], describes the interaction between tumour cells and CAR-T cells (effector and memory populations) in immunodeficient mouse models. The model incorporates tumour growth, CAR-T cell dynamics, and the formation of immunological memory. It uses three ODEs to explore the factors influencing therapy outcomes, including cytotoxicity, tumour-induced immunosuppression, and long-term memory formation. The first equation describes how effector CAR-T cells (C_T) proliferate at a certain rate (ϕ) in response to tumour burden, die or differentiate into memory cells (ρ), are activated by interaction with tumour cells (θ), and are suppressed by tumour-induced immunosuppressive mechanisms (α):

$$\frac{dC_T}{dt} = \phi C_T - \rho C_T + \theta T C_M - \alpha T C_T \tag{2}$$

The second equation describes the evolution of memory CAR-T cells (C_M), which are generated from effector CAR-T cells (ϵ), activated back into effector cells when interacting with tumour cells (θ), and decay naturally (μ):

$$\frac{dC_M}{dt} = \epsilon C_T - \theta T C_M - \mu C_M \tag{3}$$

Lastly, tumour cells (T) grow logistically with a maximum growth rate r and carrying capacity $\frac{1}{b}$. They are eliminated by effector CAR-T cells through a cytotoxic effect (γ):

$$\frac{dT}{dt} = rT(1 - bT) - \gamma C_T T \tag{4}$$

2.3 CAR-T cell therapy and macrophage-mediated Cytokine Release Syndrome

Santurio et al. have developed a multi-layer mathematical model that describes the dynamics of CAR-T cell therapy and the associated cytokine release syndrome (CRS), consisting of four compartments [21]. The first compartment represents the CAR-T cells that have been administered to the patient (injected CAR-T cells, $C_I(t)$):

$$\frac{dC_I}{dt} = -\eta F(T_P) C_I - \mu_I C_I, \tag{5}$$

The second compartment includes the CAR-T cells that have been activated and are in the process of proliferating (expanding CAR-T cells, $C_E(t)$). These cells multiply in response to the presence of tumour cells that express the specific antigen they recognize. The equation is:

$$\frac{dC_E}{dt} = \nu F(T_P) C_I + \kappa F(T_P) C_E - \epsilon(1 - F(T_P)) C_E + \theta F(T_P) C_P - \mu_E C_E, \tag{6}$$

The third compartment represents CAR-T cells that survived the expansion phase and differentiated into memory cells (Persistent CAR-T cells, $C_P(t)$), which are important for the long-term immune response, as they can be reactivated if needed:

$$\frac{dC_P}{dt} = \epsilon(1 - F(T_P))C_E - \theta F(T_P)C_P - \mu_P C_P \tag{7}$$

where the function $F(T_P)$ represents the antigen-receptor binding interaction,

$$F(T_P) = \frac{T_P}{A + T_P}. \tag{8}$$

The fourth and last compartment represents the tumour cells that express the specific antigen that the CAR-T cells are designed to recognize and attack (antigen-positive tumour cells, $T_P(t)$), which is important to understand how CAR-T cells affects the tumour burden:

$$\frac{dT_P}{dt} = \rho T_P(1 - \frac{T_P}{K}) - \gamma \frac{C_E}{B + C_E} T_P. \tag{9}$$

Model parameters include the half-saturation constant A; the death rates of injected, expanding, and persistent CAR-T cells, respectively μ_I, μ_E, μ_F; the tumour growth rate ρ; and the tumour cell carrying capacity K.

The model is extended by another layer that modifies the behaviour of the antigen-positive (T_P) tumour cells, relating them with antigen-negative (T_N) cells through the following two equations:

$$\frac{dT_P}{dt} = \rho T_P(1 - \frac{T_P + T_N}{K}) - \gamma \frac{C_E}{B + C_E} T_P \tag{10}$$

$$\frac{dT_N}{dt} = \rho T_N(1 - \frac{T_P + T_N}{K}) - g_0\gamma \frac{C_E}{B + C_E} T_N \tag{11}$$

where g_0 is a fraction that represents a reduced cytotoxicity due to the absence of antigen.

Additionally, the model introduces another layer to describe cytokine dynamics and the cytokine release syndrome (CRS). It assumes that cytokine release is mediated by macrophage activation, driven by three main mechanisms: the release of inflammatory mediators by activated CAR-T cells, the release of DAMPs (Damage-Associated Molecular Patterns) from tumour cell death, and contact-dependent activation between macrophages and CAR-T cells via CD40/CD40-ligand interaction. The equations are:

$$\frac{dM_i}{dt} = \sigma_M - h(C_E, T_P, M_a)M_i - \delta_M M_i \tag{12}$$

$$\frac{dM_a}{dt} = h(C_E, T_P, M_a)M_i - \delta_M M_a \tag{13}$$

$$\frac{dIL_6}{dt} = \sigma_I + \alpha M_a - \delta_I IL_6 \tag{14}$$

where M_i and M_a are the number of naive (monocytes) and activated macrophages at time t, IL_6 Is the IL_6 concentration, σ_M is the natural production rate of naive macrophages, δ_M is the death rate of macrophages, σ_I represents the endogenous production rate of IL_6, δ_I is the natural decay rate of IL_6, and α is the rate of IL_6 release by activated macrophages, whose activation rate is given by:

$$h(C_E, T_P, T_n, M_a) = \beta_B \frac{T_P}{A + T_P} C_E + \beta_K \frac{C_E}{B + C_E}(T_P + g_0 T_N) + \beta_C \frac{M_a}{C + M_a} C_E. \tag{15}$$

Figure 1 shows the structure of both models. Their parameters, along with their biological meaning, nominal values, and the intervals used in the calibration are listed in Table 1.

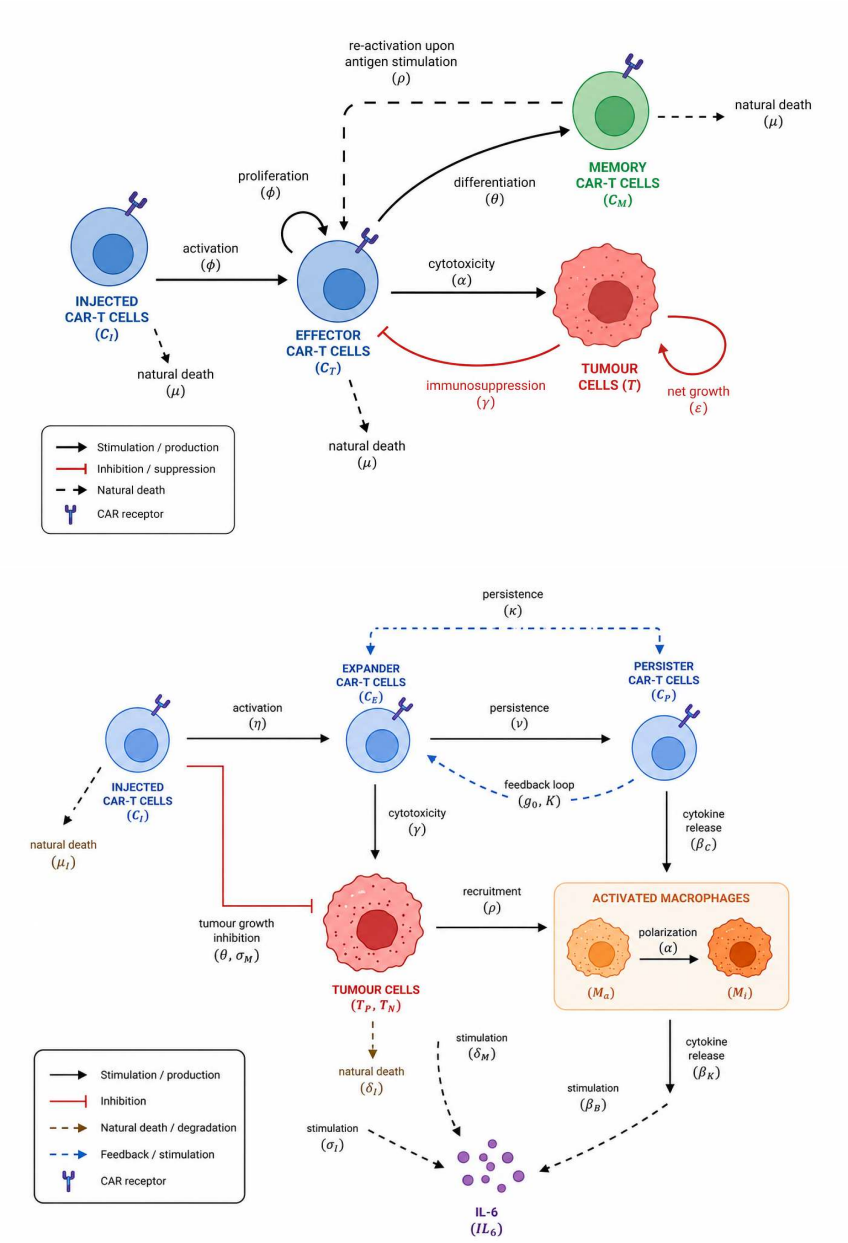

Figure 1: Top: HCART Model. Bottom: CRS Model.

Table 1: Model parameters, biological meanings, nominal values, and physical intervals used in the practical-identifiability analysis. The nominal HCART values correspond to the HDLM-2 + CAR-T 123 scenario. "Fixed" indicates that the parameter was retained at its nominal value and was not estimated. Parameter units follow the original model formulations [20, 21].

Model	Symbol	Parameter meaning	Nominal value	Calibration interval
HCART	ϕ	Proliferation rate of effector CAR-T cells	0.265	[0.25, 0.29]
HCART	ρ	Reduction rate of effector CAR-T cells, including death and differentiation into memory cells	0.350	[0.32, 0.37]
HCART	θ	Conversion coefficient from memory to effector CAR-T cells following interaction with tumour cells	6.0×10^{-6}	$[2.4 \times 10^{-6}, 9.6 \times 10^{-6}]$
HCART	α	Tumour-induced inhibition coefficient of effector CAR-T cells	4.5×10^{-8}	$[1.8 \times 10^{-8}, 7.2 \times 10^{-8}]$
HCART	ϵ	Conversion rate from effector to memory CAR-T cells	0.15	[0.13, 0.24]
HCART	μ	Death rate of memory CAR-T cells	5.0×10^{-3}	$[2.0 \times 10^{-3}, 8.0 \times 10^{-3}]$
HCART	r	Maximum tumour growth rate	5.650026×10^{-2}	Fixed
HCART	b	Inverse of the tumour carrying capacity	1.404029×10^{-12}	Fixed
HCART	γ	Cytotoxic coefficient of effector CAR-T cells	3.715843×10^{-6}	Fixed

Continued on next page

Table 1 – continued from previous page

Model	Symbol	Parameter meaning	Nominal value	Calibration interval
CRS	η	Antigen-dependent activation and depletion rate of injected CAR-T cells	1	$[10^{-8},\ 30]$
CRS	μ_I	Death and distribution rate of injected CAR-T cells	1	$[10^{-3},\ 10]$
CRS	ν	Generation rate of expanding CAR-T cells from injected CAR-T cells	0.6	$[10^{-5},\ 10]$
CRS	κ	Antigen-dependent proliferation rate of expanding CAR-T cells	2	$[10^{-3},\ 10]$
CRS	ϵ	Differentiation rate from expanding to persistent CAR-T cells	10^{-2}	$[2 \times 10^{-5},\ 2 \times 10^{-1}]$
CRS	θ	Antigen-dependent reactivation rate from persistent to expanding CAR-T cells	10^{-8}	Fixed
CRS	μ_E	Death and exhaustion rate of expanding CAR-T cells	0.5	$[10^{-3},\ 10]$
CRS	μ_P	Death rate of persistent CAR-T cells	5×10^{-3}	$[10^{-5},\ 1]$
CRS	ρ	Tumour growth rate	0.1	$[0.0069,\ 0.255]$
CRS	K	Tumour carrying capacity	4.25×10^{12}	Fixed
CRS	γ	Maximum CAR-T-mediated tumour killing rate	1	$[0.1,\ 10]$
CRS	A	Half-saturation constant for antigen–receptor binding	10^7	$[1,\ 4.25 \times 10^{12}]$
CRS	B	Half-saturation constant for CAR-T-mediated tumour killing	10^8	$[1,\ 10^{10}]$
CRS	g_0	Reduced cytotoxicity factor for antigen-negative tumour cells	10^{-2}	$[10^{-3},\ 10^{-1}]$
CRS	σ_M	Basal production rate of inactive macrophages	5×10^8	$[10^8,\ 10^{10}]$
CRS	δ_M	Macrophage death rate	0.3	$[0.1,\ 1]$
CRS	σ_I	Endogenous production rate of IL_6	6	Fixed
CRS	α	IL_6 release rate by activated macrophages	10^{-7}	$[10^{-10},\ 10^{-5}]$
CRS	δ_I	Natural decay rate of IL_6	6	$[1,\ 24]$
CRS	β_B	Macrophage-activation coefficient associated with antigen-binding-mediated inflammatory signals	10^{-11}	$[10^{-13},\ 10^{-8}]$
CRS	β_K	Macrophage-activation coefficient associated with tumour-killing-mediated DAMP release	10^{-12}	$[10^{-15},\ 10^{-10}]$
CRS	β_C	Macrophage-activation coefficient associated with CD40–CD40L contact	10^{-10}	$[10^{-13},\ 10^{-8}]$
CRS	C	Half-saturation constant for contact-mediated macrophage activation	10^{10}	Fixed

2.4 Control-oriented reformulation of the models

In the original formulation of the CRS model, the CAR-T cell dose was given as the initial condition of the state variable C_I, which described the pool of injected CAR-T cells. Thus the therapeutic input was compressed into the model at time zero, and there was no explicit dosing function. As a consequence, the model is an autonomous system in which the treatment is embedded in the initial state rather than explicitly controlled. In the original HCART model, a similar situation arises: the CAR-T cell dose is encoded either directly in the initial condition of the effector CAR-T cell population, or implicitly through parameters associated with activation. As a result, the therapeutic intervention is not represented as an explicit input, but rather as a fixed quantity that cannot be manipulated once the simulation starts. This modelling choice, although mathematically convenient, limits the analysis of how different dosing strategies affect the system dynamics. In particular, it prevents the formal study of controllability and accessibility, and it does not allow the exploration of time-dependent therapeutic protocols such as repeated infusions or continuous administration.

To overcome this limitation, we reformulated both models by introducing the CAR-T dose as an explicit control input. This reparameterization preserves the biological mechanisms described by the original models, while providing a formulation suitable for structural analysis and future control-oriented applications. To this end, we introduce an external input function $u(t)$ that represents the CAR-T cell dose administered to the patient over time, which is incorporated additively into the CAR-T cell equation. Conceptually, $u(t)$ captures the infusion profile of CAR-T cells and allows the dose to be treated as a control variable rather than as a fixed initial condition. Thus, in the HCART model the equation for the effector CAR-T cell population (2) becomes:

$$\frac{dC_T}{dt} = u(t) + \phi C_T - \rho C_T + \theta T C_M - \alpha T C_T.$$

For the CRS model, an analogous reformulation was adopted. The state variable C_I, representing the injected CAR-T cell population, was retained, but the administered dose was decoupled from its initial condition and introduced as an explicit external input $u(t)$. The first equation of the CRS model therefore becomes:

$$\frac{dC_I}{dt} = u(t) - \eta F(T_P) C_I - \mu_I C_I.$$

3 Methodologies

3.1 Structural identifiability and observability

Structural identifiability determines whether the unknown parameter vector θ can be inferred from ideal, noise-free measurements of the outputs. A model is structurally globally identifiable (SGI) if its output trajectories cannot be identical for two different parameter vectors θ_A, θ_B, i.e.,

$$y(t|\theta_A) = y(t|\theta_B),\ \forall t \quad \Longrightarrow \quad \theta_A = \theta_B. \tag{16}$$

Conversely, if infinitely many parameter values generate the same observations, the corresponding parameters are structurally non-identifiable. In this work we analyse structural *local* identifiability (SLI), which relaxes the condition (16) by requiring only that it holds in a neighbourhood of the parameter value [22]. SLI entails that at most a finite number of parameter vectors can produce identical output trajectories. In practice, SLI parameters are typically also SGI [23].

Observability addresses whether the internal state vector can be reconstructed from the measured outputs. It is often evaluated as a local property, which becomes equivalent to SLI if the initial conditions are considered as parameters:

$$y(t|x(0)_A) = y(t|x(0)_B),\ \forall t \quad \Longrightarrow \quad x(0)_A = x(0)_B. \tag{17}$$

We refer to these two properties, structural local identifiability and observability, with the acronym SIO, and determine them with a differential geometric approach. This method evaluates both properties by constructing an observability-identifiability matrix from successive Lie derivatives of the measured outputs [24, 25]. By implicitly considering parameters as constant state variables, both vectors are treated in the same way. Thus, defining an augmented state vector as $\tilde{x} = [\theta, x]$, the first-order Lie derivative of the output function is defined as:

$$L_f h(\tilde{x}, u) = \frac{\partial h(\tilde{x}, u)}{\partial \tilde{x}} f(\tilde{x}, u), \tag{18}$$

and higher-order Lie derivatives are computed recursively, i.e., $L_f^2 h(\tilde{x}, u) = \frac{\partial L_f h(\tilde{x}, u)}{\partial \tilde{x}} f(\tilde{x}, u)$. The corresponding observability-identifiability matrix is given by

$$\mathcal{O}_I(\tilde{x}, u) = \frac{\partial}{\partial(\tilde{x})} \begin{bmatrix} h(\tilde{x}, u) \\ L_f h(\tilde{x}, u) \\ L_f^2 h(\tilde{x}, u) \\ \vdots \\ L_f^{n_x + n_\theta - 1} h(\tilde{x}, u) \end{bmatrix}. \tag{19}$$

If this matrix has full rank, the model is structurally locally identifiable and observable. Otherwise, rank deficiencies indicate structurally unidentifiable parameters and/or unobservable states; for further details, see [25]. The analyses were carried out using the STRIKE-GOLDD toolbox [15], an open-source MATLAB toolbox available on GitHub. For each model, different combinations of measured outputs were analysed in order to investigate how the available experimental observations affect SIO.

The SIO properties depend exclusively on the mathematical structure of the model, and therefore represent theoretical prerequisites for successful parameter estimation and state reconstruction. They do not take into account measurement noise, limited or poorly designed sampling instants, inadequate initial conditions and inputs, or optimization uncertainties. The challenges introduced by these additional limitations are taken into account in practical identifiability analyses, which are described in the next subsection.

3.2 Practical identifiability

While structural identifiability establishes whether model parameters can be determined from ideal, continuous, and noise-free observations, practical identifiability quantifies the robustness of parameter estimation under specific observation and numerical conditions [14]. In this study, we assess this property using a bootstrap method [26]. This procedure evaluates practical identifiability by repeatedly solving the joint parameter-estimation problem using pseudo-experimental data generated under specified measurement-noise conditions. For each model, noise-free pseudo-experimental data were obtained by sampling the nominal output at the observation times t_j. Let $y_j^\star = y^{\mathrm{mod}}(t_j; \boldsymbol{\theta}^\star)$, where $\boldsymbol{\theta}^\star$ denotes

the nominal parameter vector. Pseudo-experimental data were generated by multiplying each $y_j^\star$ by an independent log-normal factor with mean one and a coefficient of variation of 5%:

$$y_j^{\text{data}} = y_j^\star \exp\left(sZ_j - \frac{s^2}{2}\right), \qquad Z_j \overset{\text{iid}}{\sim} \mathcal{N}(0,1), \qquad s^2 = \ln\left(1 + 0.05^2\right). \tag{20}$$

Thus, at each sampling time, the expected observation was $y_j^\star$ and its standard deviation was $0.05\, y_j^\star$. This analysis was implemented in the AMIGO2 toolbox [17], using the method `AMIGO_RIdent`. For the 5% noise condition, `AMIGO_RIdent` was configured to perform 500 Monte Carlo trials per model, with all selected parameters estimated simultaneously. For parameters with positive nominal values, recovery was assessed using the normalized estimates $\widehat{\theta}_k/\theta_k^\star$. Distributions concentrated around one indicate consistent recovery of the nominal parameter, whereas broad or multimodal distributions and recurrent convergence to the parameter bounds suggest limited practical identifiability under the analysed conditions.

For the HCARTr model of Section 2.2, tumour burden T was sampled at 10 time points over the first 14 days after a bolus of 5×10^5 CAR-T cells. The simulation began immediately after the bolus, with $C_T(0^+) = 5 \times 10^5$ cells and no further CAR-T administration during the observation period. All nine parameters were included in the joint estimation:

$$\boldsymbol{\theta}_{\text{HCARTr}} = (\phi, \rho, \theta, \alpha, \epsilon, \mu, r, b, \gamma). \tag{21}$$

For the CRSr model described in Section 2.3, IL_6 concentration was the sole observed output and was sampled at the 12 observation times reported for patient G01 [21]. The total dose of 4.8×10^8 CAR-T cells was administered in three fractions, with 10%, 30%, and 60% given on days 0, 1, and 3, respectively. The first fraction was included in the initial condition, $C_I(0^+) = 4.8 \times 10^7$ cells. The two subsequent fractions were represented by instantaneous increases in C_I on days 1 and 3. All 23 parameters were included in the joint estimation from the IL_6 observations:

$$\begin{aligned}\boldsymbol{\theta}_{\text{CRSr}} = (&\eta, \mu_I, \nu, \kappa, \epsilon, \mu_E, \mu_P, \rho, \gamma, A, B, \sigma_M, \delta_M,\\ &\alpha, \delta_I, g_0, \beta_B, \beta_K, \beta_C, \theta, K, C_{\text{sat}}, \sigma_I).\end{aligned} \tag{22}$$

For each CRSr parameter with a positive nominal value, the search bounds were set to 0.4 and 1.6 times that value. These bounds were chosen for the estimation study. The parameter g_0 had a nominal value of zero and was allowed to vary over $[0, 1]$. Its estimates are therefore interpreted on the original scale.

3.3 Controllability and accessibility

While structural identifiability and observability address whether parameters and states can be inferred from measurements, controllability and accessibility describe the ability to influence the system dynamics through external inputs [27]. Controllability assesses whether the system can be driven to any desired point in the space by external inputs. Accessibility, also called reachability, determines whether the system can evolve towards a full-dimensional neighbourhood of its current state under the action of external inputs, which is a necessary condition for complete controllability.

For nonlinear systems, these properties can be studied within the framework of differential geometry [27, 28, 29]. In this approach, systems are typically written in control-affine form,

$$\dot{x} = f(x) + \sum_{i=1}^{n_u} g_i(x) u_i, \tag{23}$$

where f is the drift vector field and g_i are the input vector fields. The locally reachable directions are characterized by the distribution generated by these vector fields and their successive Lie brackets,

$$[f, g] = \frac{\partial g}{\partial x} f - \frac{\partial f}{\partial x} g. \tag{24}$$

The Lie Algebra Rank Condition (LARC) evaluates whether this accessibility distribution spans the full state space; if its rank equals the state dimension n_x at a given point, the system is locally accessible. This condition provides a differential-geometric characterization of accessibility and complements the Linearization Condition (LC) used to assess local controllability around equilibrium points. We evaluated these properties using the algorithms provided in the aforementioned STRIKE-GOLDD toolbox [16]. Complementarily, we also used StrucID [30], which analyses accessibility using sensitivity matrix computations, making it computationally efficient for large-scale or stiff nonlinear systems. To this end,

it computes a sensitivity rank condition (SERC), which determines the accessibility of linearizations of the system along trajectories [19], providing numerical evidence complementary to the symbolic analysis performed by STRIKE-GOLDD. The combined use of these tools allowed controllability and accessibility to be assessed from both symbolic and numerical perspectives, which provide complementary insights.

3.4 Sensitivity

Sensitivity analysis quantifies how variations in model parameters affect the model variables and outputs, thus pinpointing the parameters with the greatest dynamical influence. For the dynamical system defined in Eq. (1), the absolute sensitivity coefficient of the output variable y_i with respect to parameter θ_j at time t_k is defined as

$$s_{ij}(t_k) = \frac{\partial y_i(t_k)}{\partial \theta_j} = \lim_{\Delta\theta_j \to 0} \frac{y_i(t_k, \theta_j + \Delta\theta_j) - y_i(t_k, \theta_j)}{\Delta\theta_j}. \tag{25}$$

Stacking these coefficients for all outputs and parameters yields the sensitivity matrix,

$$S(t_k) = [s_{ij}(t_k)]_{i=1,\dots,n_y}^{j=1,\dots,n_\theta}, \tag{26}$$

which encapsulates the local dependence of the model outputs on the parameter vector. Since both parameters and outputs may span several orders of magnitude, it is common to use normalized (relative) sensitivities defined as

$$\tilde{S}_{ij}(t_k) = s_{ij}(t_k)\frac{\theta_j}{y_i(t_k)}, \tag{27}$$

thus allowing direct comparison of the influence of different parameters on a common scale.

To compute sensitivities, we developed a custom numerical method based on the approach proposed by Mester *et al.* [18]. Our MATLAB implementation adapts the original Julia code, approximating the derivatives using the Complex-Step Derivative (CSD) method. Rather than perturbing the parameters with finite real increments, the CSD method evaluates the model in the complex plane by replacing the parameter θ_j with $\theta_j + i\Delta$, where Δ is a sufficiently small perturbation. The corresponding derivative is approximated as

$$\frac{\partial y_i}{\partial \theta_j} \approx \frac{\mathrm{Im}[y_i(\theta_j + i\Delta)]}{\Delta}, \tag{28}$$

which avoids the subtractive cancellation errors associated with classical finite-difference approximations while maintaining machine-precision accuracy.

The sensitivity analysis focused on the clinically most relevant output for each model. For the HCART model, tumour burden (T) was selected as the output of interest, whereas for the CRS model the analysis focused on the concentration of IL_6 because of its central role in Cytokine Release Syndrome. For each parameter, normalized sensitivities were evaluated over a range extending from zero to five times its nominal value, following the parameter ranges adopted in the original model publications. Sensitivity maps were generated by representing the normalized sensitivity as a function of both simulation time and parameter value, providing a comprehensive visualization of how parameter influence evolves throughout the simulation. Different numerical solvers, tolerances, and parameter scales were explored to ensure the robustness of the computed sensitivity profiles.

It is important to note that sensitivity analysis and practical identifiability address fundamentally different questions. Sensitivity analysis quantifies the local influence of infinitesimal parameter perturbations assuming that all remaining parameters remain fixed, whereas practical identifiability evaluates whether parameter values can be uniquely estimated from finite and noisy experimental data. Therefore, a parameter may exhibit high sensitivity while remaining practically non-identifiable due to parameter correlations or insufficient experimental information. Therefore, sensitivity analysis complements both the structural and practical identifiability analyses performed in this work. Besides identifying the parameters that most strongly influence clinically relevant outputs, it provides valuable guidance for model calibration, experimental design and the interpretation of practical identifiability results.

Furthermore, the obtained sensitivity profiles highlight parameters associated with therapeutic interventions whose modulation may have the greatest impact on system dynamics, thereby supporting future model-based optimization and control strategies. In this way, they provide quantitative counterparts of the qualitative (yes/no) information provided by accessibility and controllability analyses.

4 Results

4.1 Structural identifiability and observability

When assess the SIO of both models, we focused on two key aspects, namely the impacts of (I) considering different output functions, and (II) including or omitting external inputs as explained in Section 2.4.

The results for the HCART model are summarized in Table 2. It can be seen that, when the measured output is T, the reformulated model has better identifiability and observability properties than the original one, as the parameter γ is now identifiable and the whole set of state variables is observable. In contrast, if the output is the memory CAR-T cells C_M, the model remains with the same identifiability and observability properties in the original and reformulated forms.

The results for the CRS model are summarized in Table 3. As in the HCART model, the control-oriented model reformulation improves identifiability: for the same set of outputs, more parameters of the reformulated model become identifiable. This improvement is consistent across different output configurations. When the number of measured outputs is reduced, like in the case where only IL_6 is measured, both models, original and reparameterized, show a drop in identifiability. Regarding observability, while most internal states retain the same observability status in both models, the state C_i (representing the injected CAR-T cell compartment) becomes observable in the reparameterized model, whereas it is unobservable in the original formulation with every possible combination of measured outputs. This improvement is small but notable, as C_i plays a key role in early CAR-T dynamics and its observability can facilitate better inference of related parameters.

Table 2: Structural identifiability and state observability results for the HCART model under the analysed output and input configurations. HCART* denotes the formulation including the known input u. Quantities shown in red are structurally non-identifiable or non-observable, whereas quantities shown in black are structurally identifiable or observable.

Model	Outputs directly measured	Known input	Identifiability of parameters	Observability of states
HCART*	$\boldsymbol{T}$	u	ϕ, ρ, ϵ, θ, α, μ, r, b, γ	C_T, C_M, $\boldsymbol{T}$
HCART	$\boldsymbol{T}$	–	ϕ, ρ, ϵ, θ, α, μ, r, b, γ	C_T, C_M, $\boldsymbol{T}$
HCART*	$\boldsymbol{C_M}$	u	ϕ, ρ, ϵ, θ, α, μ, r, b, γ	C_T, $\boldsymbol{C_M}$, T
HCART	$\boldsymbol{C_M}$	–	ϕ, ρ, ϵ, θ, α, μ, r, b, γ	C_T, $\boldsymbol{C_M}$, T

Table 3: Structural identifiability and state observability results for the CRS model under the analysed output and input configurations. CRS* denotes the reparameterized formulation of the model. Quantities shown in red are structurally non-identifiable or non-observable, whereas quantities shown in black are structurally identifiable or observable. Boldface indicates the outputs that are directly measured in each configuration.

Model	Outputs directly measured	Known input	Identifiability of parameters	Observability of states
CRS*	$\boldsymbol{C_E, M_i, M_a, IL_6}$	u	η, μ_I, ν, κ, ϵ, θ, μ_E, μ_P, ρ, K, γ, A, B, C, σ_M, δ_M, σ_I, α, δ_I, g_0, β_B, β_K, β_C	C_I, $\boldsymbol{C_E}$, C_P, T_P, T_N, $\boldsymbol{M_a}$, $\boldsymbol{M_i}$, $\boldsymbol{IL_6}$
CRS	$\boldsymbol{C_E, M_i, M_a, IL_6}$	–	η, μ_I, ν, κ, ϵ, θ, μ_E, μ_P, ρ, K, γ, A, B, C, σ_M, δ_M, σ_I, α, δ_I, g_0, β_B, β_K, β_C	C_I, $\boldsymbol{C_E}$, C_P, T_P, T_N, $\boldsymbol{M_a}$, $\boldsymbol{M_i}$, $\boldsymbol{IL_6}$
CRS*	$\boldsymbol{C_E, M_i + M_a, IL_6}$	u	η, μ_I, ν, κ, ϵ, θ, μ_E, μ_P, ρ, K, γ, A, B, C, σ_M, δ_M, σ_I, α, δ_I, g_0, β_B, β_K, β_C	C_I, $\boldsymbol{C_E}$, C_P, T_P, T_N, M_a, M_i, $\boldsymbol{IL_6}$
CRS	$\boldsymbol{C_E, M_i + M_a, IL_6}$	–	η, μ_I, ν, κ, ϵ, θ, μ_E, μ_P, ρ, K, γ, A, B, C, σ_M, δ_M, σ_I, α, δ_I, g_0, β_B, β_K, β_C	C_I, $\boldsymbol{C_E}$, C_P, T_P, T_N, M_a, M_i, $\boldsymbol{IL_6}$
CRS*	$\boldsymbol{IL_6}$	u	η, μ_I, ν, κ, ϵ, θ, μ_E, μ_P, ρ, K, γ, A, B, C, σ_M, δ_M, σ_I, α, δ_I, g_0, β_B, β_K, β_C	C_I, C_E, C_P, T_P, T_N, M_a, M_i, $\boldsymbol{IL_6}$
CRS	$\boldsymbol{IL_6}$	–	η, μ_I, ν, κ, ϵ, θ, μ_E, μ_P, ρ, K, γ, A, B, C, σ_M, δ_M, σ_I, α, δ_I, g_0, β_B, β_K, β_C	C_I, C_E, C_P, T_P, T_N, M_a, M_i, $\boldsymbol{IL_6}$

4.2 Practical Identifiability

The practical identifiability analysis was performed using the `AMIGO_RIdent` configuration and the pseudo-experimental observation schemes described in Section 3.2. Each analysis comprised 500 RIdent trials.

The first trial fitted the supplied pseudo-experimental dataset, and subsequent trials used datasets generated by the Monte Carlo procedure, with the preceding trial's estimates used as the initial parameter guess. Consequently, the resulting parameter distributions reflect the combined effects of measurement uncertainty, parameter interactions, and numerical optimization.

Figure 2 compares the pseudo-experimental observations with the model outputs obtained using the parameter estimates from the first trial. The close-ups in the right side show the differences between predictions and observations at the sampling times with the largest absolute relative discrepancies.

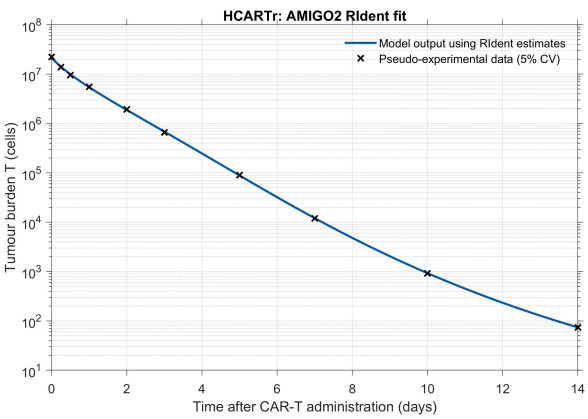

(a) HCARTr fit.

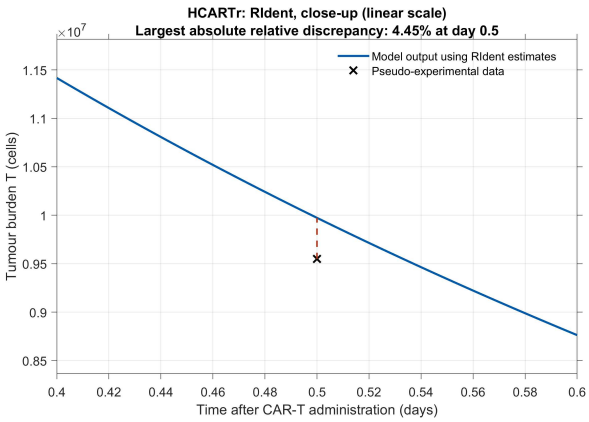

(b) HCARTr close-up.

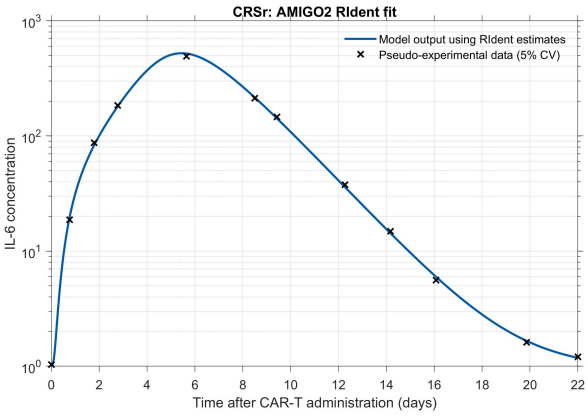

(c) CRSr fit.

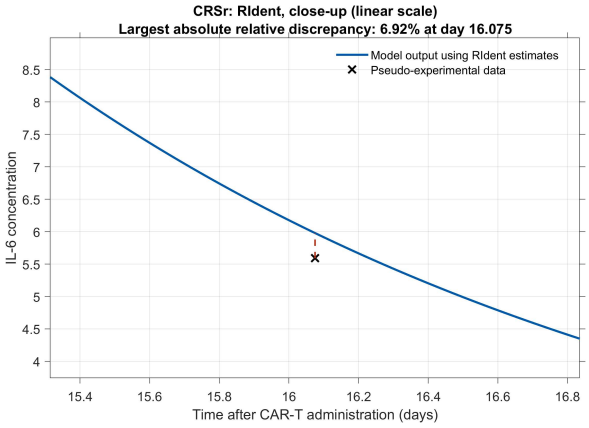

(d) CRSr close-up.

Figure 2: Model outputs obtained from `AMIGO_RIdent` and pseudo-experimental observations generated with a coefficient of variation of 5%. Blue lines represent model outputs and black crosses represent observations. Complete trajectories are shown on logarithmic scales, whereas close-ups use linear scales. Dashed segments indicate prediction-observation differences.

The HCARTr model yielded different identification patterns for the estimated parameters. Figure 3 illustrates four representative patterns of parameter recovery across the 500 Monte Carlo trials: a comparatively narrow distribution within restricted search bounds (ϕ), a dispersed distribution with lower-bound accumulation (ϵ), pronounced accumulation at both search bounds (μ), and a broad distribution with a near-nominal median (γ). The normalized estimates of ϕ and ρ had medians of approximately 0.983 and 1.020, with empirical central 95% ranges of 0.943–1.041 and 0.965–1.057, respectively. These comparatively narrow distributions should be interpreted in light of the restricted search bounds imposed on both parameters. The remaining parameters showed considerably worse recovery. Estimates of θ, α, μ, r, and b extended across most of their admissible intervals. The distribution of ϵ showed substantial dispersion, with accumulation near its lower bound and a considerable fraction of estimates above its nominal value. Its median normalized estimate was approximately 1.291, and its 97.5th percentile was close to the upper bound of 1.6. For γ, the median normalized estimate was approximately 0.993, with a central 95% range of 0.649–1.352. Thus, its median was close to the nominal value, although the estimates remained substantially dispersed. The complete set of parameter-estimate distributions is provided in Supplementary Figure S1.

Furthermore, the analyses revealed strong dependencies among pairs of parameters, as shown in

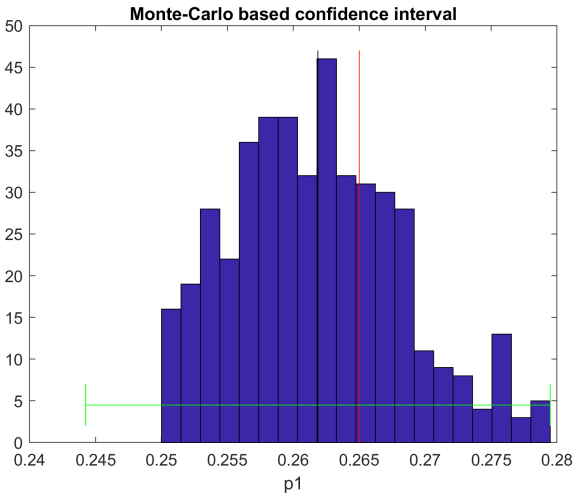

(a) Comparatively narrow distribution (ϕ).

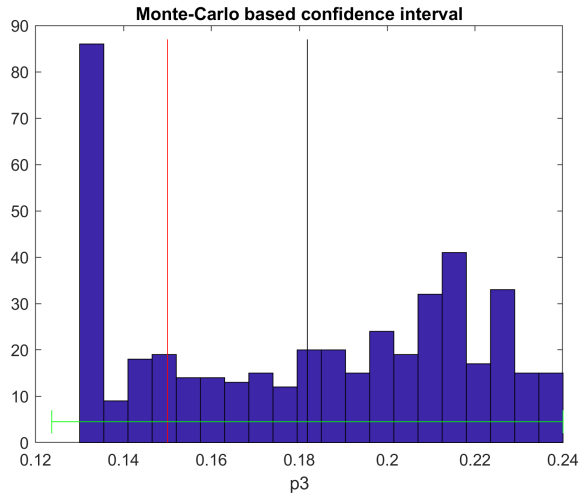

(b) Dispersion and lower-bound accumulation (ϵ).

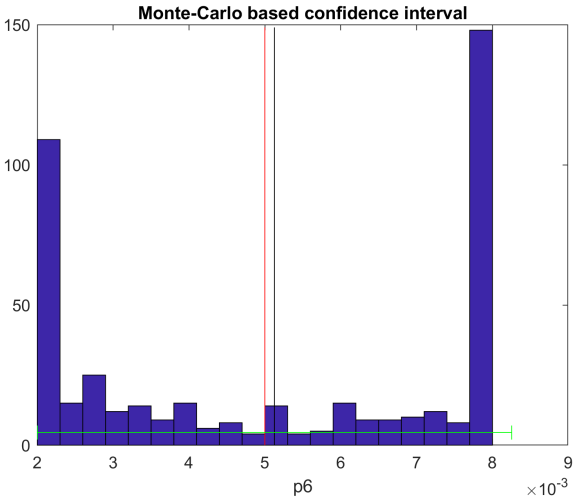

(c) Accumulation at both search bounds (μ).

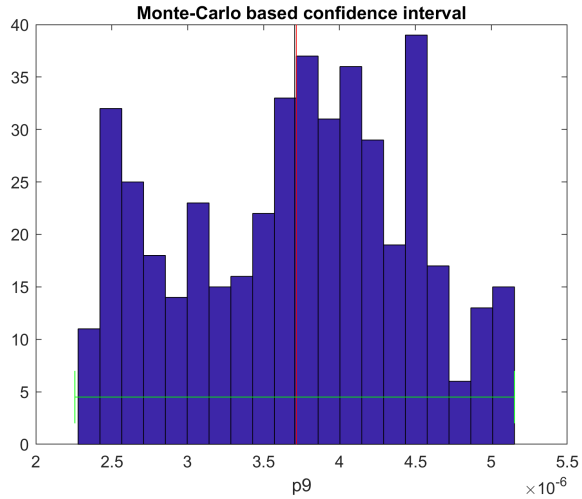

(d) Broad distribution with a near-nominal median (γ).

Figure 3: Representative distributions of HCARTr parameter estimates obtained from 500 joint Monte Carlo estimation trials. The horizontal axes show parameter values on their original scales, and bar heights indicate the number of estimates in each bin. The parameter labels p_1, p_3, p_6, and p_9 correspond to ϕ, ϵ, μ, and γ, respectively. The comparatively narrow distribution of ϕ should be interpreted in the context of its restricted search bounds and its dependence on ρ.

supplementary figure S2. Pearson correlations calculated from all 500 estimates were approximately -0.984 between ϕ and ρ, and 0.965 between α and γ. Moreover, ϕ and ρ enter the model equations through their difference, $\phi - \rho$, which limits their independent determination.

Overall, the HCARTr results indicate that observations of tumour burden constrain some parameter combinations more effectively than the individual parameters. Although ϕ and ρ exhibited comparatively narrow marginal distributions within their restricted search bounds, the remaining parameters showed substantial dispersion, displacement from their nominal values, or accumulation near the admissible boundaries.

For CRSr, the parameter-estimate distributions also revealed substantial differences in practical recovery from IL_6 measurements. The normalized estimates of θ, μ_P, and K occupied comparatively narrow regions near their nominal values. However, the distributions of μ_P and K showed separated groups of estimates rather than a single concentrated peak. These patterns indicate comparatively restricted marginal variability, but do not by themselves establish independent or robust parameter recovery. Most remaining parameters exhibited broad distributions, displacement from their nominal values, or accumulation near the search bounds. The estimates of η and μ_I accumulated near both bounds, whereas ν, σ_I, and δ_I were concentrated near the lower bound, with tails extending towards larger values. Conversely,

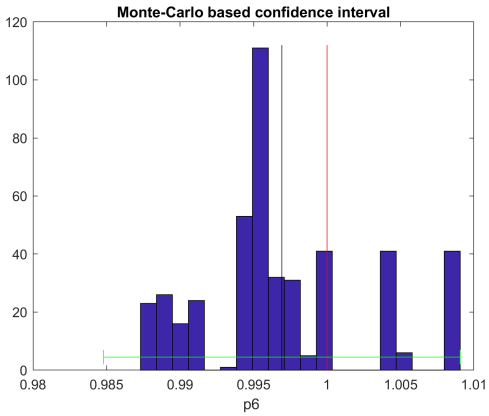

(a) Comparatively narrow distribution (θ).

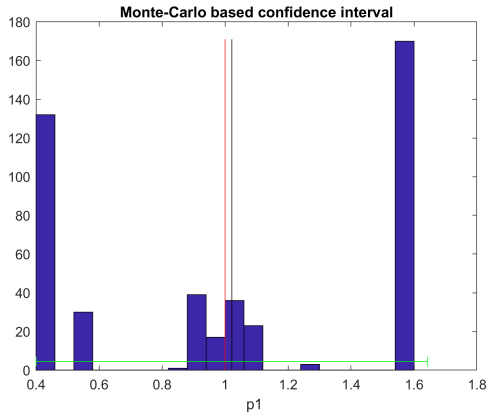

(b) Accumulation near both search bounds (η).

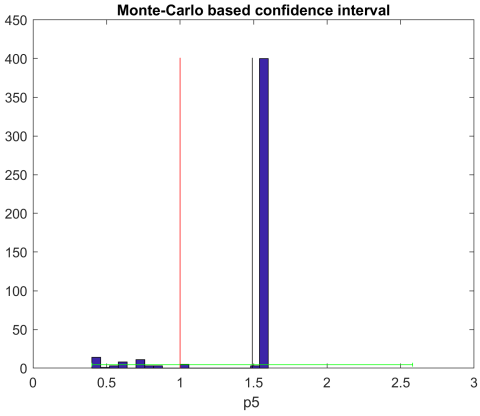

(c) Upper-bound accumulation (ϵ).

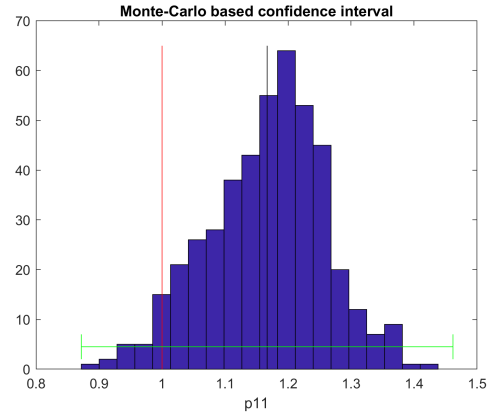

(d) Distribution displaced above the nominal value (γ).

Figure 4: Representative distributions of CRSr parameter estimates obtained from the Monte Carlo estimation procedure. The horizontal axes show parameter values normalized by their nominal values, so that unity denotes the nominal value. Bar heights indicate the number of estimates in each bin; the displayed histograms each contain 451 estimates. Each panel uses its own axis scales. The parameter labels p_6, p_1, p_5, and p_{11} correspond to θ, η, ϵ, and γ, respectively.

ϵ, ρ, and C showed pronounced accumulation near the upper bound. The distribution of γ was more broadly unimodal, with its main peak above the nominal value. Thus, a sharply concentrated histogram did not necessarily indicate accurate recovery: several prominent peaks occurred at the admissible boundaries rather than near the nominal parameter values. Figure 4 illustrates four representative patterns: a comparatively narrow distribution near the nominal value (θ), accumulation near both search bounds (η), pronounced upper-bound accumulation (ϵ), and a broader distribution displaced above the nominal value (γ). The complete set of parameter-estimate distributions is provided in Supplementary Figures S3–S5. Notably, the distribution of g_0 was concentrated at zero. Under the G01 initial conditions, the antigen-negative tumour population remains zero, making the observed output independent of g_0. Its constant estimated value therefore provides no evidence of identification from these observations.

Overall, the CRSr results show comparatively restricted marginal variability for θ, μ_P, and K, whereas most other parameters exhibit substantial dispersion, displacement, or boundary accumulation. These findings support limited individual parameter recovery under the observation scheme considered.

Table 4 compares structural local identifiability and practical parameter recovery for HCARTr and CRSr. The structural assessment considers the known-input formulation of each model, with tumour burden T as the sole measured output for HCARTr and IL_6 concentration as the sole measured output for CRSr.

Overall, the displayed fits show that close agreement with the observed outputs can coexist with substantial variability in parameter estimates. The RIdent populations reveal limitations in individual parameter recovery.

Model	Assessment	Parameters and interpretation
HCARTr	Structurally locally identifiable	ϵ, θ, α, μ, r, b, and γ
HCARTr	Practical RIdent result	ϕ and ρ exhibited comparatively narrow marginal distributions within restricted search bounds, but their strong dependence and occurrence through $\phi - \rho$ prevent interpreting this as independent recovery. γ had a near-nominal median but substantial dispersion. ϵ, θ, α, μ, r, and b showed broad distributions and/or boundary accumulation.
CRSr	Structurally locally identifiable	η, μ_I, κ, ϵ, θ, μ_E, μ_P, ρ, γ, δ_M, σ_I, δ_I, and g_0
CRSr	Practical RIdent result	θ, μ_P, and K exhibited comparatively narrow normalized distributions near their nominal values, although μ_P and K showed separated groups of estimates. Most remaining parameters exhibited broad distributions, displacement, or boundary accumulation. The concentration of g_0 at zero does not indicate recovery because the observed output is independent of this parameter under the initial conditions considered.

Table 4: Structural local identifiability and practical parameter recovery for HCARTr and CRSr. The structural results refer to the known-input formulations. Practical recovery depends on the initial conditions, dosing schedule, sampling times, and parameter search bounds used for calibration.

4.3 Controllability and accessibility

The controllability and accessibility of the reformulated HCART model (HCARTr) were focused on two biologically relevant equilibrium points: the tumour-free equilibrium ($E_0 = (0, 0, 0)$) and the tumour-bearing equilibrium ($E_T = (0, 0, 1/b)$), corresponding to the tumour carrying capacity in the absence of CAR-T cells. The results consistently showed that the controllability and accessibility of HCARTr depend on the equilibrium considered. At (E_0), LC and LARC yielded rank 2 in the three-dimensional state space. In contrast, at (E_T), all three conditions consistently yielded full rank (3/3). Consequently, LC supported short-time local controllability around the tumour-bearing equilibrium, while and LARC established local accessibility. Overall, these findings indicate that the ability to influence the model dynamics through CAR-T administration depends on the operating point, with full-rank controllability and accessibility conditions satisfied at the tumour-bearing equilibrium.

The analysis of the reformulated CRS model (CRSr) was considerably more computationally demanding because of its eight-dimensional state space and nonlinear interactions. STRIKE-GOLDD obtained LC ranks of 1/8 at the tumour-free equilibrium and 6/8 at the tumour-bearing equilibrium, indicating that the linearization condition for controllability was not satisfied at either point. However, these rank deficiencies do not establish nonlinear uncontrollability, since the LC test provides a sufficient but not necessary condition. The LARC analysis remained inconclusive because successive Lie-bracket computations became computationally prohibitive, even after increasing the maximum computation time. Hence, to complement these results, we investigated accessibility numerically using StrucID. Several initial conditions were examined, including mostly biologically admissible equilibrium points. The numerical input value assigned by StrucID was retained during these analyses. All cases were classified as locally unreachable (i.e., not accessible), although the number of numerically significant directions depended strongly on the initial condition. The tumour-free equilibrium exhibited only one numerically significant direction, whereas each equilibrium containing the tumour population T_P and the mixed tumour equilibrium exhibited seven out of eight. In these cases, the numerically deficient direction was predominantly associated with the inactive and activated macrophage populations, M_i and M_a. This realization led us to introduce an additional state variable representing the total macrophage population, $M_{ia} = M_i + M_a$, the equation of which is given by

$$\frac{dM_{ia}}{dt} = \sigma_M - \delta_M M_{ia}. \tag{29}$$

The dynamics of this combined variable is independent of the CAR-T input, since the macrophage activation terms cancel. The additional StrucID analysis did not identify the combined macrophage variable as locally reachable either. Importantly, M_{ia} is a redundant state rather than an additional independent degree of freedom.

Overall, the HCARTr model satisfied full-rank controllability and accessibility conditions at the tumour-bearing equilibrium. In contrast, the CRSr model exhibited incomplete LC ranks, computationally inconclusive Lie-bracket analyses, and numerical reachability limitations across the tested initial conditions.

4.4 Sensitivity

The sensitivity analyses revealed substantial heterogeneity in the magnitude, sign, and temporal distribution of the computed sensitivities. These features should be considered jointly: a parameter may exhibit a pronounced effect within a narrow region of the explored domain, whereas another may retain a consistent sensitivity sign but very small magnitudes.

For HCARTr, four qualitative sensitivity patterns were distinguished according to the magnitude of the response and its dependence on time and parameter value. These groups summarize the dominant features of the maps rather than defining mutually exclusive behaviours. All parameters exhibited temporal variation, with their largest relative sensitivity magnitudes occurring towards the end of the simulated interval. Figure 5 illustrates these four patterns using ϕ, α, θ, and b, respectively.

1. The first group comprised ϕ and ρ, which exhibited pronounced sensitivities with a strong temporal dependence across their explored parameter ranges. The map for ϕ provides a representative example: sensitivity magnitudes increased progressively with time over the entire range, with additional variation along the parameter axis. The response to ρ followed a similar temporal pattern but had the opposite sign. Increasing ϕ locally decreased tumour burden, whereas increasing ρ increased it.

2. The second group comprised α, ϵ, and γ, whose pronounced sensitivities also depended substantially on the parameter value. This pattern was particularly evident for α: large responses were concentrated towards higher parameter values and later times, while lower values produced much smaller responses. The sensitivity to ϵ showed a similar tendency towards larger magnitudes at higher parameter values. The map for γ displayed a more complex dependence on parameter value, indicating that the strength of its effect could not be described by a simple monotonic pattern across the entire domain. The response to α was positive, whereas those to ϵ and γ were negative.

3. The third group comprised θ and r, which exhibited intermediate sensitivity magnitudes relative to the preceding groups. The map for θ was dominated by a gradual increase in sensitivity magnitude over time, with comparatively limited variation across its parameter range. The response to r also increased with time but showed a clearer dependence on parameter value. Their signs differed: increasing θ locally reduced tumour burden, whereas increasing r increased it.

4. The fourth group comprised μ and b, which exhibited weak relative sensitivities throughout the explored domains. The map for b provides the clearest example: although it displayed visible temporal and parameter-dependent variation, the corresponding magnitudes remained very small. The response to μ was larger than that to b but remained weak compared with the other parameters. Neither sensitivity was identically zero. Their visible patterns should therefore be interpreted alongside the individual greyscales, which can make weak responses appear as prominent as much stronger ones.

Overall, the HCARTr maps distinguish parameters with pronounced, temporally increasing effects from those whose influence depends more strongly on the explored parameter values or remains weak throughout the domain. The complete set of sensitivity maps is provided in Supplementary Figure S7.

For CRSr, the sensitivity responses could again be classified into four qualitatively distinct patterns, according to the temporal extent of the response, its localization within the explored parameter range, and its magnitude. As for HCARTr, these groups describe the dominant features of the maps rather than mutually exclusive behaviours. Figure 6 illustrates the four patterns using σ_I, α, γ, and g_0, respectively.

1. The first group comprised σ_I and δ_I, which exhibited sustained sensitivities over broad regions of the parameter–time domain. The map for σ_I provides a representative example: after an initial transient, substantial sensitivity persisted throughout most of the simulated interval and across much of the explored parameter range. The response to δ_I was also sustained, although its magnitude varied during the initial transient. Their signs were opposite: increasing σ_I increased IL_6, whereas increasing δ_I decreased it.

2. The second group comprised η, μ_I, ν, ϵ, μ_E, ρ, A, B, σ_M, α, and β_B. Their largest responses were generally concentrated in the earlier or intermediate part of the simulation and diminished towards its end. The map for α illustrates this pattern particularly clearly, with a pronounced response at early times and higher parameter values. Similar patterns were observed for β_B, ν, and σ_M, although the timing and magnitude differed. Other members of this group, including η, μ_I, and μ_E, exhibited more localized regions along the parameter axis. The response to B peaked later than that to α, while ϵ showed a similar temporal localization but much smaller magnitudes. Thus, membership of this group indicates a shared temporal pattern rather than comparable sensitivity strength.

3. The third group comprised γ, κ, and δ_M, whose responses were strongly dependent on particular regions of the parameter–time domain or extended into later times. The map for γ provides the clearest example of a late localized response: pronounced negative sensitivity emerged towards the end of the simulation within a restricted range of low parameter values. The response to δ_M was also stronger at lower parameter values but extended over a broader temporal region. In contrast, κ exhibited sharp,

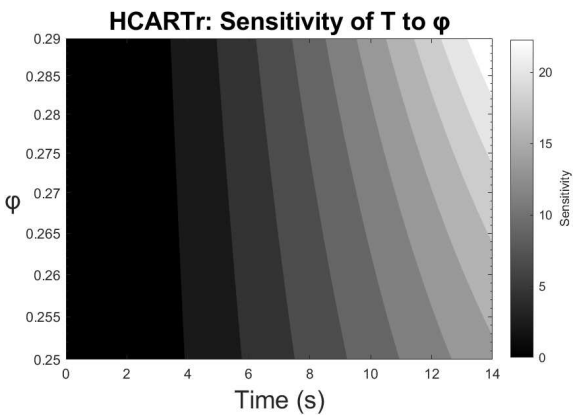

(a) Pronounced sensitivity with strong temporal variation (ϕ).

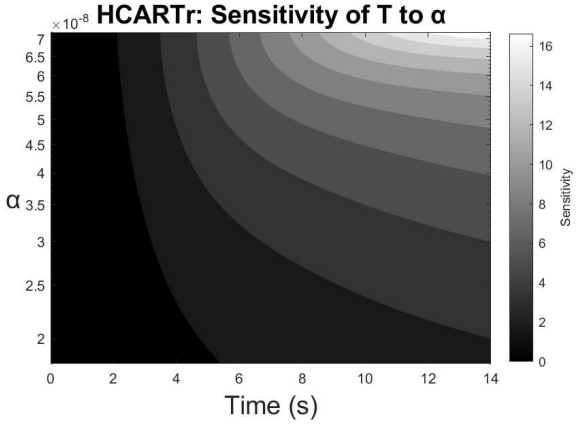

(b) Pronounced sensitivity dependent on parameter value (α).

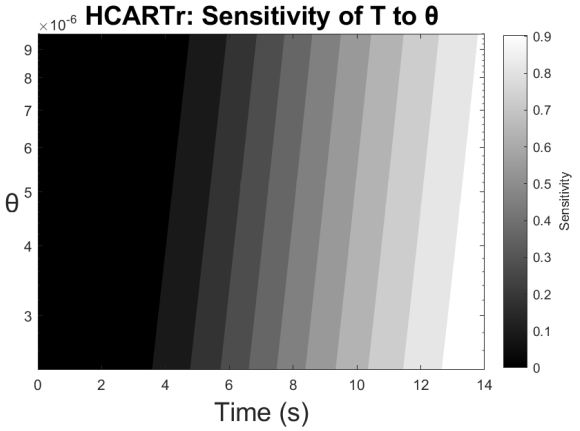

(c) Intermediate sensitivity with gradual temporal variation (θ).

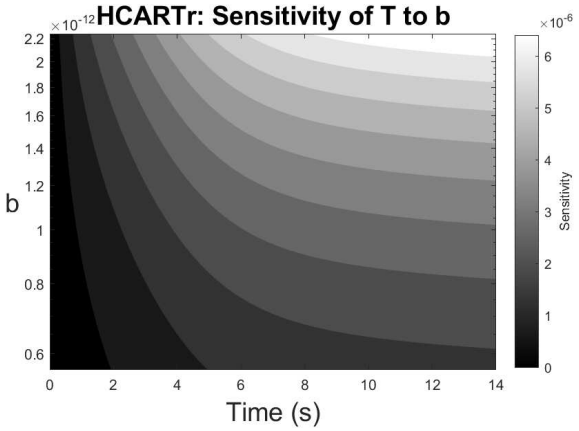

(d) Very weak sensitivity throughout the explored domain (b).

Figure 5: Representative relative sensitivity magnitude maps for tumour burden T in HCARTr. The panels illustrate the dominant patterns used to group the parameters qualitatively. Each panel uses an independent greyscale, with lighter regions indicating larger magnitudes within that panel. Shading does not indicate the sign of the sensitivity, and similar shading across panels does not imply comparable magnitudes.

narrow sensitivity regions and changes of sign. Its large computed magnitudes therefore describe localized responses and should not be interpreted as a uniformly strong influence throughout the domain.

4. The fourth group comprised θ, μ_P, K, C, g_0, β_K, and β_C, which exhibited small or vanishing relative sensitivities under the analysed conditions. Among these parameters, θ, μ_P, and K produced particularly small responses throughout their explored ranges. The sensitivities to C, β_K, and β_C were larger but remained comparatively weak and were mainly concentrated in the earlier part of the simulation. The parameter g_0 represents the limiting case of this group: its sensitivity was identically zero because the antigen-negative tumour population remained absent under the initial conditions considered.

The signs of the responses provided an additional distinction within these groups. The sensitivities to θ, ρ, K, σ_M, σ_I, α, β_B, β_K, and β_C were non-negative, whereas those to ϵ, μ_E, μ_P, γ, C, δ_M, and δ_I were non-positive. The remaining non-zero responses changed sign within the explored domain, although η, μ_I, and A were predominantly negative, while ν and B were predominantly positive. These signs were obtained from the signed sensitivity matrices; the maps themselves display magnitudes.

Overall, the CRSr maps distinguish sustained parameter effects from transient or localized responses and from weak or absent effects on IL_6. These differences indicate that the information supplied by an observation depends on both its timing and the parameter region considered. However, neither a pronounced sensitivity nor a distinctive temporal pattern alone establishes independent parameter recovery. The complete set of CRSr sensitivity maps is provided in Supplementary Figures S8–S10.

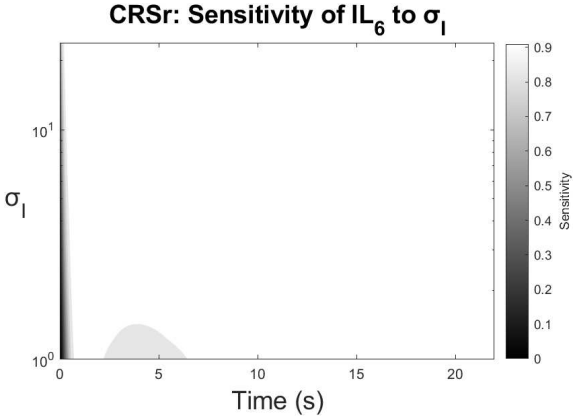

(a) Sustained sensitivity over a broad domain (σ_I).

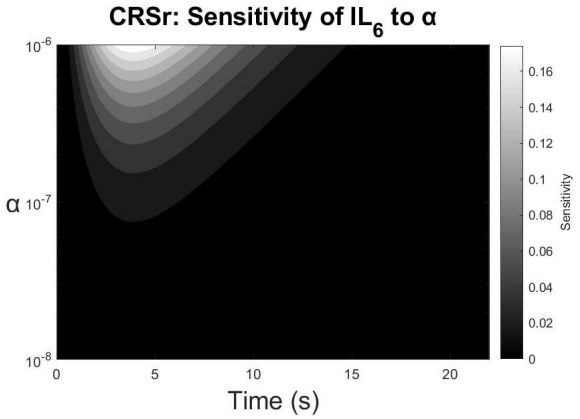

(b) Sensitivity concentrated at earlier times (α).

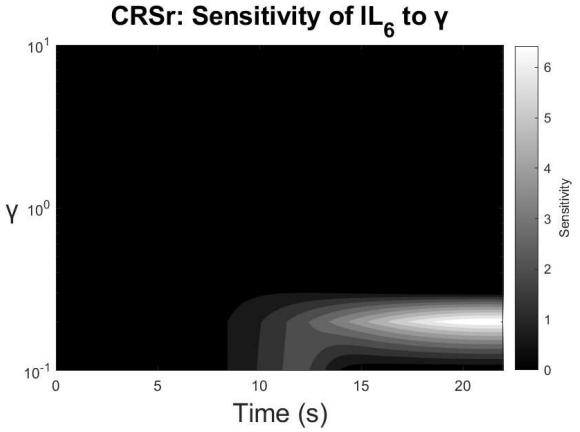

(c) Late response within a restricted parameter region (γ).

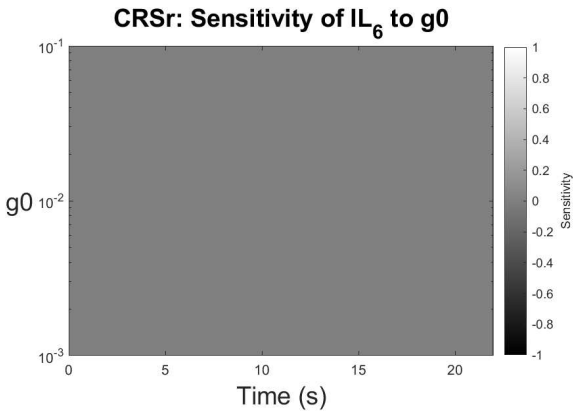

(d) Zero sensitivity throughout the explored domain (g_0).

Figure 6: Representative relative sensitivity magnitude maps for IL_6 in CRSr. The panels illustrate the dominant patterns used to group the parameters qualitatively. Parameter axes are logarithmic, and each panel uses an independent greyscale. For non-zero maps, lighter regions indicate larger magnitudes within a panel; shading does not indicate the sign of the sensitivity. The uniform shading in panel (d) represents an identically zero response; its displayed colour limits do not indicate non-zero sensitivity values.

5 Discussion

In this work, we have presented a comprehensive assessment of our ability to identify and control CAR-T cell therapy by leveraging mechanistic modelling. To this end, we have analysed their structural identifiability, observability, controllability, accessibility, sensitivity, and practical identifiability, examining the results jointly and systematically. While each methodology provides different information about the mathematical properties of the models, their combined interpretation offers a more complete overview. The results highlight both the theoretical capabilities of the models and the practical limitations that arise when estimating parameters from available observations.

Several limitations should be acknowledged. Practical identifiability was assessed using model-generated pseudo-experimental data with a coefficient of variation of 5%, under one measurement configuration for each model. These conditions do not capture model discrepancy, inter-patient variability, or the full limitations of biological measurements. The estimated distributions also depend on parameter bounds and the numerical estimation procedure, while sensitivity patterns depend on the explored parameter ranges.

5.1 From structural to practical identifiability

A first important observation is the distinction between structural and practical identifiability. The former assesses whether model parameters can be uniquely determined, locally or globally, from ideal,

noise-free observations of the specified outputs under the assumed input conditions. In contrast, the results of the practical analyses showed that structural local identifiability does not necessarily translate into precise parameter estimates. Several structurally identifiable parameters exhibited broad, multimodal or boundary-concentrated distributions. These patterns are consistent with limited information in the measured outputs and possible parameter compensation, although optimization behaviour and the imposed parameter bounds may also contribute.

For HCARTr, the comparatively narrow marginal distributions of ϕ and ρ must be interpreted alongside their restricted search bounds and strong dependence. Their occurrence through the combination $\phi - \rho$ limits their independent determination, illustrating why narrow marginal distributions alone cannot establish parameter identifiability. The remaining parameters generally exhibited substantial dispersion, displacement from their nominal values, or boundary accumulation. Even for γ, whose median was close to its nominal value, the broad distribution indicated considerable uncertainty. Thus, proximity of a central estimate to the nominal value should be distinguished from consistent recovery across trials.

These limitations were also apparent in CRSr. Although fixing the parameters recommended in the original publication substantially improved structural identifiability, several parameters remained poorly constrained in the practical analysis. The comparatively narrow distributions of θ, μ_P, and K contrasted with widespread dispersion and boundary accumulation among the other parameters. Furthermore, the separated groups of estimates observed for μ_P and K indicate that restricted marginal variability need not imply a single well-defined estimation region. The greater complexity of CRSr may contribute to parameter dependencies, but differences between the models cannot be attributed to complexity alone: the measured outputs, initial conditions, dosing schedules, and estimation settings also determine the information available for parameter recovery.

The behaviour of g_0 further illustrates the importance of the experimental conditions. Although this parameter was structurally locally identifiable in the structural assessment, the antigen-negative tumour population remained zero under the initial conditions used for practical estimation. Consequently, the measured output contained no information about g_0 in this experiment. Its concentration at zero therefore reflects the absence of informative dynamics rather than successful recovery, highlighting the distinction between structural results under generic conditions and parameter estimation under a specific experimental setting.

The close agreement between the fitted trajectories and the pseudo-experimental observations should therefore be interpreted solely as evidence of agreement for the displayed trials, rather than unique parameter determination. Such agreement can coexist with substantial uncertainty in individual parameters and does not establish predictive reliability under different treatment schedules or initial conditions. Moreover, the sequential use of the preceding trial's estimates as the next initial guess means that the RIdent distributions reflect both data perturbations and the estimation procedure; they should not be interpreted as posterior parameter distributions.

Together, these findings motivate considering identifiable parameter combinations and additional informative observations when refining the models. In particular, measurements of CAR-T or immune-cell populations, alongside tumour burden or IL_6, could help distinguish mechanisms that remain difficult to separate using a single observed output. The benefit of such measurements would need to be assessed through further analysis and experimental design studies.

5.2 Controllability and accessibility

The controllability and accessibility results highlight the importance of the operating point when assessing the influence of CAR-T administration.

For HCARTr, full-rank conditions at the tumour-bearing equilibrium supported local controllability and accessibility, whereas the tumour-free equilibrium exhibited rank deficiencies in LC and LARC. Thus, the favourable results at the tumour-bearing equilibrium cannot be extended to the entire state space.

For CRSr, incomplete linearization ranks and computationally inconclusive Lie-bracket analyses limited the symbolic assessment of controllability and accessibility via the LC and LARC tests, respectively. The numerical computations performed with StrucID classified all tested initial conditions as locally unreachable. The macrophage dynamics provide a mechanistic explanation for at least part of this limitation: the total population, $M_i + M_a$, evolves independently of the CAR-T input. Consequently, the input cannot independently prescribe this total, even if it influences the distribution between inactive and activated macrophages. Introducing the total as a redundant state clarified this constraint without adding an independent degree of freedom. These findings distinguish limitations in full-state reachability from the possibility of influencing clinically relevant outputs, such as tumour burden or cytokine concentration.

5.3 Sensitivity

The sensitivity analyses complements the practical identifiability assessment by indicating how strongly, and at which times, parameter variations affect the measured outputs. Weak or temporally localized responses may restrict the information available for estimation, whereas pronounced responses can provide informative observations when the relevant intervals are sampled. However, sensitivity magnitude alone does not establish practical identifiability, since parameters with similar effects may compensate for one another. Moreover, the patterns depend on the explored parameter ranges and experimental conditions.

For HCARTr, the pronounced sensitivities to ϕ and ρ coexist with their limited independent determination, illustrating the distinction between influencing an output and being identifiable from it. The weak responses to μ and particularly b are consistent with limited information about these parameters in tumour-burden observations. Although sensitivity magnitudes generally increased towards the end of the simulation, this does not by itself establish that later measurements are more informative, since their value also depends on measurement uncertainty and the ability to distinguish parameter effects.

For CRSr, the sustained responses to σ_I and δ_I contrasted with the earlier responses to parameters such as α and β_B and the late, localized response to γ. These differences support considering complementary sampling times rather than a single observation window. Notably, the very small sensitivities to θ, μ_P, and K caution against interpreting their comparatively narrow RIdent distributions as evidence of robust recovery. The zero sensitivity to g_0 further confirms the absence of information about this parameter under the initial conditions considered. Overall, these analyses help guide measurement selection and calibration, while neither biological importance nor the suitability of fixing a parameter can be inferred from sensitivity magnitude alone.

5.4 Outlook

The main contribution of this work is the joint discussion and interpretation of key identification aspects of CAR-T cell therapy models, including structural properties, practical parameter recovery, and parametric sensitivity. This approach distinguishes intrinsic model limitations from restrictions associated with particular observations and experimental conditions. The findings motivate further investigation of identifiable parameter combinations, complementary measurements, and informative sampling schedules, followed by evaluation against experimental data. When tailored to patient-specific data and protocols, our results would provide a stronger basis for subsequent model-based treatment optimization and control.

Data and code availability

We provide scripts of the analyses described in this paper, as well as the synthetically generated data, in a Zenodo repository that can be accessed at: `https://doi.org/10.5281/zenodo.24194637`.

Funding

This work was supported by grant PID2023-146275NB-C21 funded by MICIU/AEI/10.13039/501100011033 and ERDF/EU.

Author contributions

Conceptualization, AFV; methodology, AFV, AGV; investigation, AGV; writing—original draft preparation, AFV, AGV; supervision, AFV; funding acquisition, AFV. All authors have read and agreed to the published version of the manuscript.

Conflict of interest

The authors declare no conflicts of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

References

[1] Alexander RA Anderson and Philip K Maini. Mathematical oncology. *Bulletin of mathematical biology*, 80(5):945–953, 2018.

[2] Joshua Adam Bull and Helen Mary Byrne. The hallmarks of mathematical oncology. *Proceedings of the IEEE*, 110(5):523–540, 2022.

[3] Angela M Jarrett, Danial Faghihi, David A Hormuth, Ernesto ABF Lima, John Virostko, George Biros, Debra Patt, and Thomas E Yankeelov. Optimal control theory for personalized therapeutic regimens in oncology: Background, history, challenges, and opportunities. *Journal of clinical medicine*, 9(5):1314, 2020.

[4] Lisette G de Pillis, Weiqing Gu, and Ami E Radunskaya. Mixed immunotherapy and chemotherapy of tumors: modeling, applications and biological interpretations. *Journal of theoretical biology*, 238(4):841–862, 2006.

[5] Carl H June and Michel Sadelain. Chimeric antigen receptor therapy. *New England Journal of Medicine*, 379(1):64–73, 2018.

[6] Shannon L Maude, Theodore W Laetsch, Jochen Buechner, Susana Rives, Michael Boyer, Henrique Bittencourt, Peter Bader, Michael R Verneris, Heather E Stefanski, Gary D Myers, et al. Tisagenlecleucel in children and young adults with b-cell lymphoblastic leukemia. *New England Journal of Medicine*, 378(5):439–448, 2018.

[7] Lunan Liu, Chao Ma, Zhuoyu Zhang, Matthew T Witkowski, Iannis Aifantis, Saba Ghassemi, and Weiqiang Chen. Computational model of car t-cell immunotherapy dissects and predicts leukemia patient responses at remission, resistance, and relapse. *Journal for immunotherapy of cancer*, 10(12):e005360, 2022.

[8] Ujwani Nukala, Marisabel Rodriguez Messan, Osman N Yogurtcu, Xiaofei Wang, and Hong Yang. A systematic review of the efforts and hindrances of modeling and simulation of car t-cell therapy. *The AAPS journal*, 23:1–20, 2021.

[9] Sattva S Neelapu. How i manage: pathophysiology and management of toxicity of chimeric antigen receptor t-cell therapies. *Journal of Clinical Oncology*, 39(5):456–466, 2021.

[10] Theodoros Giavridis, Sjoukje JC van der Stegen, Justin Eyquem, Mohamad Hamieh, Alessandra Piersigilli, and Michel Sadelain. Car t cell–induced cytokine release syndrome is mediated by macrophages and abated by il-1 blockade. *Nature medicine*, 24(6):731–738, 2018.

[11] Andreas Raue, Marcel Schilling, Julie Bachmann, Andrew Matteson, Max Schelke, Daniel Kaschek, Sabine Hug, Clemens Kreutz, Brian D Harms, Fabian J Theis, et al. Lessons learned from quantitative dynamical modeling in systems biology. *PloS one*, 8(9):e74335, 2013.

[12] Franz-Georg Wieland, Adrian L Hauber, Marcus Rosenblatt, Christian T"onsing, and Jens Timmer. On structural and practical identifiability. *Current opinion in systems biology*, 25:60–69, 2021.

[13] Alejandro F Villaverde. Observability and structural identifiability of nonlinear biological systems. *Complexity*, 2019(1):8497093, 2019.

[14] Nicholas N Lam, Paul D Docherty, and Rua Murray. Practical identifiability of parametrised models: A review of benefits and limitations of various approaches. *Mathematics and Computers in Simulation*, 199:202–216, 2022.

[15] Sandra Díaz-Seoane, Xabier Rey Barreiro, and Alejandro F Villaverde. Strike-goldd 4.0: user-friendly, efficient analysis of structural identifiability and observability. *Bioinformatics*, 39(1):btac748, 2023.

[16] Sandra Díaz-Seoane, Antonio Barreiro Blas, and Alejandro F Villaverde. Controllability and accessibility analysis of nonlinear biosystems. *Computer methods and programs in biomedicine*, 242:107837, 2023.

[17] Eva Balsa-Canto, David Henriques, Attila Gábor, and Julio R Banga. Amigo2, a toolbox for dynamic modeling, optimization and control in systems biology. *Bioinformatics*, 32(21):3357–3359, 2016.

[18] Rachel Mester, Alfonso Landeros, Chris Rackauckas, and Kenneth Lange. Differential methods for assessing sensitivity in biological models. *PLoS computational biology*, 18(6):e1009598, 2022.

[19] L Gerard Van Willigenburg, Johannes D Stigter, and Jaap Molenaar. Establishing local strong accessibility of large-scale nonlinear systems by replacing the lie algebraic rank condition. In *2021 European Control Conference (ECC)*, pages 2645–2650. IEEE, 2021.

[20] Luciana RC Barros, Emanuelle A Paixão, Andrea MP Valli, Gustavo T Naozuka, Artur C Fassoni, and Regina C Almeida. Cart math—a mathematical model of car-t immunotherapy in preclinical studies of hematological cancers. *Cancers*, 13(12):2941, 2021.

[21] Daniela S Santurio, Luciana RC Barros, Ingmar Glauche, and Artur C Fassoni. Mathematical modeling unveils the timeline of car-t cell therapy and macrophage-mediated cytokine release syndrome. *PLOS Computational Biology*, 21(4):e1012908, 2025.

[22] Oana-Teodora Chis, Julio R Banga, and Eva Balsa-Canto. Structural identifiability of systems biology models: a critical comparison of methods. *PloS one*, 6(11):e27755, 2011.

[23] Xabier Rey Barreiro, Alejandro Fernandez Villaverde, et al. On the origins and rarity of locally but not globally identifiable parameters in biological modeling. *IEEE Access*, 2023.

[24] E Tunali and Tzyh-Jong Tarn. New results for identifiability of nonlinear systems. *IEEE Transactions on Automatic Control*, 32(2):146–154, 1987.

[25] Alejandro F Villaverde, Antonio Barreiro, and Antonis Papachristodoulou. Structural identifiability of dynamic systems biology models. *PLoS computational biology*, 12(10):e1005153, 2016.

[26] Milind Joshi, Andreas Seidel-Morgenstern, and Andreas Kremling. Exploiting the bootstrap method for quantifying parameter confidence intervals in dynamical systems. *Metabolic engineering*, 8(5):447–455, 2006.

[27] Robert Hermann and Arthur Krener. Nonlinear controllability and observability. *IEEE Transactions on automatic control*, 22(5):728–740, 1977.

[28] Eduardo D Sontag et al. *Mathematical control theory*, volume 2. Springer-Verlag New York, 1998.

[29] Mathukumalli Vidyasagar. *Nonlinear systems analysis*. SIAM, 2002.

[30] JD Stigter and D Joubert. Computing measures of identifiability, observability, and controllability for a dynamic system model with the strucid app. *IFAC-PapersOnLine*, 54(7):138–143, 2021.

Supplementary Information

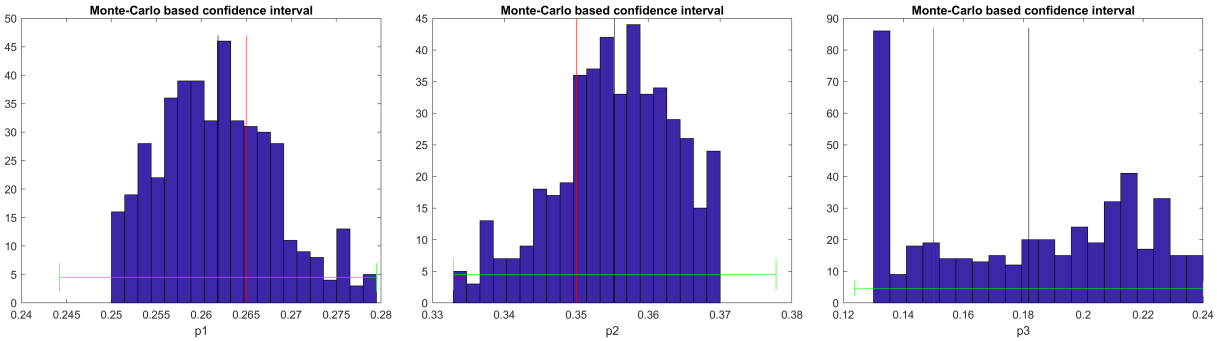

(a) ϕ (p_1).

(b) ρ (p_2).

(c) ϵ (p_3).

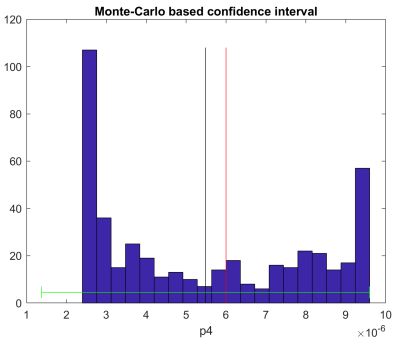

(d) θ (p_4).

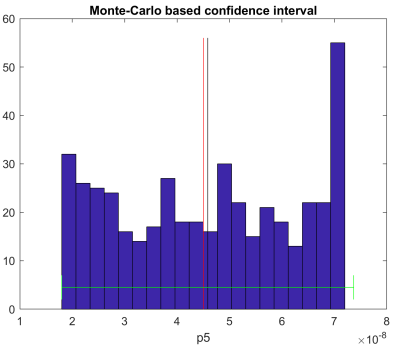

(e) α (p_5).

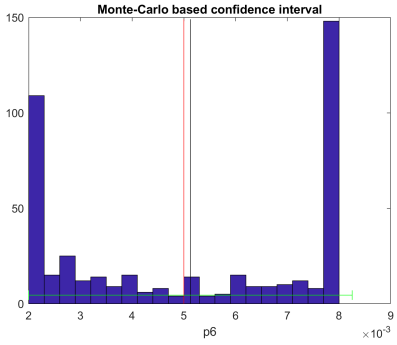

(f) μ (p_6).

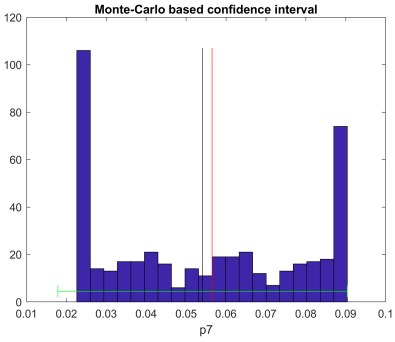

(g) r (p_7).

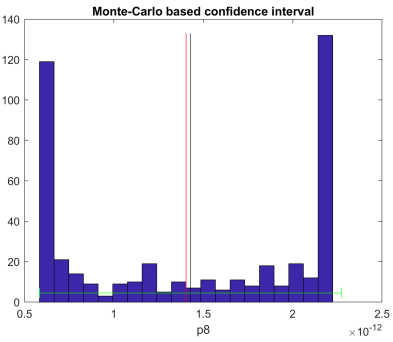

(h) b (p_8).

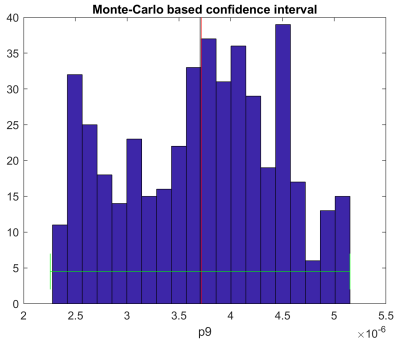

(i) γ (p_9).

Figure S1: Distributions of HCARTr parameter estimates obtained from 500 joint Monte Carlo estimation trials. Horizontal axes show parameter values on their original scales; bar heights indicate the number of estimates in each bin. Each panel uses its own axis scales.

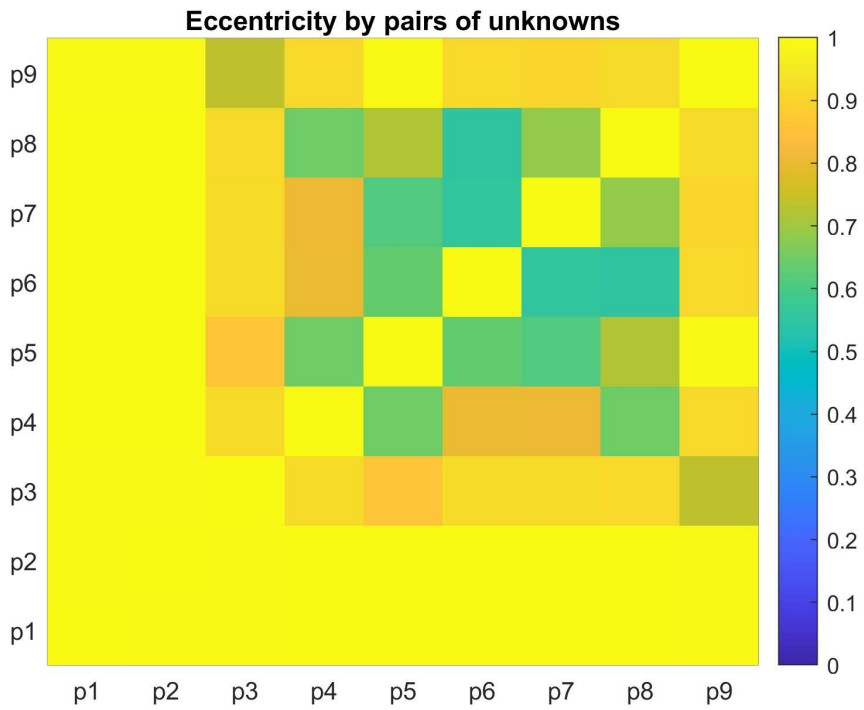

Figure S2: Pairwise eccentricity matrix reported by AMIGO_RIdent for the HCARTr parameter estimates. Parameter indices correspond to $(p_1, \ldots, p_9) = (\phi, \rho, \epsilon, \theta, \alpha, \mu, r, b, \gamma)$.

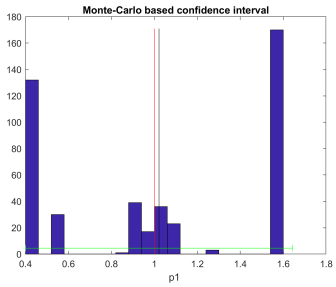

(a) η

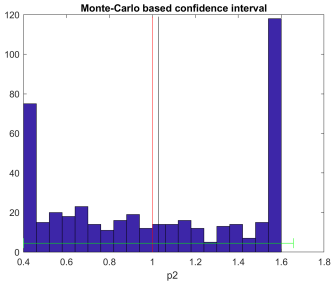

(b) μ_I

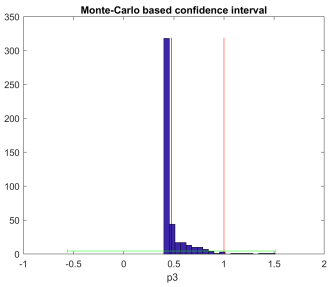

(c) ν

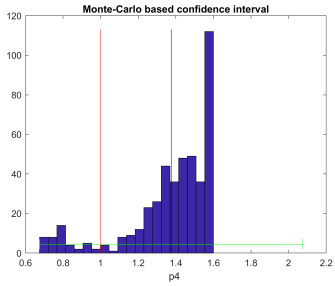

(d) κ

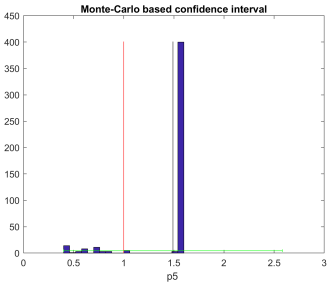

(e) ϵ

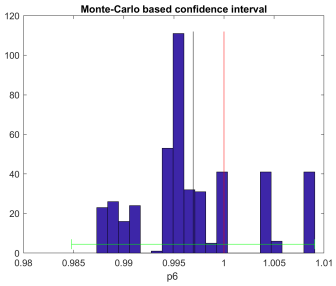

(f) θ

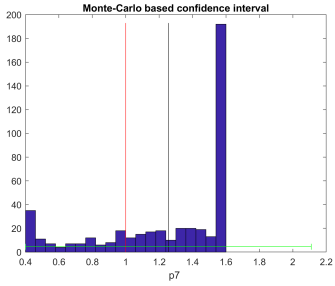

(g) μ_E

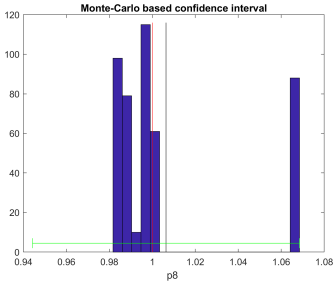

(h) μ_P

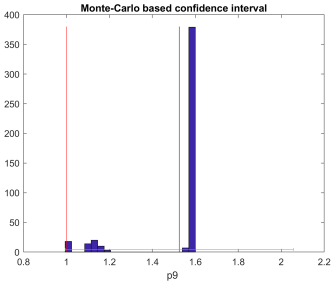

(i) ρ

Figure S3: Distributions of CRSr parameter estimates obtained from 500 joint Monte Carlo estimation trials, for parameters 1–9. Horizontal axes show estimates normalized by their nominal values; bar heights indicate the number of estimates in each bin. Each panel uses its own axis scales.

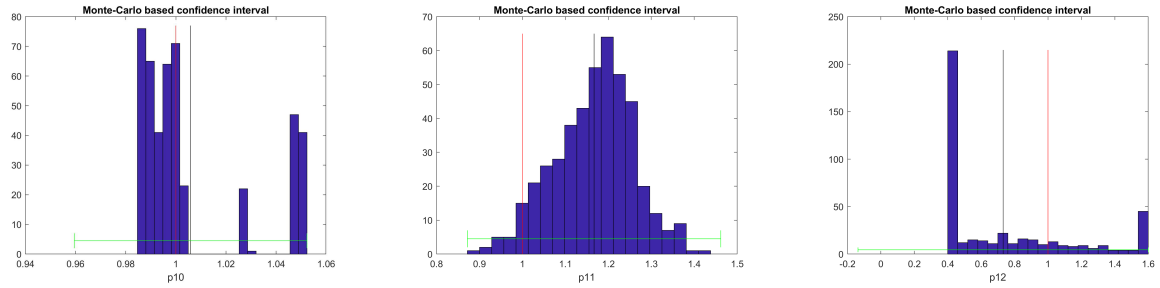

(a) K (b) γ (c) A

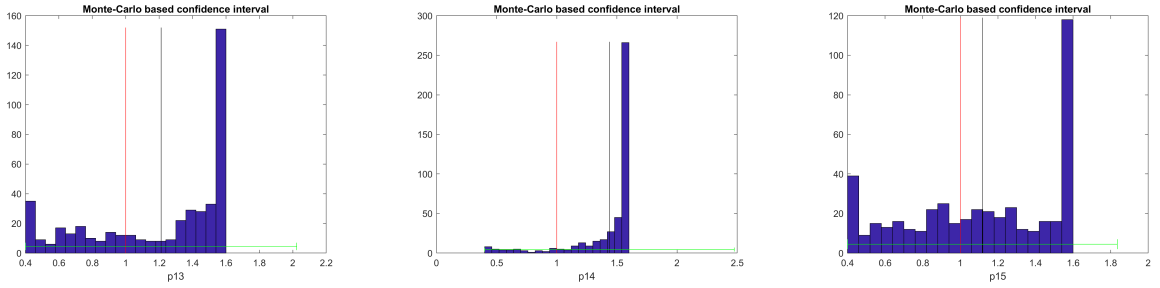

(d) B (e) C (f) σ_M

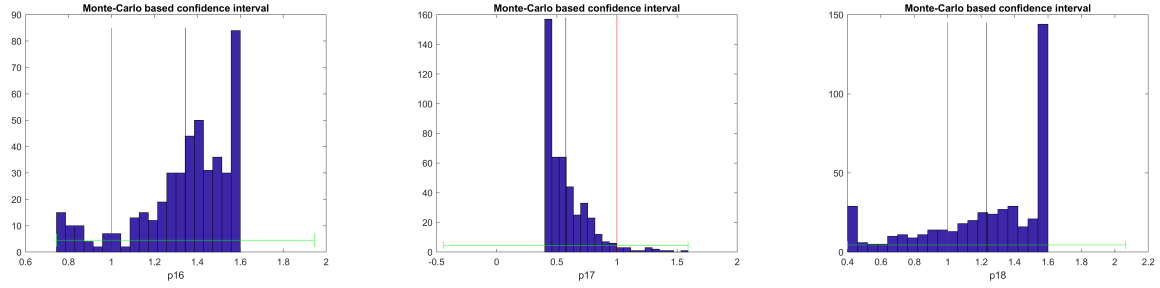

(g) δ_M (h) σ_I (i) α

Figure S4: Distributions of CRSr parameter estimates obtained from 500 joint Monte Carlo estimation trials, for parameters 10–18. Horizontal axes show estimates normalized by their nominal values; bar heights indicate the number of estimates in each bin. Each panel uses its own axis scales.

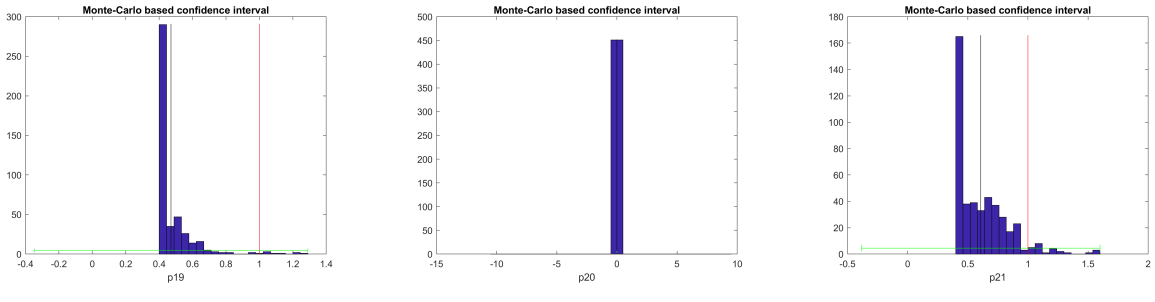

(a) δ_I (b) g_0 (c) β_B

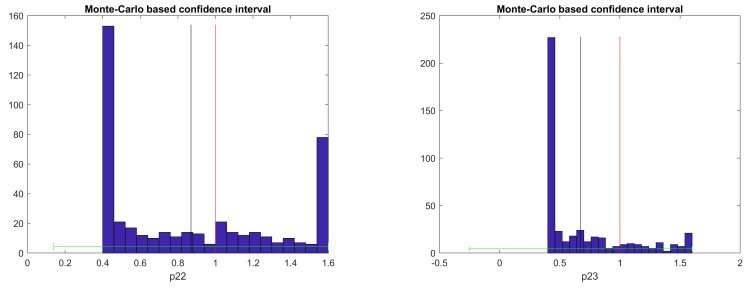

(d) β_K (e) β_C

Figure S5: Distributions of CRSr parameter estimates obtained from 500 joint Monte Carlo estimation trials, for parameters 19–23. Horizontal axes show estimates normalized by their nominal values, except for g_0, which is shown on its original scale because its nominal value is zero. Bar heights indicate the number of estimates in each bin. Each panel uses its own axis scales.

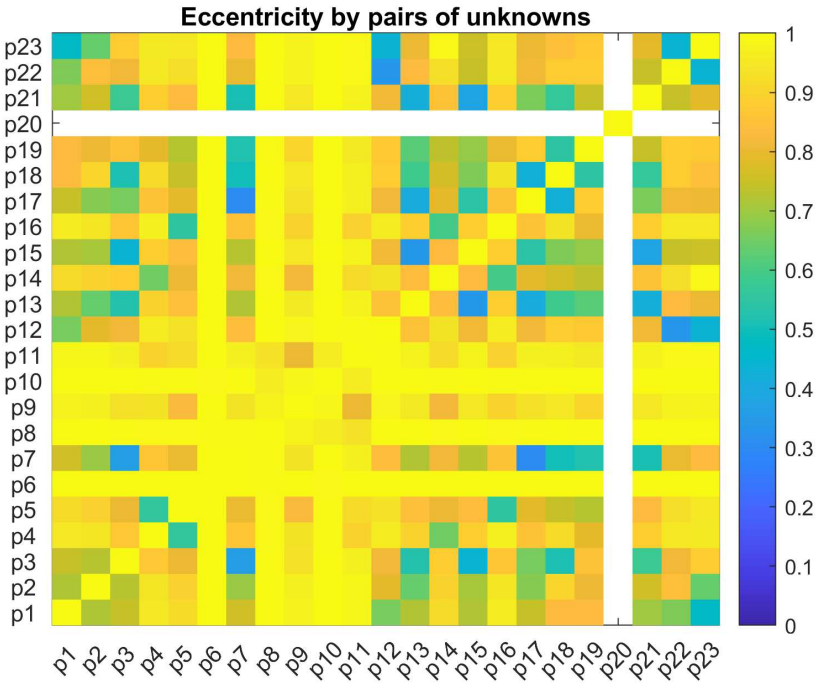

Figure S6: Pairwise eccentricities computed by `AMIGO_RIdent` for the CRSr parameter estimates. Parameter indices follow the order shown in Figures S3–S5.

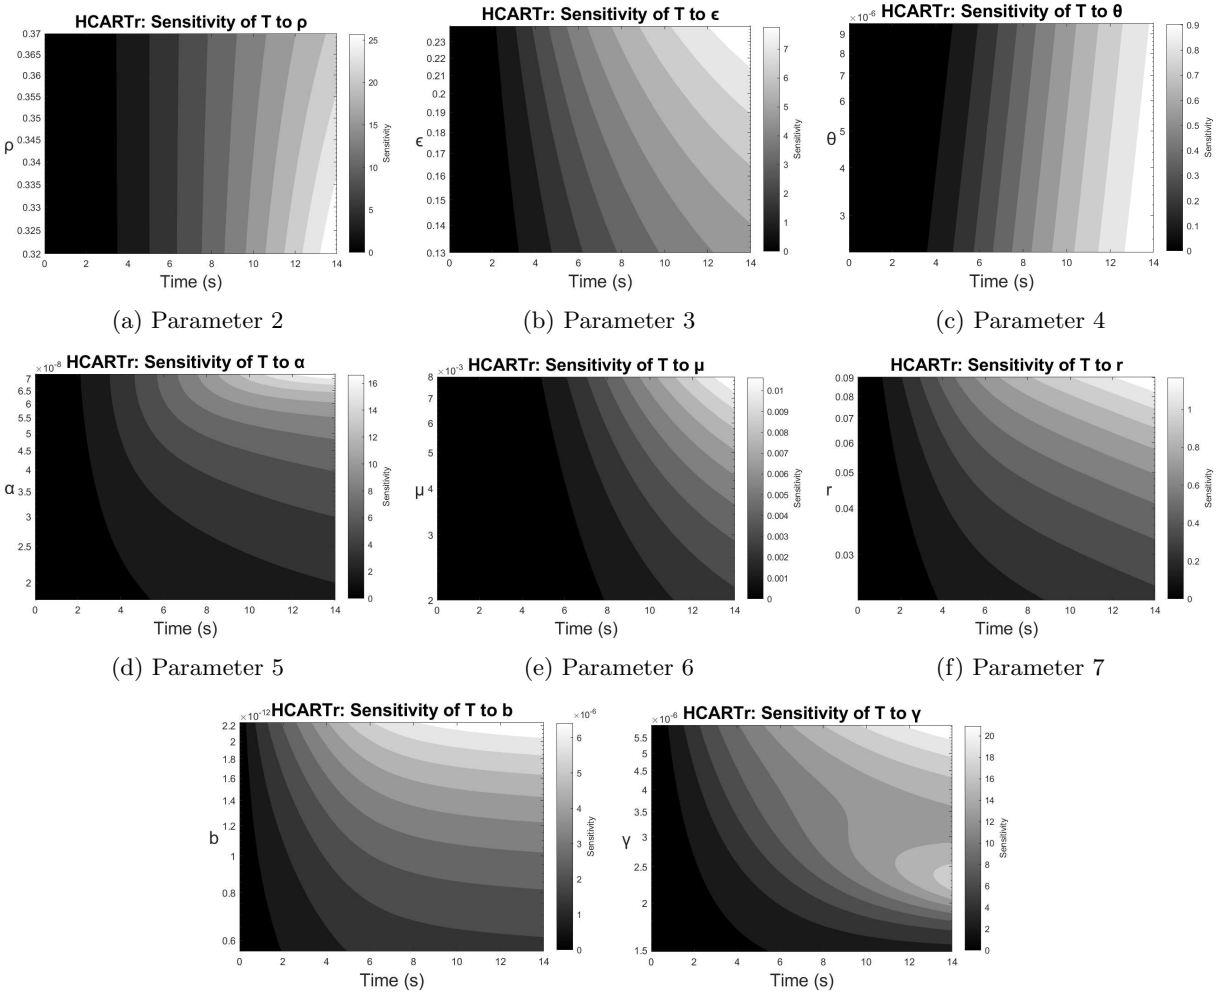

(a) Parameter 2 (b) Parameter 3 (c) Parameter 4

(d) Parameter 5 (e) Parameter 6 (f) Parameter 7

(g) Parameter 8 (h) Parameter 9

Figure S7: Sensitivity of tumour burden T (state 3) over time for parameter indices 2–9.

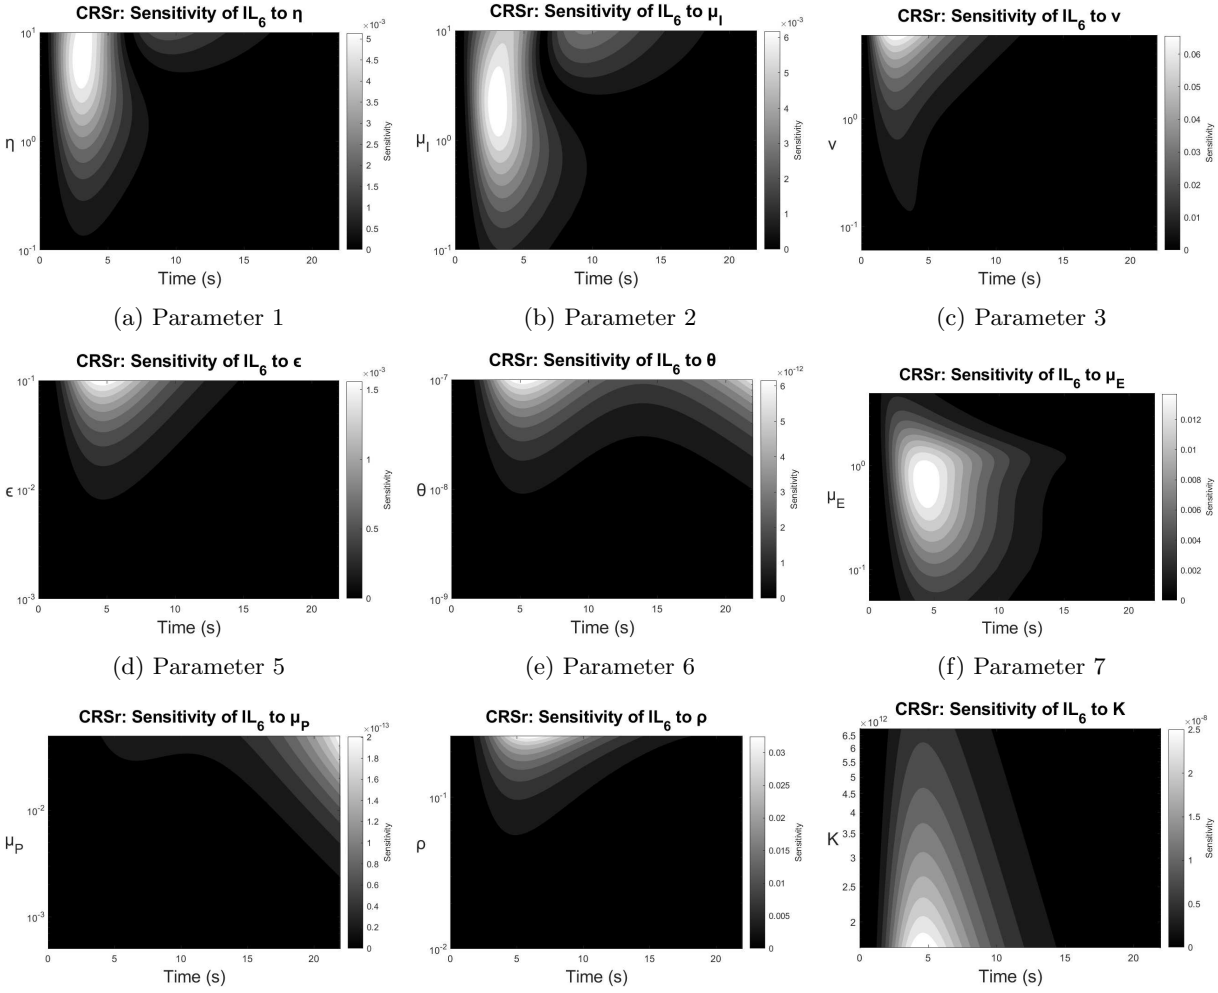

(a) Parameter 1 (b) Parameter 2 (c) Parameter 3

(d) Parameter 5 (e) Parameter 6 (f) Parameter 7

(g) Parameter 8 (h) Parameter 9 (i) Parameter 10

Figure S8: Sensitivity of IL_6 concentration (state 8) over time for parameter indices 1–3 and 5–10.

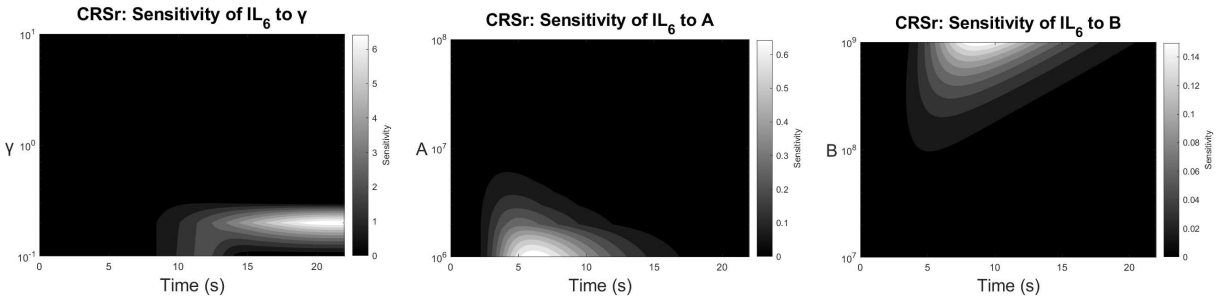

(a) Parameter 11

(b) Parameter 12

(c) Parameter 13

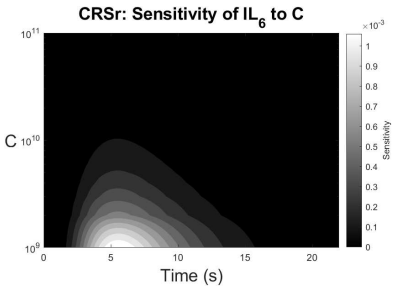

(d) Parameter 14

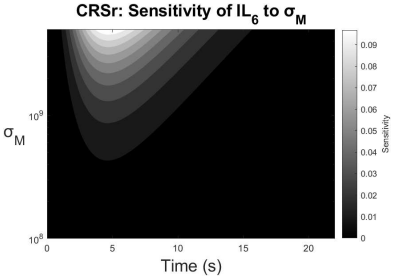

(e) Parameter 15

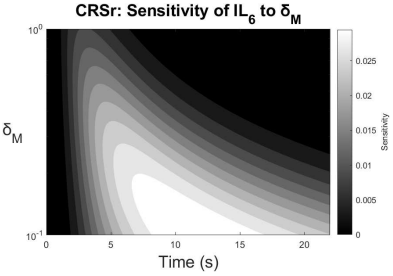

(f) Parameter 16

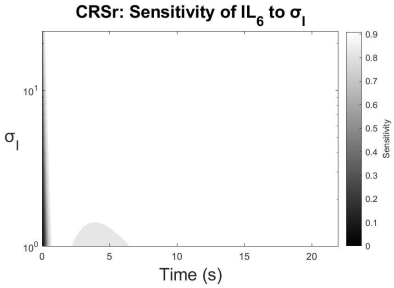

(g) Parameter 17

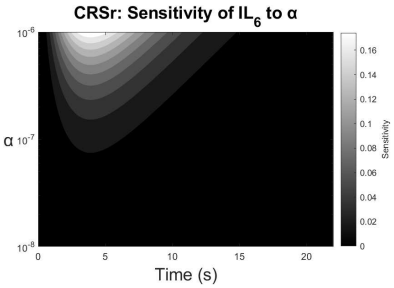

(h) Parameter 18

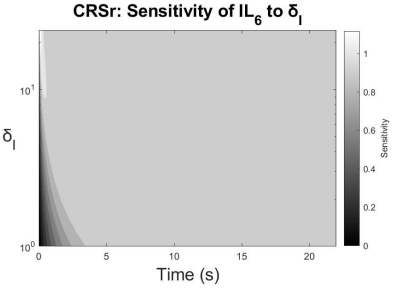

(i) Parameter 19

Figure S9: Sensitivity of IL_6 concentration (state 8) over time for parameter indices 11–19.

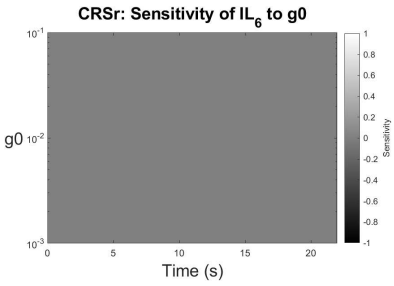

(a) Parameter 20

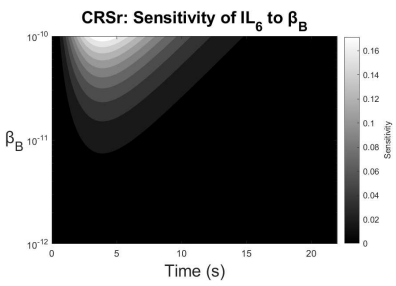

(b) Parameter 21

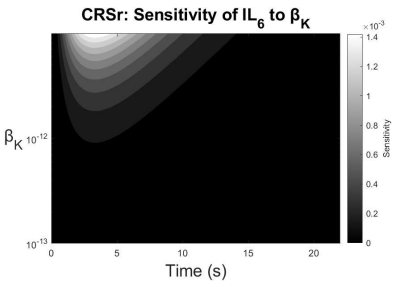

(c) Parameter 22

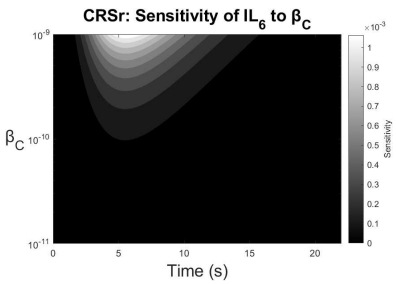

(d) Parameter 23

Figure S10: Sensitivity of IL_6 concentration (state 8) over time for parameter indices 20–23.