Evolutionary Dynamics of CAR T Cell Therapy

Gregory J. Kimmel, Frederick L. Locke[§], Philipp M. Altrock*[§]

H. Lee Moffitt Cancer Center and Research Institute, Tampa, FL 33612, USA.

[§]These authors contributed equally

*Corresponding authors: frederick.locke@moffitt.org, philipp.altrock@moffitt.org
H. Lee Moffitt Cancer Center, and Research Institute
12902 USF Magnolia Drive, Tampa, FL 33612, USA

ABSTRACT

Chimeric antigen receptor (CAR) T cell therapy is a remarkably effective immunotherapy that relies on *in vivo* expansion of engineered CAR T cells (after lymphodepletion by chemotherapy) to overcome the patient's tumor burden. The quantitative laws underlying this expansion and subsequent eradication of tumors remain unknown. Still, about 60% of CAR T-treated patients are likely to progress; their tumors are not eradicated. There is a critical need to understand and disentangle the multiple processes that contribute to CAR expansion and tumor eradication. Our goal is a quantitative understanding and ranking of the sources of variability that determine tumor response and the ultimate success of CAR T cell therapy. We here develop a mathematical model of T cell-tumor cell co-evolution, and demonstrate that CAR T cell expansion is shaped by immune reconstitution, antigen load, and cytokine signaling after lymphodepletion. We find that success of therapy is a stochastic event. Even if durable response is highly likely, its timing can be highly variable. As a consequence, we predict that clinical interventions to decrease tumor mass before CAR T could improve likelihood and timing of durable response. Changing the composition of the CAR population toward more memory cells could also improve survival, but we predict that there are diminishing returns for such efforts. All together, we provide a cell population-based approach to model and understand CAR T cell therapy, which can be leveraged to suggest new dosing strategies, and assess the need for multiple doses and their outcomes.

INTRODUCTION

Relapsed and refractory Large B cell Lymphoma (LBCL) is the most common subtype of Non-Hodgkin Lymphoma, the most common hematologic malignancy in the US with 72,000 new cases (4.3% of all cancer) and 20,000 deaths (3.4% of all cancer deaths) in 2017[1]. In LBCL patients that do not respond to chemotherapy, the median overall survival is <7 months[2]. These patients could benefit from autologous chimeric antigen receptor (CAR) T cell therapy, consisting of genetically engineered T cells re-targeted to CD19 a protein specific to B lineage cells from which LBCL arises[3]. ZUMA-1 was a pivotal, multicenter, phase 1-2 trial of axicabtagene ciloleucel (axi-cel), a second generation CD19 directed CAR T cell therapy with CD28 co-stimulation, for chemorefractory LCBL (n=101 patients treated)[4-6]. Overall response rate (ORR) and Complete Response (CR) rate were 82% and 54% compared to robust historical controls suggesting that the respective responses would have been 26% and 7% with standard chemotherapy[7]. With 29 months median follow-up on ZUMA-1, responses are ongoing in 39%, with a plateau in the duration of the response curve after 6 months. Similar results were demonstrated with other CD19 CAR T cell therapies for LBCL[8,9]. The US FDA approved axi-cel for relapsed/refractory LBCL. Response to CAR T may be delayed, but patients not achieving CR or partial response (PR) by 90 days fare poorly, with over 60% of LBCL patients eventually progressing. Clinically relevant quantitative models of CAR T dynamics and response are needed to predict patient outcomes and the impact of combinatorial therapies with CAR T.

Tumor response and survival in CAR T cell therapy are linked to maximum CAR T cell density in peripheral blood after treatment[4]. Furthermore, *in vivo* CAR T cell expansion, tumor response, and toxicity are impacted by CAR T differentiation states[10], T cell-competition dynamics[11], and homeostatic signals[12] during homeostatic repopulation of normal and CAR T cells after lympho-depleting chemotherapy and CAR T infusion. The need for improved understanding of the evolutionary dynamics of lymphocytes and tumor cells within patients during CAR T cell therapy could be leveraged to explore new dosing and combination therapies.

T cell proliferation, differentiation (from naïve to memory to effector cells), and apoptosis (following activation, cell killing, or aging) are the key mechanisms that influence overall cell population dynamics. The rates of these mechanisms are influenced by three qualitatively different processes. First, exogenously secreted cytokines support memory cell proliferation and homeostasis (e.g.IL-7, IL-15), especially after CAR T cell therapy[12]. Second, T cells secrete cytokine signals in an autocrine fashion (e.g. IL-2) regulating their own proliferation in a density-dependent way[13], and stimulating or inhibiting differentiation [14]. Third, the amount of cognate

antigen encountered, including the CD19+ tumor in the context of CD19 CAR T, plays a dynamic role in T cell turnover because it directly influences T cell differentiation into effector cells, and thus indirectly regulates T cell life span[15]. A better understanding of CAR T cell therapy requires modeling these complex processes using mathematical principles[16].

When cell population densities are influenced both by intrinsic (cell-autonomous) and extrinsic (non-cell-autonomous) forces, eco-evolutionary dynamics provides an appropriate mathematical modeling framework[17]. Ecological dynamics describe density dependent effects, e.g. how population size impacts growth rate[18], or how growth is influenced by age- or differentiation structure[16,19,20]. Evolutionary dynamics describe how relative population abundances (frequencies) influence and change the per capita growth rates or the overall fitness landscape, and the population's ability to adapt due to mutation or phenotypic switching[21]. Here, we combine such population dynamics mechanisms to better understand CAR T therapy and to predict novel outcomes under alternative treatment scenarios. We integrate normal and CAR T cell population densities and differentiation states (memory to effector differentiation), and tumor burden over time, in order to discover the principles that drive durable response of CAR T cell therapy.

METHODS

We developed a mathematical modeling framework that describes normal T cells, CAR T cells, and tumor cells. The model considers four cell populations in the form of four continuous-time birth and death stochastic processes and their deterministic mean-field equations. The model considers normal memory T cells N, memory CAR T cells M, effector CAR T cells E, and antigen-presenting tumor cells B. In principle, all populations can be arbitrarily large, but we assume that memory cells grow toward a homeostatic maximum, reflected by the carrying capacities K_N and K_M. We considered the case where a patient received lymphodepleting chemotherapy prior to autologous injection of CAR T cells. We set time to 0 at the time of CAR injection. Normal and CAR memory T cells then grow toward their respective carrying capacities and potentially influence each other dynamically. This mutual influence gives rise to co-evolution at the level of memory T cells; normal memory cells expand at the rate r_N*log[K_N/(N+M)], CAR memory cells expand at the rate r_M*log[K_M/(N+M)]. Further, memory CAR T cells differentiate into effector CAR T cells, which occurs at a rate r_E(B) and is dependent on the amount of antigen (CD19) presented by the tumor. While we assume that memory CAR T cells could persist indefinitely, effector CAR T cells have a maximum life-span determined by their intrinsic

death rate d_E, and the fact that tumor killing upon contact with B cells also leads to exhaustion and effector cell death at rate γ_E, and is proportional to tumor mass B. The tumor cell population B grows autonomously at the net growth rate r_B, and experiences tumor killing at rate γ_B, proportional to the number of effector cells E. This dynamical system can be written in a mean-field limit as a set of four deterministic ordinary differential equations

$$\frac{dN}{dt} = -r_N \log\left[\frac{N+M}{K_N}\right] \tag{1}$$

$$\frac{dM}{dt} = -r_M \log\left[\frac{N+M}{K_M}\right] \tag{2}$$

$$\frac{dE}{dt} = r_E(B)M - \gamma_E\,B\,E - d_E E \tag{3}$$

$$\frac{dB}{dt} = r_B B - \gamma_B\,B\,E \tag{4}$$

For a derivation of these equations and their associated complex stochastic birth and death process see the Supplement. This system was used to find suitable parameter value ranges using clinical data from multiple sources (see Results and Supplement), and to make numerical and simulation predictions. Note that Equations (1) and (2) decouple from equations (3) and (4), but not vice versa.

Clinical data. We used data from LBCL patients treated on the multicenter ZUMA-1 trial[6], or with standard of care axi-cel at Moffitt Cancer Center. The ZUMA-1 trial reported quartile CAR T cell levels in peripheral blood over five consecutive time points (days 7, 14, 28, 90, 180), which we used to parameterize the model (see Supplement). We also used normal complete blood count-derived lymphocyte counts (assumed to be T cells) of patients receiving CAR T cell therapy (days 5, 7, 14, 28, 90, 180). All peripheral measurements of cells/μL were converted to total T cell counts assuming 5L of blood and multiplying by 10^8, which results from the assumption that 1% of cells are in periphery at all times and from conversion of μL to L. The normal T cell carrying capacity, K_N = 500 cells/μL was estimated from data by Turtle *et al.*[22]. Median tumor volume at day 0 was assumed to be B_0=5.269 x10^{10} cells as per our own estimates from patients treated at Moffitt. We defined the time of cure as $B(t_{cure}) = 0$, and defined progression when $B(t_{progression}) \geq 5B_0$. Patients not in either of these categories after 1000 days are placed in a third category, which we called undetermined. No simulations of patients led to patients in the undetermined category.

Model fitting. We obtained median parameter values and parameter value distributions using a nonlinear least squares regression function with the variance as weights. This optimization problem is solved using the BFGS algorithm from the optimization package in Julia, which is wrapped around solving the differential equation[23,24]. For patient data that was only available in quartiles, we assumed an approximately normal distribution and estimated the variance from the difference between the quantiles. We generated 243 patient trajectories by utilizing all combinations of the five time points and three available quantiles. We fit parameters in equations (1)-(2) first as they are uncoupled from equations (3)-(4). This created a range of parameter values and we used the median values to fit equations (3)-(4).

Model analysis and simulations. We analyzed the deterministic mean field model using linear stability analysis and showed that in the mean-field limit our parameters lead to eventual cancer progression. However, it was clear that many of the trajectories spent large amounts of time near the tumor-free ($B = 0$) state. It is well known that small population dynamics cannot be adequately captured with a deterministic model (see Supplement). In contrast, a full stochastic simulation with a system comprised of billions of cells would take days or weeks to run a single patient. This is computationally prohibitive as we desire many patient trajectories to gather statistics. To alleviate this issue, while maintaining the accuracy of the stochastic model, we developed a hybrid model, which exploited the fact that fluctuations should only be relevant when a population went below a threshold. The model operates in the deterministic limit until a population goes below the threshold of 100 cells, and from then on update that population stochastically, while the remaining populations were updated deterministically. In principle, the CAR and tumor populations should both be allowed to switch since their respective populations can become small. However, we are only interested in whether the tumor is eradicated and so our model only switches the tumor to and from the stochastic model, all other populations remain in the deterministic setting. We ran the deterministic system using the built-in solver Rosenbrock23 in the *DifferentialEquation* package of the software Julia, and simulated the stochastic master equation using a Gillespie algorithm[25]. For a typical simulation result with stochastic outcome, we simulated 1000 virtual patients, each drawn with their own, slightly different set of parameters. These parameter variations were achieved using a normal distribution with mean chosen as the median parameter value estimated from model fitting, and variance resulting as a small fraction of that mean (typically between 5-15%). A patient was counted as cured when their tumor reached 0 cells, which is an absorbing state of the stochastic birth and death process. A patient was counted as progressed when their tumor mass hit the

threshold of 5 times their mass prior to treatment. These two endpoints led to Kaplan-Meier curves counting progression free survival (PFS) as a function of time post CAR T cell injection.

RESULTS

The patient dynamics after lymphodepletion and CAR T cell infusion can be categorized into three characteristic scenarios. All patients receive lymphodepleting chemotherapy followed by autologous CAR T cell infusion. In the first scenario a patient's tumor shows no response (Fig. 1 A). In the second scenario, a patient shows a transient response after CAR infusion, characterized by a rapid decline in tumor burden during CAR T cell expansion, followed later by a growth of the tumor mass after the CAR T cell population has decreased (Fig. 1 B). In the third scenario, a patient has complete disappearance of tumor and has a long-term response, which most likely is a result of tumor eradication (Fig. 1 C). These three fundamentally different outcomes each emerge as a solution under a specific set of parameter choices in our mathematical model.

Figure 1: Dynamics of clinical states, CAR density over time, and rate of cure. A: Schematic of no response–the treatment is ineffective and the tumor grows. **B**: Schematic of transient response–the treatment initially shrinks the tumor, but is unable to eradicate it leading to eventual relapse. **C**: Schematic of durable response–the tumor is undetectable 180 days post infusion and has likely been eliminated. **D**: Distribution of total CAR T cells in periphery over time (assuming a normal distribution). **E**: Probability of durable response (tumor eradication) could be captured by a probability of tumor extinction conditioned on presence of CAR T cells.

The degree of CAR T cell expansion, as measured in the peripheral blood, has been linked to response and durable response to CAR T cell therapy in LBCL[4]. A key component of this model is a mechanistic description of the rate of expansion and duration of persistence of the CAR T cell population, which then could be responsible for driving patient outcomes. The dynamics of normal T cell density determines characteristics of the CAR T cell population expansion, and treatment success. To this end, we built a mathematical model that incorporates normal and CAR T cell dynamics and explains observed tumor dynamics during CAR T cell therapy. Data from the ZUMA-1 trial demonstrate[5,6] that the overall CAR T cell density can range over a wide interval at various clinically important time points, as measured in periphery at days 7, 14, 28, 90 and 180 after infusion (Fig. 1 D). We hypothesized that malignant B cell extinction is a stochastic event. This means that, even if durably response is highly likely (and progression highly unlikely), the times to these events follows a potentially broad distribution with non-vanishing variance and skewness. This stochastic process is governed by individual birth and death events in critically small cell populations.

Following CAR T therapy, tumor cells over time can be modeled by a birth and death process. Birth events are driven by cell intrinsic proliferation, whereas death events are driven by intrinsic events (apoptosis) and extrinsic predation proportional to the number of effector CAR T cells present. Then, the probability of extinction can be calculated as a function of time (see Supplement). Next, as a first approximation, one can quantify the overall success probability of CAR T cell therapy as a tumor extinction event conditioned upon survival of a sufficient number of effector CAR T cells. If these two events are assumed to be independent, the likelihood of cure is the product of tumor extinction probability and CAR T cell survival probability, which can be seen to be a monomodal function of time since CAR infusion (Fig. 1 E). In general, one might want to relax these strong assumptions. To this end, we developed and parameterized a more comprehensive dynamical system that models the dynamics of normal memory T cells, memory CAR T cells that can persist or differentiate, effector CAR T cells that can kill tumor cells, and tumor cells that express the CAR target CD19 (Fig. 2 A). This system represents the minimally complex interaction patterns among these four cell populations, modeled as a deterministic mean field model when cell densities are large (e.g. during expansion), or as a stochastic individual based model, when cell densities become critically low and random events dominate the time scales (see Methods).

Figure 2: Modeling the success of CAR T cell therapy using memory T cell interactions and an individual based complex/stochastic dynamical system. A: Model schematic of eco-evolutionary dynamics and corresponding cell-level birth/death events. **B**: The normal T cell compartment showing the return to homeostatic levels. The points are the median patient values. **C**: The model and patient quartiles of the total CAR in the body with the assumption that only 1% is present in periphery (error bars show the quartiles). **D**: Two trajectories of tumor burden over time. Both enter the stochastic region (<100 tumor cells), but only one is cured, the other simulated patient's tumor spends time in the region before eventually growing which leads to progression. **E**: The impact of the initial amount of CAR memory T cells present at infusion on cure rate. We observe diminishing returns for highly specialized bag containing only memory cells. **F**: Conditioning chemotherapy is vital to effective treatment. If lymphodepletion knocks out less than 95% of the resident T cell population, treatment outcome is poor.

In order to accurately describe CAR T cell expansion and tumor killing, it is necessary to describe CAR T cells as a heterogeneous population. Within the ZUMA-1 trial, the CAR T cell infusion product was described as containing naïve and central memory, as well as effector memory and effector cells[26]. To model this heterogeneity, we binned these four T cell-states into two CAR T cell compartments of memory cells (CCR7+, naïve and central memory), and effector cells (CCR7-, effector memory and effector) (Fig. 2 A). Memory CAR T cells divide

symmetrically into memory cells while competing with normal memory T cells, and divide asymmetrically to differentiate into tumor-killing effector cells. This differentiation step can depend on the number of target cells, or antigen load, specific to the CAR T cell population[27,28], in our case CD19. Effector cells can autonomously die (apoptosis), and interact with the tumor by killing antigen-presenting B and lymphoma cells. The process of cell killing is assumed to also reduce an effector CAR T cell's life-span[29]. With this model we can statistically explain the CAR T cell density over time measured during the ZUMA-1 trial (Fig. 2 C). To recapitulate both peak and tail of the measured CAR T cell densities over time, we implemented both T and CAR T cell co-evolution (Eq. (1)-(2), Methods), as well as a time-sensitive shift in memory CAR T cell proliferation. We allowed for a flexible CAR T cell proliferation. Around day 21, our model predicts that memory CAR T cells gradually switch from a fast to a slow proliferation rate, which could be the result of an overall reduction in concentration of homeostatic cytokines as the normal T cells have almost returned to their pre-lymphodepleted population size.

To incorporate the dynamics of normal T cells after lymphodepletion and CAR T infusion, we used longitudinal data from patients at Moffitt Cancer Center and an estimate of the median long-term homeostatic T cell density[22]. Patients can be seen to quickly return to near-homeostatic values, with 32% normal T cells at day 7, and 94-98% after day 28 (Fig. 2 B, assuming normal density of about 500 cells per μL blood). In order to parameterize our dynamical system, we used model fitting and nonlinear optimization together with normal and CAR T cell counts from two independent sources (Methods), as well as the clinical end-point tumor burden reported by ZUMA-1. To this end, we defined a statistical goal function that would identify a satisfactory set of parameters that enter in Eqs. (1)-(4), which led to characteristic parameter distributions (Supplement), on which we based all further modeling, simulations, and predictions. The median values and ranges of all parameters used are summarized in Supplementary Table 1. Importantly, we found that the CAR T cell carrying capacity is systematically lower than its normal T cell counterpart, that memory CAR T cell dynamics are best explained using a transition from a rapid to a slow growth phase around day 21 post CAR injection, and that memory CAR to effector CAR T cell differentiation is negatively, but weakly affected by tumor mass.

The dynamical system, illustrated in Fig. 2 A, is hierarchical in the CAR T cell compartment. Memory (memory) T cell compartments influence each other and impact more differentiated compartments, but not vice versa. This means that we can establish the quantitative rules of normal and CAR T cell interactions independent of effector CAR T cell and tumor properties.

We first fit a generalized logistic model and found that the optimization routine was selecting a Gompertzian co-evolutionary growth model in the non-tumor killing T cell compartments. We established that T cell expansion toward homeostatic levels is dominated by a Gompertzian growth law[30]. The Gompertzian growth law observed here was recently found in a cell population growth model due to the reduction in available phenotypes as the population increases in size and diversity. In an analogous fashion, as the T cell population returns to homeostatic levels, antigen-experienced phenotypes are filled, leading proportional slowing of growth during immune reconstitution. The resulting decrease in available cell states leads to the characteristic logarithmic approach of the T cells' carrying capacity[31].

Normal T cell dominance is established by a reduced CAR T cell carrying capacity; we estimate that the CAR T cell carrying capacity would be reduced by about 33%. Due to rapid expansion into effector CAR T cells, and an initially higher CAR T cell proliferation rate, the overall CAR T cell populations is non-monotonic and exhibits a peak in the first two weeks after injection. This non-monotonic behavior is critical for tumor eradication.

The four-dimensional deterministic system (1), (2) can be analyzed regarding its steady states and their stability. As we show in the Supplement using linear stability analysis, two long-term outcomes are important. In the first outcome, normal T cells return to homeostatic levels while CAR T cells vanish, and the tumor eventually grows uncontrolled. This single state is stable if the normal T cells' carrying capacity is larger than the CAR T cells' carrying capacity, which is most likely the case. In the second type of outcome, the tumor is brought to a stable value or it is eradicated. However, the two corresponding steady states are stable if CAR T cells' carrying capacity is larger than that of normal cells, which is highly unlikely. Hence, in biologically meaningful settings where CAR T cells are slightly or strongly maladapted in the long run, tumor eradication is unstable. Nonetheless, in many cases the non-monotonic CAR signal indicates that the tumor mass shrinks at least for some time. This implies that treatment success is purely a stochastic event. Tumor eradication is possible if there is a sufficient number of CAR T cells present at the time when the tumor is small enough to be dominated by stochastic effects.

The probability of tumor eradication by effector CAR T cells depends on the amount of CAR T cells left at the intermediate time point, at which the tumor cell population enters a stochastic regime—provided that the tumor responded to CAR injection. In Fig. 2 E, we show two example tumor trajectories of the same initial conditions and tumor growth rate prior to treatment, whereby in one example, the tumor goes extinct around day 20 post injection, while in the other example, the tumor persists and eventually escapes the stochastic regime around day 90, and

starts to grow back, leading to relapse. In order to estimate the probability of tumor extinction (durable response, or cure) we ran simulations of the four-dimensional system, such that the system was integrated in time using deterministic dynamics at high densities, and switching to stochastic (Gillespie, Monte-Carlo) simulations once the tumor cell population was below 100 cells. This revealed that the fraction of CAR memory T cells, and lower tumor growth cell proliferation rates increase the chances of tumor eradication (Fig. 2 F). However, increasing the fraction of memory cells in the product above 80% could yield diminishing returns, which indicates that other CAR T cell properties are the rate limiting factors, such as effector cell life span and the tumor killing rate. Similarly, for a fixed number of infused CAR T cells, higher normal T cell counts in the peripheral blood at infusion, possibly due to less efficient lymphodepleting chemotherapy, have a strong detrimental effect on the probability of tumor eradication (Fig. 2 G). These two examples demonstrate the utility of the stochastic approach. Even if tumor extinction is very likely, but its timing could be critically determined by both intrinsic tumor properties, as well as by variations in properties of the CAR T cell treatment.

The stochastic representation of the CAR T cell evolutionary dynamics allows us to calculate survival curves from 'virtual patients', each defined by a unique set of randomly chosen parameter values. We focused on the two scenarios of transient and long-term response (sketched in Figs. 1 B, C). That is, we excluded non-responder cases from our simulations, and we assumed that even those cases that eventually progressed showed some response. In addition, we did not model the impact of treatment toxicity, and we focused on progression free survival (PFS), defined as the time the tumor stayed below five times its value at treatment initiation (day 0). Thus, any comparison of our PFS curves to previously reported overall survival curves, must be caveated with our lack of inclusion of treatment related deaths which are about 3%[6]. Then, the 'typical' patient, defined by using all median parameter values (TABLE S1) does not progress until day 180 post CAR infusion. We were thus interested in the overall impact of parameter variation on survival outcomes (Fig 3 A), which we approached by assuming all parameters to be normally distributed around the median (mean) with a small variance calculated as a small fraction of the mean. Increasing parameter variance led to more heterogeneous progression free survival (PFS) curves. All PFS curves showed an initial plateau, which stems from the aforementioned fact that we assumed at least some tumor response to CAR T cell treatment, which is generally the case in clinical settings (ie. quantifiable decreased tumor size, if not objective response criteria, is typically seen). All differences between simulated PFS curves were statistically significant to an arbitrary degree, as including

more virtual patients could always be used to establish significance, e.g. using a log-rank test. Hence, we focus on effect sizes.

Figure 3: The mathematical model recapitulates and predicts progression free survival (PFS), and can suggest actionable therapy improvements. 1000 simulated patients were used to generate each PFS curve. A: Impact of parameter variation on the PFS. All subsequent panels have $\sigma = 0.15$. **B**: Impact of tumor growth rate on the PFS. **C**: Impact of initial tumor size on the PFS. The much larger tumors lead to some patients progressing earlier since the CAR could not expand fast enough. **D**: Impact of CAR T infusion phenotype composition. In general, a higher memory fraction lead to better PFS rates. **E**: Impact of lymphodepletion on PFS. Similar to Figure 2G, a sizable impact on PFS is observed by doubling or tripling the amount of normal T cells after depletion. **F**: The distribution of cure times for the median parameters. Most patients are curved before day 40. **G**: The distribution of progression times for the median parameters. Most patients progress between days 20-400.

The intrinsic tumor properties are modeled by tumor growth rate, r_B, and initial tumor burden (B_0, where we report values relative to the median tumor size). The tumor growth rate had a strong detrimental effect on PFS early on (Fig. 3 B). At day 100, PFS of slow-growing tumors (r_B=0.075/day) could be above 90%, whereas PFS of fast-growing tumors (r_B=0.175/day) could be above 70%. In comparison to the median tumor size, ten-fold larger tumors (B_0=10) showed a drop 70% in PFS at day 100 (from 90% PFS), and to 55% for twenty-fold larger tumors (B_0=20, Fig. 3 C). Our model thus predicts that changes in intrinsic tumor properties can have drastic effects on PFS. Especially changes in tumor burden could be observed in patient cohorts early on.

Next, we asked how PFS is influenced by properties of the CAR and normal T cell population prior to injection. To this end, we considered changes in the fraction of CAR memory fraction at day 0, and variation in the efficacy of lymphodepleting chemotherapy. In concordance with the results in Fig. 2 F, larger differences in PFS can be detected when the memory CAR T cell fraction is low (Fig. 3 D). The potential failure to effectively lymphodeplete prior to CAR treatment could have drastic effects; doubling the initial normal T cell population size could halve PFS at day 100 and thereafter. This confirms the hypothesis that sufficient lymphodepletion is pivotal for treatment success. However, we also predict that attempts to increase memory CAR T cell fractions past a certain point might not be necessary, as they would lead to diminishing returns.

Our stochastic simulations show that cure occurs early. Most virtual patients were cured between days 10 and 35, and we rarely found late cure events up to day 100 (Fig. 3 F). This, in turn, could indicate that those rare patient events with cure way beyond 100 days are driven by other mechanisms, such as a vaccine effect. Meanwhile, progression times were distributed over a broad range of time points. Typical progression times occurring anywhere between day 20 and day 400 (Fig. 3 G). These drastic differences in characteristic time scales occur because cure as a stochastic tumor extinction event is much more likely to occur while the effector CAR T cell population persists at sufficiently large densities, which is the case up to day 30 (see Fig. 2 D).

The mathematical model allows us to explore alternative scenarios that were not part of previous clinical studies. It is especially useful to explore alternative treatment strategies. For example, a second CAR T cell treatment at a specified point in time. For this setting, we assumed that a second dose, equal to the first dose in total CAR T cell amount and composition, was given together with an appropriately timed second lymphodepleting

chemotherapy. To assess the amount of survival improvement, we focused on long-term progression free survival (PFS), estimated at day 700 post injection of the second CAR. We then explored this statistical outcome variable as a function of the time point of the second treatment. We focused on variation in the two tumor intrinsic properties, parameterized by initial tumor size, B_0, and background tumor growth rate, r_B, and simulated the second CAR T cell dose identical to the first one. Note that we assumed that the number of CAR T cell injected was chosen to match ZUMA-1. We simulated 1000 virtual patients. With baseline (single) treatment long-term cure rate of 45%, we anticipated that out of these 1000, 550 would progress eventually. Note that our results presented in Fig. 3 G suggest that less than 1% of these progress after day 600, but mostly before day 700. Thus, we considered this time frame a good time to end our simulations, and asses any additional treatment benefit 700 days after the initial CAR was given.

We expected that the additional benefit of a second dose decreases monotonically with timing, and approaches a minimal added benefit asymptotically, which would be similar to a second 'independent' CAR T cell treatment. Indeed, we see that the highest long-term benefit could be achieved if the second CAR were given immediately after the first (Fig. 4 A, B). We do not consider cytotoxicity here, which would be a deterrent for an immediate second dose. Thus, one needs to ask whether there is a trade-off between later second dose and associated loss of long-term benefit. Smaller than median tumor burden did not make a large difference across timing of the second dose, whereas tumors larger than ten times the median initial tumor would benefit from earlier second dose (Fig. 4 A). Patients with a median sized tumor could benefit from second dose at day 200, as this would increase their PFS from 45% baseline to 60%. Patients with a ten-fold larger tumor (same tumor growth rate) could expect the same PFS with a second dose at day 100 or earlier. These results point to the critical need to accurately assess tumor burden, should a second CAR T cell treatment be considered within a desired time frame.

Tumor growth rate had an overall larger effect on long term PFS of a second CAR dose. Assuming a second dose at day 50 could result in long-term PFS of 60% for average tumor growth rates (r_B=0.175/day). This value would decrease to less than 50% in much faster growing tumors (r_B=0.275/day), and to above 65% in slower growing tumors (r_B=0.075/day), assuming the same tumor mass and the same amount of CAR T cells per dose per patient (Fig. 4 B). Our results highlight the impact of the tumor growth rate on survival, a potentially unknown between-patient variability. Faster growing tumors, less likely to be cured by a single dose, could benefit from second dose as early as possible, logistic and toxicity constrains permitting.

The added benefit of a second dose, of as a function of both tumor size and tumor growth rate, is non-monotonic. Especially for intermediate and higher tumor growth rates, there can be a tumor size that yields maximal benefit of a second dose (Fig. 4 C). For very large tumors that grow rapidly, second CAR T cell treatment at day 100 (evaluated as percent increase in PFS at day 200) could be detrimental. On the one hand, these complex relationships emerge because the single dose is already very efficient. Slow growing tumors cannot improve with a second dose, especially those with large tumor mass that progress slowly. On the other hand, large tumors that would progress and are not treated twice within a short period of time will not show improvement for a wide range of tumor growth rates. Hence, average and smaller tumors that grow at median or faster rates could see a survival improvement of 40% or higher relative to the baseline of a single treatment, if the second dose is given not immediately.

Figure 4: The impact of tumor size and growth rate on a second dose of CAR T cells applied after 100 days of the initial dose in regards to long-term PFS (defined at 700 days). PFS=Progression free survival. A: The influence of initial tumor burden on the effect of a second infusion of CAR is minimal. **B**: The influence of tumor growth rate on the effect of a second infusion of CAR is minimal. **C**: The sensitivity of the PFS to initial tumor size and growth rate 200 days after initial dose and 100 days after the second dose.

DISCUSSION

CAR T cell expansion (CAR peak), durability (CAR decay), tumor burden, the effects of lymphodepleting chemotherapy, and the composition of the CAR (CCR7 positive memory cells) are associated with patient outcomes. To model these effects, we devised a cell population dynamic model that describes the co-evolution of CAR T, normal T and tumor cells in patients. When cell densities are large, a deterministic mean field approach can be used to describe dynamics and long-term outcomes. However, we established that durable response, as a result of tumor extinction due to CAR T cell killing, cannot be captured by the mean field model, and requires a stochastic individual based model description. Using a hybrid simulation approach, we confirmed that our model predicts durable response in the event of tumor extinction, or progression. This approach then led to the prediction and sensitivity of progression free survival (PFS) under variation of the composition of the CAR T cell population, the tumor burden, and other patient specific parameters.

Our model confirms the hypothesis that sufficient lymphodepletion is an important factor in determining durable response. The model also predicts that increases in memory CAR T cell fractions (e.g. determined by the fraction of CCR7 positive cells) leads to diminishing returns in additional gains in PFS. This points to other necessary dynamical quantities that affect CAR expansion and tumor killing, likely impacted by upregulation of inflammatory cytokines, such as IL-7 or IL-15[12]. Thus, future modeling that follows from our hybrid stochastic approach developed here should integrate other available signals in patients, such as the dynamics and upregulation of homeostatic and inflammatory cytokines.

Several of the mechanistic assumptions and practical simplifications we made in our model may be oversimplified in the broader biological context of CAR T cell therapy. First, we assume that immune reconstitution varies minimally across patients, but our results demonstrate that this adds sufficient feed-back to explain CAR spike and decay. A different approach would be to re-parameterize normal T cell dynamics over time, which could reveal whether variability in normal T cell growth is necessary to explain outcomes. Alternatively, CAR density over time could be driven by unknown external stimuli, resulting in a birth and death process with a time-dependent tumor killing rate, which could be used to explain the distribution of times progression, applying computational strategies that integrate time-dependence [32,33].

Second, our model assumes that tumor cell proliferation rate is independent of tumor burden. However, the tumor growth rate might depend on tumor burden, compared to carrying capacity. In this context, one could explore other sources of tumor burden variability that originates from a

logistic dependence of tumor cell proliferation on tumor volume, called proliferation-saturation [34,35]. This dependence leads to high proliferation rate in small tumors. Future studies of the stochastic process dominating tumor killing in proportion to tumor cell-proportional antigen-load should carefully consider potential non-linear relationships between tumor size and growth as a potential dynamic biomarker for the success of CAR T cell therapy.

Our probabilistic measure of PFS does not include the evolution of resistance to CAR T cell therapy by genetic or epigenetic mechanisms[16]. Thus, our results can be seen as conditional on non-resistance, which would add an additional probabilistic modeling layer.

CAR T cell therapy can cause severe cytokine release syndrome (CRS), or a characteristic immune cell-associated neurologic syndrome (ICANS). Both are associated with significant morbidity and mortality[36-38], and are characterized by high levels of inflammatory cytokines and, to a lesser extent, by high numbers of CAR T cells in the blood[6]. Both forms of toxicity occur in predictable distributions across patient populations[39]. Published data from the ZUMA-1 trial demonstrated an association between CAR T cell expansion and grade 3-4 (severe) ICANS associated toxicities, but no association with grade 3-5 (severe) CRS[6]. Elevation of key inflammatory cytokines (IFN- γ, IL-6, IL-1) are also associated with both severe ICANS and CRS, however a homeostatic cytokine (IL-15) seen at higher levels in responding patients was also associated with toxicity. Several biomarkers were associated with severe ICANS associated toxicity but not CRS (ferritin, IL2-Rα and GM-CSF)[40,41]. Thus, one could hypothesize that CAR T cell expansion and the occurrence of toxicities can be explained by modeling dynamic cytokine/biomarker levels in periphery together with CAR T cell repopulation dynamics and toxicity in relation to homeostatic (IL-15) and inflammatory (IL-6, IFN-γ, IL-1) cytokines. Computational and mathematical tools similar to the ones we developed here could be used to calibrate the model, assess parameter sensitivities, and make stochastic modeling predictions of patient outcomes under varying conditions of the factors that determine toxicity of CAR T cell therapy, in order to understand how these events interfere with durable response, and to elucidate the basic mechanism of T cell expansion-related toxic events. Here we present an important first step towards integration of these mechanisms into a quantitative understanding of adoptive therapy of CAR T cell therapy.

TABLES

Biological parameter	Symbol	Median value
Normal T cell carrying capacity	K_N	5.00×10^2 cells/µL[22]
CAR memory T cell carrying capacity	K_M	1.71×10^2 cells/µL
Normal T cell net growth rate	r_N	1.60×10^{-1} day^{-1}
Initial CAR memory net growth rate	$r_{M,max}$	4.28×10^{-1} day^{-1}
Final CAR memory net growth rate	$r_{M,min}$	1.07×10^{-2} day^{-1}
CAR memory net growth rate switch	τ	2.08×10^1 days
Basal CAR asymmetric diff. rate	$r_E(0)$	1.30×10^0 day^{-1}
Effector death rate	d_E	2.93×10^{-1} day^{-1}
Tumor net growth rate	r_B	7.44×10^{-2} day^{-1}
Effector exhaustion rate	γ_E	1.25×10^{-3} day^{-1}* cm^{-2}
Tumor killing rate	γ_B	5.00×10^{-3} µL/(days*cells)
Impact of tumor volume	α_1	3.54×10^{-2} cm^{-2}
Impact of tumor volume	α_2	$> 5.00 \times 10^0$ cm^{-2}
Initial median normal T cell number	N_0	3.00×10^9 cells
Initial CAR T cell number	$M_0 + E_0$	1.20×10^8 cells
Initial median tumor cell number	B_0	5.27×10^{10} cells

Table 1: Parameter and initial condition values as identified by our machine learning procedure (Supplement).

REFERENCES

Howlader, N. *et al.* SEER Cancer Statistics Review, 1975-2014. (2014).

Crump, M., et al. Outcomes in refractory aggressive diffuse large B-cell lymphoma (DLBCL): Results from the international SCHOLAR-1 study. *Journal of Clinical Oncology* **34**, 7516-7516 (2016).

Schroeder HW, R. R.   Ch. Antigen receptor genes, gene products, and co-receptors, 47-51 (2013).

Locke, F. L. *et al.* Long-term safety and activity of axicabtagene ciloleucel in refractory large B-cell lymphoma (ZUMA-1): a single-arm, multicentre, phase 1-2 trial. *Lancet Oncol* **20**, 31-42, doi:10.1016/S1470-2045(18)30864-7 (2019).

Locke, F. L. *et al.* Phase 1 Results of ZUMA-1: A Multicenter Study of KTE-C19 Anti-CD19 CAR T Cell Therapy in Refractory Aggressive Lymphoma. *Mol Ther* **25**, 285-295, doi:10.1016/j.ymthe.2016.10.020 (2017).

Neelapu, S. S. *et al.* Axicabtagene Ciloleucel CAR T-Cell Therapy in Refractory Large B-Cell Lymphoma. *N Engl J Med* **377**, 2531-2544, doi:10.1056/NEJMoa1707447 (2017).

Crump, M. *et al.* Outcomes in refractory diffuse large B-cell lymphoma: results from the international SCHOLAR-1 study. *Blood* **130**, 1800-1808, doi:10.1182/blood-2017-03-769620 (2017).

Schuster, S. J. *et al.* Chimeric Antigen Receptor T Cells in Refractory B-Cell Lymphomas. *N Engl J Med* **377**, 2545-2554, doi:10.1056/NEJMoa1708566 (2017).

Abramson, J. S. *et al.* High durable CR rates and preliminary safety profile for JCAR017 in R/R aggressive b-NHL (TRANSCEND NHL 001 Study): A defined composition CD19-directed CAR T-cell product with potential for outpatient administration. *Journal of Clinical Oncology* **36**, 120-120, doi:10.1200/JCO.2018.36.5_suppl.120 (2018).

Fraietta, J. A. *et al.* Determinants of response and resistance to CD19 chimeric antigen receptor (CAR) T cell therapy of chronic lymphocytic leukemia. *Nat Med* **24**, 563-571, doi:10.1038/s41591-018-0010-1 (2018).

Turtle, C. J. *et al.* Immunotherapy of non-Hodgkin's lymphoma with a defined ratio of CD8+ and CD4+ CD19-specific chimeric antigen receptor-modified T cells. *Sci Transl Med* **8**, 355ra116, doi:10.1126/scitranslmed.aaf8621 (2016).

Kochenderfer, J. N. *et al.* Lymphoma Remissions Caused by Anti-CD19 Chimeric Antigen Receptor T Cells Are Associated With High Serum Interleukin-15 Levels. *J Clin Oncol* **35**, 1803-1813, doi:10.1200/JCO.2016.71.3024 (2017).

Hart, Y. *et al.* Paradoxical signaling by a secreted molecule leads to homeostasis of cell levels. *Cell* **158**, 1022-1032, doi:10.1016/j.cell.2014.07.033 (2014).

Zhu, J. & Paul, W. E. Peripheral CD4+ T-cell differentiation regulated by networks of cytokines and transcription factors. *Immunol Rev* **238**, 247-262, doi:10.1111/j.1600-065X.2010.00951.x (2010).

Wong, H. S. & Germain, R. N. Robust control of the adaptive immune system. *Seminars in immunology* **36** (2018).

Altrock, P. M., Liu, L. L. & Michor, F. The mathematics of cancer: integrating quantitative models. *Nature reviews. Cancer* **15**, 730-745, doi:10.1038/nrc4029 (2015).

Vincent, T. L. & Brown, J. S. *Evolutionary game theory, natural selection, and Darwinian dynamics*. (Cambridge University Press, 2005).

Kimmel, G. J., Gerlee, P., Brown, J. S. & Altrock, P. M. Neighborhood size-effects shape growing population dynamics in evolutionary public goods games. *Communications Biology* **2**, 53 (2019).

Werner, B. *et al.* The cancer stem cell fraction in hierarchically organized tumors can be estimated using mathematical modeling and patient-specific treatment trajectories. *Cancer Research* **76**, 1705-1713 (2016).

Wodarz, D. & Komarova, N. *Computational biology of cancer: lecture notes and mathematical modeling*. (World Scientific, 2005).

Nowak, M. A. *Evolutionary Dynamics*. (2006).

Turtle, C. J. *et al.* CD19 CAR–T cells of defined CD4+: CD8+ composition in adult B cell ALL patients. *The Journal of Clinical Investigation* **126**, 2123-2138 (2016).

Mogensen, P. K. & Riseth, A. N. Optim: A mathematical optimization package for Julia. *Journal of Open Source Software* **3**, 615 (2018).

Rackauckas, C. & Nie., Q. Differentialequations. jl–a performant and feature-rich ecosystem for solving differential equations in julia. *Journal of Open Research Software* **5** (2017).

Gillespie, D. T. Exact Stochastic Simulation of Coupled Chemical Reactions. *The Journal of Physical Chemistry* **81**, 2340–2361 (1977).

Locke, F. L. *et al.* Immune signatures of cytokine release syndrome and neurologic events in a multicenter registrational trial (ZUMA-1) in subjects with refractory diffuse large B cell lymphoma treated with axicabtagene ciloleucel (KTE-C19) [abstract]. *Cancer Research* **77**, CT020, doi:doi:10.1158/1538-7445.AM2017-CT020 (2017).

Wherry, E. J. & Ahmed, R. Memory CD8 T-cell differentiation during viral infection. *Journal of Virology* **78**, 5535-5545 (2004).

Farber, D. L., Yudanin, N. A. & Restifo, N. P. Human memory T cells: generation, compartmentalization and homeostasis. *Nat Rev Immunol* **14**, 24-35, doi:10.1038/nri3567 (2014).

Zhan, Y., Carrington, E. M., Zhang, Y., Heinzel, S. & Lew, A. M. Life and Death of Activated T Cells: How Are They Different from Naive T Cells? *Front Immunol* **8**, 1809, doi:10.3389/fimmu.2017.01809 (2017).

Gerlee, P. The model muddle: in search of tumor growth laws. *Cancer research* **73**, 2407-2411 (2013).

West, J. & Newton, P. K. Cellular interactions constrain tumor growth. *Proceedings of the National Academy of Sciences* **116**, 1918-1923 (2019).

Höhna, S. The time-dependent reconstructed evolutionary process with a key-role for mass-extinction events. *Journal of Theoretical Biology* **380**, 321-331 (2015).

Dinh, K. N. & Sidje, R. B. An adaptive Magnus expansion method for solving the chemical master equation with time-dependent propensities. *Journal of Coupled Systems and Multiscale Dynamics* **5**, 119-131 (2017).

Prokopiou, S. *et al.* A proliferation saturation index to predict radiation response and personalize radiotherapy fractionation. *Radiation Oncology* **10**, 159 (2015).

Poleszczuk, J. *et al.* Predicting patient-specific radiotherapy protocols based on mathematical model choice for Proliferation Saturation Index. *Bulleting of Mathematical Biology* **80**, 1195-1206 (2017).

Neelapu, S. S. *et al.* Chimeric antigen receptor T-cell therapy - assessment and management of toxicities. *Nat Rev Clin Oncol* **15**, 47-62, doi:10.1038/nrclinonc.2017.148 (2018).

Lee, D. W. *et al.* ASTCT Consensus Grading for Cytokine Release Syndrome and Neurologic Toxicity Associated with Immune Effector Cells. *Biol Blood Marrow Transplant* **25**, 625-638, doi:10.1016/j.bbmt.2018.12.758 (2019).

Dholaria, B. R., Bachmeier, C. A. & Locke, F. Mechanisms and Management of Chimeric Antigen Receptor T-Cell Therapy-Related Toxicities. *BioDrugs* **33**, 45-60, doi:10.1007/s40259-018-0324-z (2019).

Jacobson, C. A. CD19 Chimeric Antigen Receptor Therapy for Refractory Aggressive B-Cell Lymphoma. *J Clin Oncol* **37**, 328-335, doi:10.1200/JCO.18.01457 (2019).

Turtle, C. J. *et al.* CD19 CAR-T cells of defined CD4+:CD8+ composition in adult B cell ALL patients. *J Clin Invest* **126**, 2123-2138, doi:10.1172/jci85309 (2016).

Park, J. H. *et al.* Long-Term Follow-up of CD19 CAR Therapy in Acute Lymphoblastic Leukemia. *N Engl J Med* **378**, 449-459, doi:10.1056/NEJMoa1709919 (2018).

Evolutionary Dynamics of CAR T-Cell Therapy: Supplemental Information

Gregory J. Kimmel   Frederick L. Locke
Philipp M. Altrock

July 29, 2019

Contents

1   Model motivation

A mechanistic mathematical model which explains the likelihood of successful treatment would be very useful in clinical applications. We will present a game-theoretic model, which describes a mechanism that leads to the characteristic behavior observed in the patients in the ZUMA-1 trial [1].

Figure 1 shows the inner quartiles of the trial. The spike in CAR concentration occurs around a week after infusion. In addition, we observe that three months post-infusion, the CAR population is still largely present in the peripheral blood in more than 50% of patients.

Figure 1: Quartile data reproduced from the ZUMA-1 trial [1].

We motivate our model via the following observations:

- If chemo depletion of normal T-cells does not occur, the CAR T-cells do not expand *in vivo*. **Implication:** *A carrying capacity must exist which is related to the total population of T-cells. If this wasn't the case, CAR T-cells would always expand, which is not observed.*

- The CAR T-cells have an initial spike in concentration before decreasing and leveling out to a constant (which may or may not be zero). **Implication:** *The amount of space available to the CAR T-cells must be changing over time.*

We propose a potential mechanism which drives the initial spike and decay of CAR T concentration back to lower levels (or potentially zero). We postulate that the mechanism that drives the dynamics of the CAR population is competition between wild-type (or normal) T-cells and CAR T-cells. This model is particularly appealing for its biological simplicity, the size of the normal population can be thought of as a proxy for the microenvironment, that contains information about the cytokine levels.

1.1   Reemergence of normal T-cells

This model was motivated from two key observations. First, chemo-depletion of T-cells was required for proliferation of CAR T-cells. Second, after the spike at around 7 days, the CAR population levels dropped (in some cases, below level of detection). This was observed even in the case of disease remission. The first observation implies that there is an innate carrying capacity shared by both types of T-cells (wild-type and CAR). The initial depletion is necessary to give CAR cells an ability to proliferate. The second observation implies that the wild-type population has an innate advantage. If this was not the case, the long-term levels of CAR would not tend to zero.

First, let us state our model assumptions. There are up to four subtypes of T-cells relevant to the CAR: Naive, central memory, effector memory and effector cells. We first reduce the complexity by consolidating these four subtypes into two, which we call: juvenile and effector. By juvenile, we mean naive and central memory, those cells which, when presented with their respective antigen, differentiate into the effector memory and effector cell types. We combine these two subtypes into an "effector" compartment that target their respective antigen. Next, we introduce the competition that occurs between CAR and normal T-cells at the juvenile level. That is, we assume a negative feedback on these two populations as they fill the ecological niche. We assume that the asymmetric differentiation rate can be antigen-dependent and that the effector cells have a finite life-span. We introduce a predator-prey interaction term, which defines the killing rate of the tumor cells when presented with an effector cell. For simplicity, we assume a linear

(Holling type I) predation rate, but other more complex interaction terms can be considered. Finally, we assume simple exponential growth of the tumor.

Mathematically, these assumptions and relationships may be cast as the following 4D ordinary differential equation (ODE):

$$\dot{N} = N(r_N - a_{11}N - a_{12}M), \tag{1a}$$

$$\dot{M} = M(r_M - a_{21}N - a_{22}M), \tag{1b}$$

$$\dot{E} = r_E(B)M - \gamma_E BE - d_E E, \tag{1c}$$

$$\dot{B} = r_B B - \gamma_B BE. \tag{1d}$$

With initial concentrations N_0, M_0, E_0, B_0, respectively. Here $r_{N,M}$ is the net growth rate of normal and CAR memory T-cells, respectively. The coefficients a_{ij} are the interaction rates that determine the magnitude of negative feedback that the j subtype has on the i subtype. r_E is the asymmetric differentiation rate and d_E is the CAR effector death rate. The $\gamma_{E,B}$ is the effector-tumor interaction rate where the effector and tumor cell dies, respectively. Finally, r_B is the tumor net growth rate.

1.1.1 Carrying capacity formulation

A general formulation of interaction involves a competitive-Lotka-Volterra model with interaction coefficients a_{ij}. However, it is very difficult to infer these parameters from biological data. A simplifying assumption involves scaling these coefficients $a_{ij} = 1/K_i$, the carrying capacity that each subtype would have in the absence of the other subtype. Hence, by redefining the parameters $a_{ij} = r_i/K_i$, we arrive at

$$\dot{N} = r_N N \left[1 - \left(\frac{N + M}{K_N} \right)^{\beta} \right], \tag{2a}$$

$$\dot{M} = r_M M \left[1 - \left(\frac{N + M}{K_M} \right)^{\beta} \right], \tag{2b}$$

$$\dot{E} = r_E(B)M - \gamma_E BE - d_E E, \tag{2c}$$

$$\dot{B} = r_B B - \gamma_B BE. \tag{2d}$$

Preliminary optimization of the model parameters, shows that $\beta \to 0$ and that r_i is getting large. Rescaling $r_i = r_i/\beta$ ($i = N, M$) and letting $\beta \to 0$

we obtain the final form for our model

$$\dot{N} = -r_N N \ln\left(\frac{N+M}{K_N}\right), \tag{3a}$$

$$\dot{M} = -r_M M \ln\left(\frac{N+M}{K_M}\right), \tag{3b}$$

$$\dot{E} = r_E(B)M - \gamma_E BE - d_E E, \tag{3c}$$

$$\dot{B} = r_B B - \gamma_B BE. \tag{3d}$$

Steady states. Since N and M are only dependent on their 2D compartment, we can consider them separately. If $M^* = 0$, then $N^* = 0, K_N$. Suppose $M^* \neq 0$, then if $N^* = 0$ we have $M^* = K_M$. A coexistence state is only observed in the degenerate case $K_N = K_M$, as indefinite persistence of CAR T-cells was not observed, we can conclude that $K_N > K_M$. This helps limit the parameter region.

We note that in this formulation there are *no* coexistence states. The stability and states can be read off from the previous section and we summarize the results here:

- **Unphysical solution**: $(0, 0, 0, e^{r_B t})$ - the T-cell-absent state (unphysical - there should be a nonzero wild-type concentration). This state is in fact *always* unstable.

- **Failed treatment**: $(K_N, 0, 0, e^{r_B t})$ - the CAR cells are depleted and the tumor is not eliminated, leading to exponential tumor growth. This state is stable if $K_N > K_M$.

- **Cured patient**: $\left(0, K_M, \frac{K_M r_E(0)}{d_E}, 0\right)$ - the CAR cells eliminate the tumor and outcompete the wild-type to become the resident population in the body (potentially unrealistic - we expect the wild-type to have an innate advantage). This state is stable if

$$K_M > \max\left(K_N, \frac{d_E r_B}{\gamma_B r_E(0)}\right).$$

- **Stable patient**: $\left(0, K_M, \frac{r_B}{\gamma_B}, B^*_{\text{stable}}\right)$, where B^*_{stable} satisfies Eq. (**??**) - the CAR cells dynamically stabilize the tumor (potentially unrealistic -

we expect the wild-type population to be present). This state is stable
if
$$K_M > \max\left(K_N, \frac{\gamma_E r_B}{\gamma_B r'_E(B^*_{\text{stable}})}\right)$$

- **Tumor growth**: This state is slightly different than the rest in that it
  is not a "fixed point". Ultimately, we are concerned with tumor growth
  and if the maximum value E can take is such that $\dot{B} > 0$, then we have
  an unsuccessful treatment outcome. Hence if $E(t) < E_{\max} = \frac{r_B}{\gamma_B}$ the
  tumor will grow (note that this bounds one away from the stable patient
  outcome).

We see that with the assumption that $K_N > K_M$, the only long-term states
are failed treatment and eventual tumor growth. In this deterministic setting
however, the tumor cell population gets very small and spends a lot of time
in a regime where random cell death events would lead to the elimination of
tumor. These small-population size effects cannot be adequately captured
with an ODE model and so we will eventually turn to a stochastic model
formulation.

2 Data analysis

Here, we use the median, lower and upper and quartiles of the 101 patients
on the ZUMA-1. This data set contains total peripheral CAR T-cell concen-
trations (quantiles), from peripheral blood measurements, recorded at days
7, 14, 28, 90, and 180 post CAR injections. These 15 data points, in combi-
nation with normal CAR T-cell dynamics obtained independently, were used
to fit the parameters of our complex T-cell/CAR T-cell and tumor cell dy-
namics model in its deterministic limit. Note that over time, the breakdown
of the total CAR into memory/juvenile and effector CAR T-cells was not
available. Hence, when fitting these data, one must either make assumptions
on the phenotype breakdown at each available time point, or compare the
model solution $(M + E)(t)$ concentration against the total CAR T data.

To fit the data, we implemented a BFGS constrained optimization rou-
tine using Julia's built in optimization and differential equation solvers to
determine the optimal parameters which minimize our predefined loss func-
tion [2, 3]. We used a weighted-least squares function where the weights are
the variance at each time point. Since we only have the quartiles, we assume

that the distribution of CAR T-cell levels at each time point is normally distributed about the median and calculate an estimate of the variance from the interquartile range

$$\sigma^2 \approx \left(\frac{Q_3 - Q_1}{1.35} \right)^2.$$ (4)

We first fit the normal and memory CAR T-cells since they are uncoupled from the other two populations. This loss function is

$$\theta_{N,M} = \lambda_0 \sum \frac{[M(t_i) - \hat{M}_i]^2}{\sigma_{M,i}^2} + \lambda_1 \sum [N(t_i) - \hat{N}_i]^2,$$ (5)

where the λ's are tunable weights that can be used to adjust the loss function landscape. Different λ's lead to different optimal parameters. To generate multiple trajectories we took the 5 time points and 3 quartiles to generate $3^5 = 243$ trajectories. We then made the assumption that the memory fraction f_{mem} was fixed over time. We chose $f_{\mathrm{mem}} = [0.1, 0.5, 0.9]$. This leads to a total of 729 trajectories. It was clear quickly, that the initial 14 days (spike and initial decline) could be well captured by three fitted parameters r_N, r_M, K_M. However the decline from the quartile showed a slowing down sometime after day 14 - some patients still had observable CAR levels 6 months post-infusion. To account for this, we introduced two additional parameters that govern a switch in the size of r_M. We could have also modified K_M, but as this quantity is intricately linked to the long-term stability, it made more sense from a fitting perspective to put this dependence in rates into r_M. This could also be tied to changing signal levels, since by day 28, the normal T-cell count had returned to pre-chemo depleted levels. Hence, we introduced the function

$$r_M(t) = \frac{r_{M,\max} - r_{M,\min}}{1 + e^{t-\tau}} + r_{M,\min}.$$ (6)

This was a useful modification as figure 2A shows a clear difference between $r_{M,\max}$ and $r_{M,\min}$. For a given trajectory, we have 11 data points (six for normal T-cells and five for CAR memory) and we are fitting six parameters. For the median trajectory, this led to an $R^2 = 0.976$. One could also ask how much improvement the additional two parameters (τ and partitioning r_M into $r_{M,\min}$ and $r_{M,\max}$) gives. This is known as the coefficient of partial determination and is calculated in an analogous way and we denote it

as $R^2_{\text{partial}} = 0.69$. This shows that an additional 69% of the variation is explained by introducing these extra parameters.

Many functional forms can be given for $r_E(B)$. Clinical trial data and *in vitro* experiments show that more antigen presentation leads to an increase in differentiation into effector cells [4, 5, 6]. However, very large tumors had a mechanism, which negatively impacted this rate. We modeled this via the function

$$r_E(B) = r_E(0) \left[1 + e^{-\alpha_1 B} \left(1 - e^{-\alpha_2 B} \right) \right] . \tag{7}$$

To obtain the parameters in the E and B compartment, we introduced a second loss function

$$\theta_{E,B} = \lambda_0 \sum \frac{[E(t_i) - \hat{E}_i]^2}{\sigma^2_{E,i}}, \tag{8}$$

where the fitted parameters for M, N are used to solve for the remaining parameters. The interquartile range (IQR) is shown in table 1.

Figure 2: Violin plots showing the distribution of the parameters $r_{M,\text{max}}, r_{M,\text{min}}, K_M, \tau$ (green, blue, red, gold) using 729 trajectories generated from the five time points, three quartiles and three memory frequencies.

Biological parameter	Symbol	Value	IQR/Reference
Normal T-cell carrying capacity	K_N	5.00×10^2 cells/μL	[7]
CAR memory T-cell carrying capacity	K_M	1.71×10^2 cells/μL	$[1.22, 2.30] \times 10^2$
Normal net growth rate	r_N	1.60×10^{-1} day^{-1}	
Initial CAR memory net growth rate	$r_{M,\max}$	4.28×10^{-1} day^{-1}	$[3.25, 6.00] \times 10^{-1}$
Final CAR memory net growth rate	$r_{M,\min}$	1.07×10^{-2} day^{-1}	$[0, 2.76] \times 10^{-2}$
CAR memory net growth rate switch	τ	2.08×10^1 days	$[1.53, 3.03] \times 10^1$
Basal CAR asymm. diff. rate	$r_E(0)$	1.30×10^0 day^{-1}	$[0.2303, 3.661] \times 10^0$
Effector death rate	d_E	2.93×10^{-1} day^{-1}	$[1.10, 15.24] \times 10^{-1}$
Tumor growth rate	r_B	7.44×10^{-2} day^{-1}	$[0.414, 5.035] \times 10^{-1}$
Effector exhaustion rate	γ_E	1.25×10^{-3} days^{-1}*cm^{-2}	$[2.11, 28.30] \times 10^{-2}$
Tumor killing rate	γ_B	5.00×10^{-3} (days*cells/μL)$^{-1}$	$[2.11, 28.30] \times 10^{-2}$
Impact of tumor volume	α_1	3.54×10^{-2} cm^{-2}	$[0, 28.89] \times 10^{-2}$
Impact of tumor volume	α_2	$> 5 \times 10^0$ cm^{-2}	$[3.75, \infty) \times 10^0$

Table 1: Median parameter fitting obtained by minimizing the loss functions (5) and (8).

3 Small fluctuations are relevant in the small tumor limit

We have discussed the stability and fixed points of Eq. (1). The treatment is initially successful and the spike in CAR is observed. As the wild-type reemerges and the CAR reaches a homeostatic level, the tumor eventually grows and relapse occurs. Though this scenario is plausible, it was also observed that the patient had a complete response. The culprit revolves around the deterministic system governed by our set of ODEs. When the tumor population becomes small, stochastic effects (fluctuations) become relevant. These chance extinction events is what could lead to the cure observed in patients. In fact, our model dictates that this is a *necessary* requirement if CAR T-cells do not persist indefinitely.

In what follows, we will propose a framework which in the large population limit will converge to the deterministic dynamics.

3.1 Stochastic dynamics

Let W, X, Y, Z be the size of wild-type, CAR memory, CAR effector and tumor cells, respectively. We define the following birth events,

$$N \xrightarrow{r_N} N + N, \tag{9}$$

$$M \xrightarrow{r_M} M + M, \tag{10}$$

$$M \xrightarrow{r_E(B)} M + E, \tag{11}$$

$$B \xrightarrow{r_B} B + B, \tag{12}$$

and the following death events

$$N \xrightarrow{\delta_N(N,M)} \emptyset, \tag{13}$$

$$M \xrightarrow{\delta_M(N,M)} \emptyset, \tag{14}$$

$$E \xrightarrow{\delta_E(B)} \emptyset, \tag{15}$$

$$B \xrightarrow{\gamma_B(E)} \emptyset. \tag{16}$$

We can now calculate the transition rates to move state i to state j. We choose a time interval small enough such that only a single event occurs (all other events occur on the order $O(\Delta t)^2$). Defining $T(N\pm1, M, E, B|N, M, E, B) = T_N^{\pm}$ and noting that r_i can depend on other populations we have the population transition rates

$$T_i^+ = r_i i, \quad i = w, x, y, z, \tag{17}$$

$$T_N^- = \delta_N N = r_N N \ln\left(\frac{N + M}{K_N}\right), \tag{18}$$

$$T_M^- = \delta_M M = r_M M \ln\left(\frac{N + M}{K_M}\right), \tag{19}$$

$$T_E^- = \delta_E E = \gamma_E E B + d_E E \tag{20}$$

$$T_B^- = \delta_B B = \gamma_B E B. \tag{21}$$

Denoting $P(N - 1, M, E, B, t) = P_{N-1}$, the corresponding master equation is given by

$$\frac{\partial P}{\partial t} = T_{i-1}^+ P_{i-1} + T_{i+1}^- P_{i+1} - (T_i^+ + T_i^-)P, \tag{22}$$

where $i = N, M, E, B$. A short coming of an exact simulation is the computing time required for the typical tumor volume. The number of cells in this environment is on the order of millions to billions. The number of events to observe a complete response is when the number of tumor cells goes to 0. The minimum number of events needed then is clearly B_0, where B_0 is the initial number of tumor cells. Of course this requires each event to lead to the removal of one tumor cell so in practice the total number of events needed will exceed B_0 by a substantial amount. The amount of time it takes for a single run can vary from hours to a day. The issue is that the rates are so high, which leads to the time to next event to occur almost instantaneously. A way to overcome this obstacle was proposed by Gillespie himself, who suggested the method known as "tau-leaping" [8]. This method is similar to forward Euler for a continuous system except the update is a Poisson distributed random variable. The time τ is given as a fixed step-size, which alleviates the events occurring rapidly, but error is introduced because the simulation is no longer exact (populations are assumed to be constant in the time interval $[t, t+\tau]$. This method works nicely, but what if the populations exhibit substantially different scales. That is, suppose there are two populations and one is small (stochastic fluctuations are relevant) and the other is very large (fluctuations are not relevant). In this case, tau-leaping could potentially lead to large error if a step is too large, but could be computationally costly if it is too small. The issue is that one populations update rates occur on vastly different scales than the other and so the timing of the events are separate. This issue is analogous to our problem where the tumor population decreases rapidly, but both the normal and CAR T-cells remain relatively high for some time. Therefore we propose a method that exploits this separation of time scales to produce a fast algorithm, while maintaining the fluctuations necessary for complete response.

3.2   Hybrid model

Hybrid models that combine the speed of deterministic models with the accuracy of stochastic models has been used before in many different contexts [9, 10, 11]. Here, we develop a hybrid model for our four population system. The process involves communicating information about the population between both models and using them when appropriate. We will need to keep track of the discrete and continuous population sizes (to convert from continuous to discrete, we used the ceiling function). The final piece needed

to connect the two is to define the threshold at which we consider fluctuations important. We set the threshold to be at 100 cells. A outline of our procedure is given below:

- Define patient-level parameters (e.g. tumor growth rate, initial levels of CAR and normal T-cells)

- Initialize the discrete and continuous population vectors.

- Set $t = 0.0$ and define the final time to run the simulation (we typically used $t = 300 - 1000$ measured in days).

- While $t < t_{\text{final}}$:

- Solve the ODE forward in time until the tumor population goes below the stochastic threshold (e.g. 100 cells) or the final time is reached.

- Update the new continuous and discrete populations.

- If a discrete (stochastic) event occurred, update the populations.

- Determine the time till next *stochastic* event τ and simulate the deterministic system until this next event (or t_{final}) is reached.

This method allows us to utilize the speed of simulating an ODE for the large populations and allows us to catch the complete response observed when the tumor population stochastically goes extinct.

3.3 Figure generation

To generate figure 3 and 4 of the main text we ran the hybrid model for 1000-10000 computer-simulated patients. We recorded the time and rate of cure or progression for each patient and parameter set. The Kaplan-Meier (K-M) curves were generated in *Mathematica*, using the function *SurvivalModelFit*. We introduced a perturbation parameter σ that governs the variability of all fitted parameters. In this sense, σ can be regarded as a sensitivity analysis parameter. Consider a parameter α, then the perturbed values can be in the range $\alpha[1 + \sigma U(-1, 1)]$, where $U(a, b)$ is a uniform distribution with bounds a, b. From a clinical trial standpoint, we can consider σ as a patient variability measure. For example, if we wish to simulate a *single* patient to determine that individuals optimal treatment, we would choose $\sigma = 0$.

However if we wanted to generate a *cohort*, we would consider $\sigma \neq 0$. For most data analysis, we chose $\sigma = 0.05 - 0.15$.

References

[1] Sattva S Neelapu, Frederick L Locke, Nancy L Bartlett, Lazaros J Lekakis, David B Miklos, Caron A Jacobson, Ira Braunschweig, Olalekan O Oluwole, Tanya Siddiqi, Yi Lin, et al. Axicabtagene ciloleucel car t-cell therapy in refractory large b-cell lymphoma. *New England Journal of Medicine*, 377(26):2531–2544, 2017.

[2] Christopher Rackauckas and Qing Nie. Differentialequations. jl–a performant and feature-rich ecosystem for solving differential equations in julia. *Journal of Open Research Software*, 5(1), 2017.

[3] Patrick Kofod Mogensen and Asbjørn Nilsen Riseth. Optim: A mathematical optimization package for julia. *Journal of Open Source Software*, 3(24), 2018.

[4] Andreas Mayer, Yaojun Zhang, Alan S. Perelson, and Ned S. Wingreen. Regulation of t cell expansion by antigen presentation dynamics. *Proceedings of the National Academy of Sciences*, 116(13):5914–5919, 2019.

[5] Zvi Grossman, Booki Min, Martin Meier-Schellersheim, and William E Paul. Concomitant regulation of t-cell activation and homeostasis. *Nature Reviews Immunology*, 4(5):387, 2004.

[6] Karolina Pilipow, Alessandra Roberto, Mario Roederer, Thomas A Waldmann, Domenico Mavilio, and Enrico Lugli. Il15 and t-cell stemness in t-cell–based cancer immunotherapy. *Cancer research*, 75(24):5187–5193, 2015.

[7] Cameron J Turtle, Laïla-Aïcha Hanafi, Carolina Berger, Theodore A Gooley, Sindhu Cherian, Michael Hudecek, Daniel Sommermeyer, Katherine Melville, Barbara Pender, Tanya M Budiarto, et al. Cd19 car–t cells of defined cd4+: Cd8+ composition in adult b cell all patients. *The Journal of clinical investigation*, 126(6):2123–2138, 2016.

[8] Daniel T Gillespie. Approximate accelerated stochastic simulation of chemically reacting systems. *The Journal of Chemical Physics*, 115(4):1716–1733, 2001.

[9] Fabian Spill, Pilar Guerrero, Tomas Alarcon, Philip K Maini, and Helen Byrne. Hybrid approaches for multiple-species stochastic reaction–diffusion models. *Journal of computational physics*, 299:429–445, 2015.

[10] Karla Misselbeck, Luca Marchetti, Martha S Field, Marco Scotti, Corrado Priami, and Patrick J Stover. A hybrid stochastic model of folate-mediated one-carbon metabolism: Effect of the common c677t mthfr variant on de novo thymidylate biosynthesis. *Scientific reports*, 7(1):797, 2017.

[11] Aurélien Alfonsi, Eric Cances, Gabriel Turinici, Barbara Di Ventura, and Wilhelm Huisinga. Adaptive simulation of hybrid stochastic and deterministic models for biochemical systems. In *ESAIM: proceedings*, volume 14, pages 1–13. EDP Sciences, 2005.