Identifiability of heterogeneous phenotype adaptation from low-cell-count experiments and a stochastic model

Alexander P Browning, Rebecca M Crossley, Chiara Villa, Philip K Maini, Adrianne L Jenner, Tyler Cassidy, Sara Hamis

Identifiability of heterogeneous phenotype adaptation from low-cell-count experiments and a stochastic model

Alexander P Browning[1], Rebecca M Crossley[1], Chiara Villa[2], Philip K Maini[1], Adrianne L Jenner[3], Tyler Cassidy[4], and Sara Hamis[5]

[1] *Mathematical Institute, University of Oxford, Oxford, United Kingdom*
[2] *Sorbonne Université, CNRS, Université de Paris, Inria, Laboratoire Jacques-Louis Lions UMR, Paris, France*
[3] *School of Mathematical Sciences, Queensland University of Technology, Brisbane, Australia*
[4] *School of Mathematics, University of Leeds, Leeds, United Kingdom*
[5] *Department of Information Technology, Uppsala University, Uppsala, Sweden*

November 5, 2024

Abstract

Adaptive resistance contributes significantly to treatment failure in many cancers. Despite the increased prevalence of experimental studies that interrogate this phenomenon, there remains a lack of applicable quantitative tools to characterise data, and importantly to distinguish between resistance as a discrete phenotype and a (potentially heterogeneous) continuous distribution of phenotypes. To address this, we develop a stochastic individual-based model of adaptive resistance in low-cell-count proliferation assays. That our model corresponds probabilistically to common partial differential equation models of resistance allows us to formulate a likelihood that captures the intrinsic noise ubiquitous to such experiments. We apply our framework to assess the identifiability of key model parameters in several population-level data collection regimes; in particular, parameters relating to the adaptation velocity and within-population heterogeneity. Significantly, we find that heterogeneity is practically non-identifiable from both cell count and proliferation marker data, implying that population-level behaviours may be well characterised by homogeneous ordinary differential equation models. Additionally, we demonstrate that population-level data are insufficient to distinguish resistance as a discrete phenotype from a continuous distribution of phenotypes. Our results inform the design of both future experiments and future quantitative analyses that probe adaptive resistance in cancer.

Keywords: adaptation, resistance, heterogeneity, inference, identifiability, individual-based model

[1]Corresponding author: browning@maths.ox.ac.uk

1 Introduction

Adaptive resistance is acknowledged as a significant factor in the eventual failure in the treatment of many cancers [1–4]. Such short-term phenotypic adaptation arises in isogenic populations through epigenetics such that cells quickly manifest a reversible drug-tolerant phenotype when exposed to sufficiently high doses of a therapeutic drug [5, 6]. Both experimental [6, 7] and theoretical [8, 9] studies have proposed adaptive therapy and the intermittent delivery of drug to overcome this phenomenon. Mathematical models, in particular, have been proposed to characterise this behaviour; interpret experimental studies of phenotypic adaptation; and to develop treatment schedules robust to resistance [10–16].

Complicating the characterisation of tumour-level plasticity within a given cancer are the seemingly opposing observations that resistance corresponds to a well-defined discrete phenotype [17–20] and to a continuous spectrum of phenotypes [21,22]. Indeed, many mathematical models of resistance describe a heterogeneous population comprising cells that are either firmly drug-sensitive or drug-resistant [18, 23–25]. While mathematical models that capture continuous phenotype adaptation have been proposed [26–28] and are in fact well studied in the partial differential equation (PDE) literature [21,29], they remain largely unvalidated with experimental data. Issues relating to the *identifiability* [30] of mechanisms behind adaptive plasticity from these models are likely to be further exacerbated by the heterogeneity present in even isogenic cell populations [31]. Despite the increased prevalence of experimental studies that interrogate adaptive plasticity, there remains a lack of quantitative tools to calibrate models that capture continuous heterogeneity to data. Thus key questions—such as the data requisite to identify the mechanisms behind adaptive plasticity, and the ability to distinguish between discrete and continuous heterogeneity—remain unanswered.

We are motivated by a recent study of intermittent therapy of mutant melanoma cells by Kavran et al. [6], in which the authors provide genetic evidence for the presence of at least two (reversible) phenotype states: a drug-resistant phenotype and a drug-sensitive phenotype arising within a seven-day period of drug exposure and drug removal, respectively. From reported cell fold-change data from each phenotype, we have previously quantified a dose and phenotype dependent difference in net growth rate (Fig. 1a); a characteristic of high interest in the context of the eventual development of treatments robust to adaptation [32]. Notably, Kavran et al. [6] provide compelling evidence for a continuous distribution of phenotypes present in the intermediate period between resensitisation and addiction through observations of the cell-adhesion marker L1CAM (reproduced in Fig. 1b), a protein well-known as a marker of the epithelial-to-mesenchymal transition and drug resistance in melanoma [33]. While sequence and protein data provide qualitative insight into the adaptive dynamics, their link to cell growth rate is unlikely to be direct. We must, therefore, resort to using cell count data arising from proliferation assay experiments (Fig. 1c) to quantify adaptive dynamics and the corresponding, possibly heterogeneous, net cellular growth rate.

To capture the stochasticity intrinsic to low-cell count experiments such as proliferation assays, we develop an individual-based model (IBM) of drug-based adaptation [27]. We build on

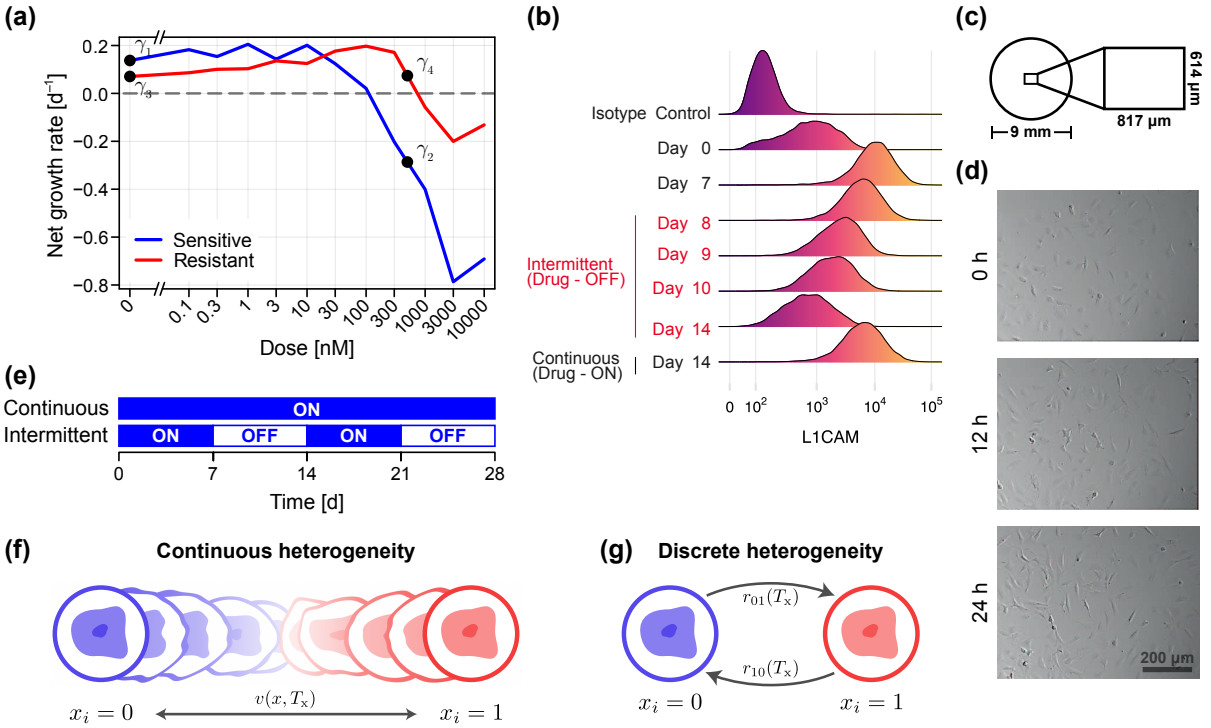

Figure 1. Experimental data of phenotype adaptation. An example suite of experimental data of phenotype adaptation. (a) Kavran et al. [6] expose WM239A melanoma cells to either a continuous treatment, or an intermittant treatment, as shown in (e). Net growth rate of cells calculated for various drug dose levels [32]. Cells that have been exposed to drug in the 7 days prior to measurement are classified as drug-resistant; cells that have not as drug-sensitive. (b) Phenotype characterised experimentally by the expression of L1CAM, a marker for cell adhesion. Day 14 intermittent data (i.e., cells that have not been exposed to drug between days 7 and 14) show a similar profile to day 0. Reprinted from [6] with permission from the author. (c) Schematic of a cell proliferation assay. (d) Example suite of cell proliferation assay data; experiments conducted with a low density of 3T3 Fibroblast cells (reprinted from [34] under a CC-BY license and further analysed in [35]). (e) Cells are subject to either *continuous treatment*, in which a drug concentration is maintained, or to *intermittent treatment*, in which treatment alternates between 7-day periods of drug exposure and drug removal. (f,g) Schematics of continuous and discrete models of phenotypic heterogeneity, respectively (see text for details).

the IBM of Hamis et al. in a stochastic differential equation (SDE) framework to present a model in a continuous phenotypic space where a heterogeneous population of cells tend toward either a drug-sensitive state (mathematically, denoted by $x = 0$) or a drug-resistant state (denoted by $x = 1$). Key to our analysis is that drug-responsive movement throughout the phenotypic space is both stochastic and described, on average, by a function of potentially unknown analytical form. The choice to work in an SDE framework means that our IBM corresponds precisely in a probabilistic sense to common PDE models of phenotypic adaptation [36].

Exploiting the tractability of the SDE and analogous PDE model, we build an inference framework that captures intrinsic noise in low-cell-count proliferation assay experiments without the pervasive, but often unjustified, assumption that experimental observations of cell count are subject to additive Gaussian noise. To do this, we derive and present a chemical master equation (CME) that describes the time-evolution of cell count, with which we construct a likelihood function [37]. We then perform inference and identifiability analysis under three data collection scenarios. First, we consider a suite of cell proliferation experiments for cells that are initially either resistant or sensitive and are exposed (or not) to a drug over a seven day period. Second, we consider a hypothetical scenario in which proliferation assays are observed continuously such that the time of cell-proliferation and cell-death events are directly observed. Lastly, we consider another hypothetical scenario in which a cell proliferation marker (i.e., L1CAM) correlates weakly, but linearly, with cell proliferation. All analysis is initially conducted in an idealised scenario where the functional form of the phenotype adaptation mechanism is correctly specified. We later relax this assumption and perform model selection.

Together, the data collection regimes we study establish the identifiability of individual model parameters and, more importantly, our ability to distinguish phenotypic heterogeneity from possible model misspecification. Given that the entrenched model of phenotypic heterogeneity is that of two well-defined discrete phenotypes, we conclude our study by investigating whether such discrete heterogeneity can be distinguished, using cell count data alone, from a model comprising a continuous phenotypic space. As we are primarily interested in establishing the theoretical identifiability of model parameters and mechanisms, in the main text we focus our analysis on regimes where cell counts are extracted from images precisely; we relax this assumption in the supplementary material by investigating where identifiability is lost if only imprecise measurements can be made.

2 Mathematical methods

2.1 Individual-based model

We assume that individuals undergo a biased random walk in phenotype space, such that the phenotype of a cell i, denoted x_i, is given by

$$\mathrm{d}x_i = v(x_i, T_{\mathrm{x}})\,\mathrm{d}t + \beta\,\mathrm{d}W_i, \tag{1}$$

where $T_{\mathrm{x}} \in \{0, 1\}$ indicates untreated and treated states, respectively; $v(x_i, T_{\mathrm{x}})$ describes the adaptation velocity; β describes the magnitude of diffusive movement thoughout the phenotype space, and W_i is a Wiener process. We further assume that, for $\beta = 0$, the system has a stable steady state at $x_i = 0$ for $T_{\mathrm{x}} = 0$ (this is referred to as the sensitive state), and likewise at $x_i = 1$ for $T_{\mathrm{x}} = 1$ (referred to as the resistant state).

We take net cellular growth rate to be a linear function of x [32], given by

$$\lambda(x, T_{\mathrm{x}}) = \begin{cases} \gamma_1 + (\gamma_3 - \gamma_1)x & \text{if } T_{\mathrm{x}} = 0, \\ \gamma_2 + (\gamma_4 - \gamma_2)x & \text{if } T_{\mathrm{x}} = 1, \end{cases} \tag{2}$$

as shown in Fig. 1a. Provided that the growth rate is monotonic in x, the functional form of λ is arbitrary since we could, in theory, rescale the phenotypic space in Eq. (1) and thus equivalently the functional form of v. Furthermore, we follow [32], and assume that $\lambda(x, T_{\mathrm{x}}) < 0$ corresponds solely to net death (apoptosis or necrosis), and $\lambda(x, T_{\mathrm{x}}) > 0$ corresponds solely to net proliferation. We further assume that both proliferation and death events occur according to a Poisson process. Upon death, a cell is removed from the population. Upon proliferation, a cell is replicated such that daughter cells are created with an (initially) identical phenotype index to the parent.

While we focus on analysis of synthetic data, we choose biologically realistic parameters based upon analysis on the emergence of reversible resistance to the BRAF-inhibitor LGX818 in $\mathrm{BRAF}^{\mathrm{V600E}}$-mutant melanoma cells [6, 32]. The growth rate parameters are chosen to be $\gamma_1 = 0.15$, $\gamma_2 = -0.3$, $\gamma_3 = \gamma_4 = 0.1$ to approximately match the mean growth rate of sensitive and resistance cells under drug and no drug conditions (see dose-response curve in Fig. 1a). Very little information is available regarding the adaptation dynamics through $v(x_i, T_{\mathrm{x}})$ and diffusivity β, other than the qualitative observation that cells move between drug-sensitive and drug-resistant states within a seven day window. For preliminary analysis we set $v(x_i, T_{\mathrm{x}}) = -\nu(x_i - T_{\mathrm{x}})$ with $\nu = 0.4$ such that x_i is an Ornstein-Uhlenbeck process. Further, we set $\beta = 0.05$ such that the stationary distribution of sensitive cells has a standard deviation of approximately 0.05. Implicit in our model is an assumption that the mechanisms behind drug-sensitisation and the reverse are identical. However, this need not be the case as we later exposit: it is sufficient to study identifiability in a single direction.

We set the initial condition in the model to a probabilistic representation of a spatially uniform low-cell count proliferation assay experiment; specifically, a cell proliferation assay conducted in a standard 9 mm well initialised with approximately 1000 cells (this is slightly larger than the initial population in [6]). The field-of-view of the imaged proliferation assay in Fig. 1c–d is 817×614 µm, and so each cell has probability $\rho = 817 \times 614/(4500^2\pi)$ of presenting in the field-of-view. The initial condition is thus set to $n_0 \sim \mathrm{Binomial}(1000, \rho)$, corresponding to a mean initial cell count of approximately 7.9 per image.

In Fig. 2a–d, we simulate a set of synthetic cell proliferation assay experiments with our IBM under both continuous and intermittent treatments; the latter is defined as alternating 7-day periods of drug and no drug (Fig. 1d). Results in Fig. 2a,b highlight emergent isogenetic heterogeneity due to white-noise driven fluctuations in the phenotype index. Results in Fig. 2c,d

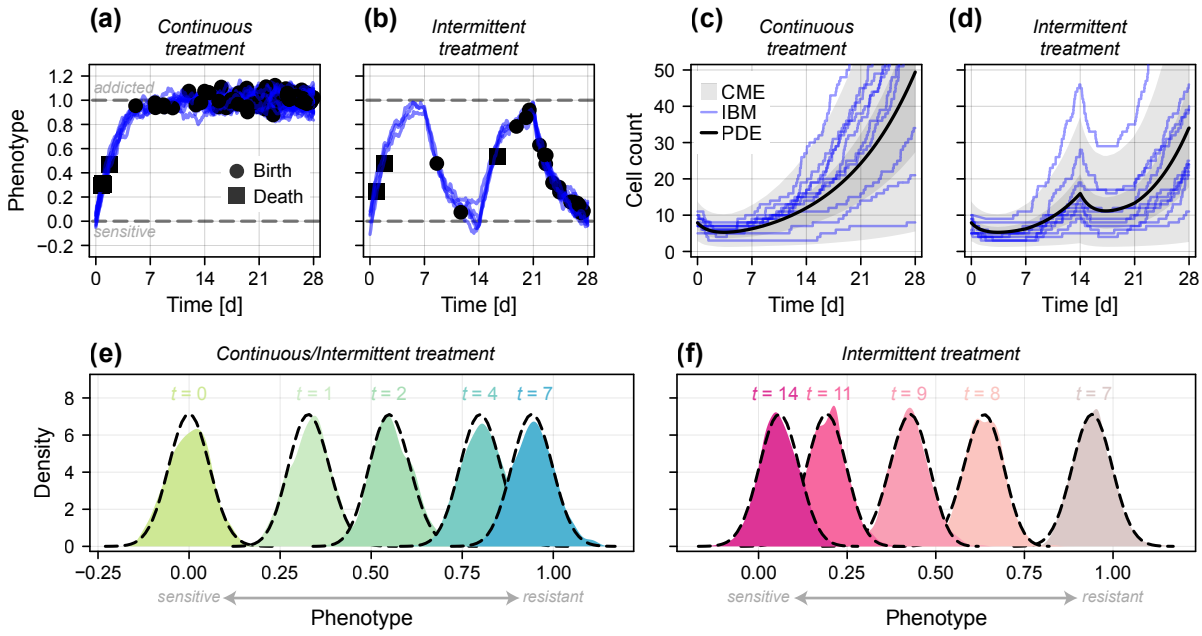

Figure 2. Model comparison. We compare realisations of the SDE-based IBM to the solution of both the corresponding Fokker-Planck PDE (Eq. (3)) and the CME (Eq. (11)) for the population size. (a–b) A single realisation of an IBM initiated with drug sensitive cells. The mean phenotype is zero in the equilibrium drug-sensitive state, and unity in the equilibrium drug-resistant state. Treatment applied is (a) continuous; and (b) intermittent (Fig. 1e). (c–d) Cell count observations from 10 realisations of the IBM (blue) under both (c) continuous and (d) intermittent treatment. Also shown is the expected population $\langle N \rangle(t)$ computed from a numerical solution of the PDE (black), and both a 50% and 95% credible region computed from a numerical solution of the CME (grey). A full comparison between the solution of the CME and the IBM is provided as supplementary material (Fig. S1). (e–f) Comparison between the phenotypic distribution computed empirically using an IBM initiated with 500 cells (coloured) and from the PDE (black dashed). Results in (e) show the phenotype distribution for both continuous and intermittent treatment for $0 \leq t \leq 7\,\mathrm{d}$ (in which both regimes are identical) and in (f) for intermittent treatment from $7 \leq t \leq 14\,\mathrm{d}$.

show high levels of stochasticity in cell count. Since the simulations are discrete, there is a non-zero probability of extinction as our model does not, in its standard formulation, consider migration into and out of the cell proliferation assay field-of-view (Fig. 1c).

2.2 Partial differential equation model

We now define $u(x,t)$ as the density of cells with phenotype x at time t, such that the dynamics of $u(x,t)$ are governed by the Fokker-Planck equation

$$\underbrace{\frac{\partial u(x,t)}{\partial t} + \frac{\partial}{\partial x}\Big(v(x,T_{\mathrm{x}}(t))u(x,t)\Big) = \beta\frac{\partial^2 u(x,t)}{\partial x^2}}_{\text{Fokker-Planck equation}} + \underbrace{\lambda(x,T_{\mathrm{x}}(t))u(x,t)}_{\text{Proliferation and Death}}\,, \tag{3}$$

subject to the usual set of no-flux and vanishing far-field boundary conditions [27, 36, 38].

Experiments are initiated with a sample of cells from a zero-net-growth stationary distribution for either a sensitive or resistant population. For the Ornstein-Uhlenbeck formulation of

$v(x, T_{\mathrm{x}})$, this corresponds approximately to

$$x_i(0) \sim \mathcal{N}(0, \beta/\sqrt{2\nu}), \tag{4}$$

which we set as the initial condition in the model.

We denote by $\langle N \rangle(t)$ the expected cell count and by $p(x,t)$ probability density function (PDF), given by

$$\langle N \rangle(t) := \int_{\mathbb{R}} u(x,t)\,\mathrm{d}x, \tag{5}$$

and

$$p(x,t) := \frac{u(x,t)}{\langle N \rangle(t)}, \tag{6}$$

respectively.

In Fig. 2e–f, we compare a finite-difference approximation to the PDE (see supplementary code) to a set of realisations of the IBM initialised with a large ($n_0 = 500$) number of initial cells. We remind the reader that we expect a close match (that converges as $n_0 \to \infty$), as the PDE is an exact probabilistic representation the IBM.

2.3 Chemical master equation

We now derive an approximate master equation for the time-evolution of the probability mass function for the cell count, defined as

$$q(n,t) := \mathbb{P}(N(t) = n). \tag{7}$$

We consider that

$$\begin{aligned} q(n, t+\delta t) = q(n,t) &+ \mathbb{P}\big(\text{proliferation in } (t, t+\delta t) | N(t) = n-1\big) q(n-1, t) \\ &+ \mathbb{P}\big(\text{death in } (t, t+\delta t) | N(t) = n+1\big) q(n+1, t) \\ &- \mathbb{P}\big(\text{proliferation in } (t, t+\delta t) | N(t) = n\big) q(n, t), \\ &- \mathbb{P}\big(\text{death in } (t, t+\delta t) | N(t) = n\big) q(n, t). \end{aligned} \tag{8}$$

Note that we can also include terms in the above that explicitly capture migration into and out of the field of view. Generally, however, we would expect these to vanish if we assume that the assay as a whole is sufficiently homogeneous such that migration out of the window occurs at the same rate as migration into the window (i.e., periodic boundary conditions).

To make progress, we assume that the phenotypic states of cells are independent. While not strictly true for very high proliferation and death rates (since cells inherit their phenotype from a parent), this is appropriate for the range of growth rates we observe (Fig. 1a). Under these assumptions, the per-capita instantaneous proliferation and death rates are given by

$$r_{\mathrm{prol}}(t) = \int_{\mathbb{R}} p(x,t) \max(0, \lambda(x,t))\,\mathrm{d}x, \tag{9}$$

and

$$r_{\text{death}}(t) = -\int_{\mathbb{R}} p(x,t)\min(0,\lambda(x,t))\,\mathrm{d}x, \tag{10}$$

respectively.

For δt sufficiently small, we can consider a Taylor expansion of the exact Poisson probability to obtain an asymptotic expression for the event probabilities in Eq. (8). These are given by

$$\mathbb{P}\big(\text{proliferation in } (t,t+\delta t)|N(t)=n\big) \sim n r_{\text{prol}}(t)\delta t,$$

and

$$\mathbb{P}\big(\text{death in } (t,t+\delta t)|N(t)=n\big) \sim n r_{\text{death}}(t)\delta t.$$

Substituting into the difference equation (Eq. (8)) and taking $\delta t \to 0$, we arrive at the CME

$$\begin{aligned}\frac{\mathrm{d}q(n,t)}{\mathrm{d}t} &= (n-1)r_{\text{prol}}(t)q(n-1,t)\\ &\quad - n\big(r_{\text{prol}}(t)+r_{\text{death}}(t)\big)q(n,t)\\ &\quad + (n+1)r_{\text{death}}(t)q(n+1),\end{aligned} \tag{11}$$

subject to absorbing boundaries such that $q(n)=0$ for $n<0$.

In Fig. 2c–d, we compare the solution of the CME to realisations of the IBM, showing that the CME captures both the average and variance of the cell count. A more detailed comparison is provided in Fig. S1.

2.4 Likelihood-based inference

We take a Bayesian approach to parameter estimation and identifiability analysis and apply the CME (Eq. (11)) to construct a likelihood for cell count data reported from proliferation assays. The advantage of this approach, compared to a more standard approach that considers an average cell count subject to additive Gaussian noise, is that we account directly for the stochasticity intrinsic to the proliferation death process. As we are primilarly interested in the identifiability of model parameters, we assume that all cell counts are exact. In the supplementary material, we investigate identifiability in the case that experimental observations are potentially subject to miscounting.

Experiments are conducted for t days, at the conclusion of which a cell count observation is taken. We denote by $n_k^{(t,T_{\text{x}},\mathcal{P})}$ the kth cell count taken from an experiment terminated at time t, conducted entirely with ($T_{\text{x}}=1$) or without ($T_{\text{x}}=0$) drug, using an initial population of sensitive (denoted $\mathcal{P}=0$) or resistant (denoted $\mathcal{P}=1$) cells, and denote by $\mathcal{D}$ the complete set of data. Further denoting the solution to the CME with conditions $(\mathcal{P},T_{\text{x}})$ and parameter values $\boldsymbol{\theta}$ by $q_{(\mathcal{P},T_{\text{x}})}(n,t;\boldsymbol{\theta})$, the log-likelihood is given by

$$\ell_{\mathcal{D}}(\boldsymbol{\theta}) = \sum_{(\mathcal{P},T_{\text{x}})}\sum_{t}\sum_{k}\log q_{(\mathcal{P},T_{\text{x}})}\big(n_k^{(t,T_{\text{x}},\mathcal{P})},t;\boldsymbol{\theta}\big). \tag{12}$$

Note that we have assumed that cell count observations are independent between time points;

effectively assuming that measurements are taken at the termination of an experiment and not as a time-series. While our approach could be trivially extended to account for time-series data, this would add significant computational cost by potentially requiring a numerical solution to the CME for each individual observation. While we focus our results on inference using cell count data, we also consider log-likelihood functions constructed for two other data types: event timing data (i.e., the exact time of proliferation or cell death observed from temporal data) and from a cell proliferation marker that may linearly correlate with the net growth rate.

Following the construction of the log-likelihood function, we can either take a frequentist approach and find the maximum likelihood estimate (MLE), or apply a Bayesian approach to quantify identifiability and parameter uncertainty. While unusual to consider both approaches, we do so in this work as the former is advantageous as it allows us to perform model selection using frequentist hypothesis tests [39].

For the latter, we assume that knowledge about model parameters is initially encoded in a *prior distribution*, $p(\boldsymbol{\theta})$. We choose $p(\boldsymbol{\theta})$ to be independent uniform over a sufficiently wide range of parameter magnitudes (full details are given as supplementary material). This choice also ensures that the maximum *a posteriori* estimate (MAP) corresponds to the MLE. Following a set of observations, denoted by $\mathcal{D}$, arising from cell proliferation assay measurements, or otherwise, we update our knowledge about the model parameters using the relevant likelihood denoted $\ell_{\mathcal{D}}$ to obtain the *posterior distribution*, given by

$$p(\boldsymbol{\theta}|\mathcal{D}) \propto \exp(\ell_{\mathcal{D}}(\boldsymbol{\theta}))p(\boldsymbol{\theta}). \tag{13}$$

When applying the Bayesian approach, we sample from the posterior using an adaptive Markov-Chain Monte Carlo algorithm [40]. To obtain MLEs we apply the DIRECT global search algorithm implemented in NLopt for Julia [41]. As we are primarily interested in parameter and model identifiability, we initiated each chain using the "true" set of parameter values that are used to generate synthetic data.

3 Results

3.1 Phenotypic heterogeneity is poorly identified from cell count data

We begin our analysis by considering a suite of synthetic cell proliferation assays conducted within a seven day period (specifically, a set of assays that terminate at $t = 1$, 3, 5, and 7 d). For each termination time, we conduct a set of four experiments: with or without drug and initiated with either a population of fully sensitive or resistant cells. We devote two 96-well plates to each termination time, such that the sample size for each condition is $M = 48$. The duration is chosen based on the observation that the population adapts or resensitizes within a seven day interval [6] (in the supplementary material, we consider a variety of termination time sets).

Applying the CME-based Bayesian inference procedure reveals that all growth rate parameters are practically identifiable. The results in Fig. 3a–b show how model predictions produced

at the MAP align with synthetic cell count data observations (additionally, the model also captures higher order moments of cell count, although this is not shown). Furthermore, results in Fig. 3c show that the adaptation speed parameter, ν, is identifiable. However, we see from results in Fig. 3d that the diffusion parameter β, which corresponds to the variance in the phenotype variable x within an adapting population, is only *one-sided identifiable*: we can establish an upper bound, but no lower bound. The parameter is, however, *structurally identifiable*: we show this in the supplementary material using a significantly larger ($M = 768$) data set, however the parameter becomes again non-identifiable when imprecise cell-count observations are made (Figs. S4 and S5). Thus, we expect, from cell count data alone, models with phenotypic heterogenity to be indistinguishable from models with deterministic adaptation (the $\beta = 0$ scenario).

To investigate the identifiability of β further, we recall that the phenotype distribution, $p(x,t)$, affects overall cell count dynamics only indirectly. Specifically, cell proliferation and death is governed at the population-level by the overall proliferation and death rates, given by Eqs. (9) and (10). For $\beta = 0$, $p(x,t)$ tends to a degenerate distribution such that $r_{\text{prol}}(t) = \lambda(\bar{x},t)$, where $\bar{x}$ is the mean phenotype (in the case of homogeneity, the only phenotype). The most obvious consequence of the $\beta = 0$ parameter regime is that proliferation and death cannot occur simultaneously: thus, we expect a sharp proliferation-death transition at $r_{\text{prol}}(t) = 0$ as the population switches between death and proliferation events, depending on the presence of drug and the mean phenotype. In contrast, the transition at $r_{\text{prol}}(t) = 0$ will be diffuse for non-zero β. In Fig. 4, we compare the event rates for various values of β. Clearly, aside from minor differences at the proliferation-death transition, rate curves are visually indistinguishable for decreasing values of β. For large β, which has very little or no posterior mass (see Fig. 3d), the proliferation rate curve becomes distinguishable.

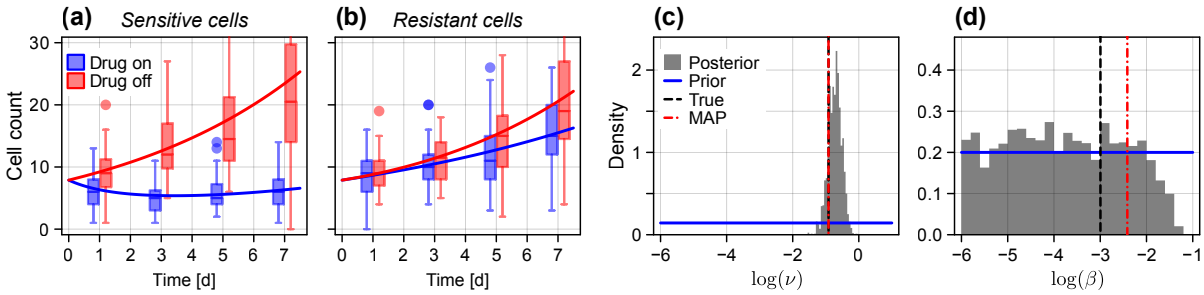

Figure 3. Proliferation assay inference. We perform Bayesian inference on a set of synthetic cell proliferation assay data using the CME as a likelihood. Independent cell count observations ($M = 48$ replicates per condition) are collected from experiments conducted with fully sensitive or fully resistant cells, with and without drug, and terminated at $t = \{1\,\text{d}, 3\,\text{d}, 5\,\text{d}, 7\,\text{d}\}$. (a–b) Synthetic proliferation assay cell count data (box plots) and the model predicted mean cell count at the MAP (solid lines). (c–d) Posterior distributions for the logarithms of v, the adaptation speed, and β, the diffusivity. Shown also is the uniform prior (blue), the true value (black dashed), and the MAP (red dashed). While the adaptation speed is identifiable (as are all other parameters; not shown), the diffusivity is only one-sided identifiable; the model cannot be distinguished from that with purely deterministic adaptation (i.e., no heterogeneity).

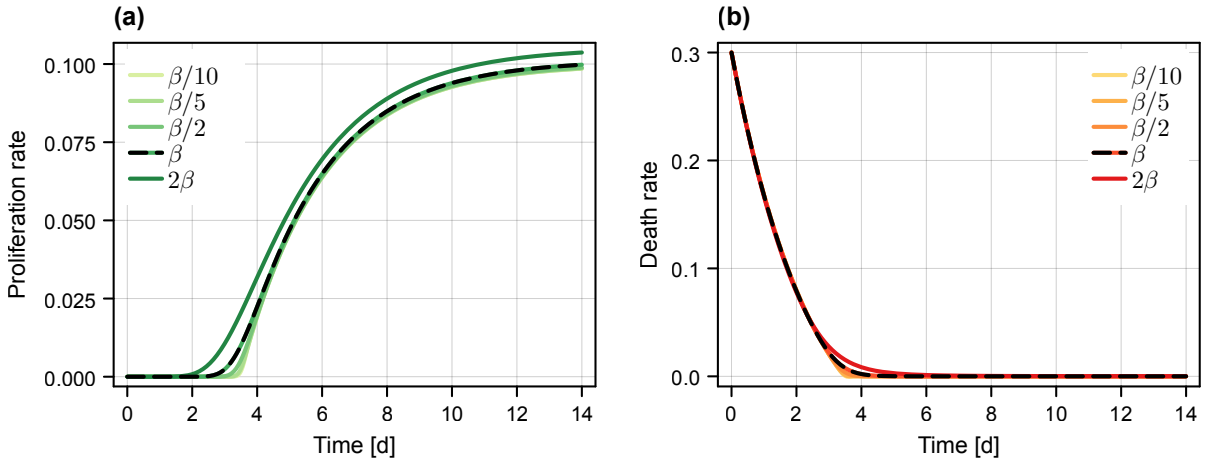

Figure 4. Practical non-identifiability from cell count data. Practical non-identifiability of the diffusion parameter β (corresponding to a measure of the heterogeneity), seen through differences in the total expected proliferation and death rate functions, $r_{\text{prol}}(t)$ and $r_{\text{death}}(t)$, respectively. All other parameters are fixed at their true values. Clear differences are seen in the proliferation rate between β and a proliferation rate constructed where $\beta \mapsto 2\beta$; we have seen previously that β is one-sided identifiable. However, reducing β shows (visually) very minor differences between the proliferation and death rate functions as both tend to the deterministic limit (in this case, corresponding to exponential decay from the negative maximum death rate through to the maximum proliferation rate. In the absence of heterogenity (i.e., for $\beta \to 0$), proliferation and death events cannot occur simultaneously in a population.

3.2 Phenotype heterogeneity is identifiable from event-timing data

Under the current model formulation, in which heterogeneity is driven soley by diffusion through the phenotypic space, it is only in the regime where $\beta > 0$ that we will ever see proliferation and death events occur simultaneously. Thus, in the constraints of our model formulation, we expect to be able to more precisely identify heterogeneity if we observe the precise timings of cellular proliferation and death events from, for example, live cell imaging.

We therefore investigate a hypothetical scenario where we have access to noise-free event timings from a set of proliferation assays that are initiated with a total of 10,000 cells. Without loss of generality, for the rest of the study we focus only on adaptation in the forward direction (i.e., from drug-sensitive to drug-resistant), since an analogous analysis could be conducted in the reverse direction. A log-likelihood function can be constructed by discretising the resultant Poisson process such that the number of proliferation and death events occuring in the interval $(t, t+\delta t)$, denoted $E^{(\delta t)}_{\text{prol}}(t)$ and $E^{(\delta t)}_{death}(t)$, respectively, are distributed according to

$$E^{(\delta t)}_{\text{prol}}(t) \sim \text{Po}\big(r_{\text{prol}}(t)N(t)\delta t\big), \qquad E^{(\delta t)}_{\text{death}}(t) \sim \text{Po}\big(r_{\text{death}}(t)N(t)\delta t\big), \tag{14}$$

where $N(t)$ is the (observed) cell population at time t. Under the well-mixed phenotype assumption for which the CME applies, Eq. (14) is exact as $\delta t \to 0$. We choose $\delta t = 0.035$ such that the seven-day experiment is subdivided into 200 observation intervals (as a consequence, one could also consider event-timing data that is not exact, but accurate to intervals of width δt that correspond to a finite imaging frequency).

The synthetic data set is shown in Fig. 5a, along with an estimate for the instanteous event rate constructed using a moving average. Visually, heterogeneity can be detected by comparing

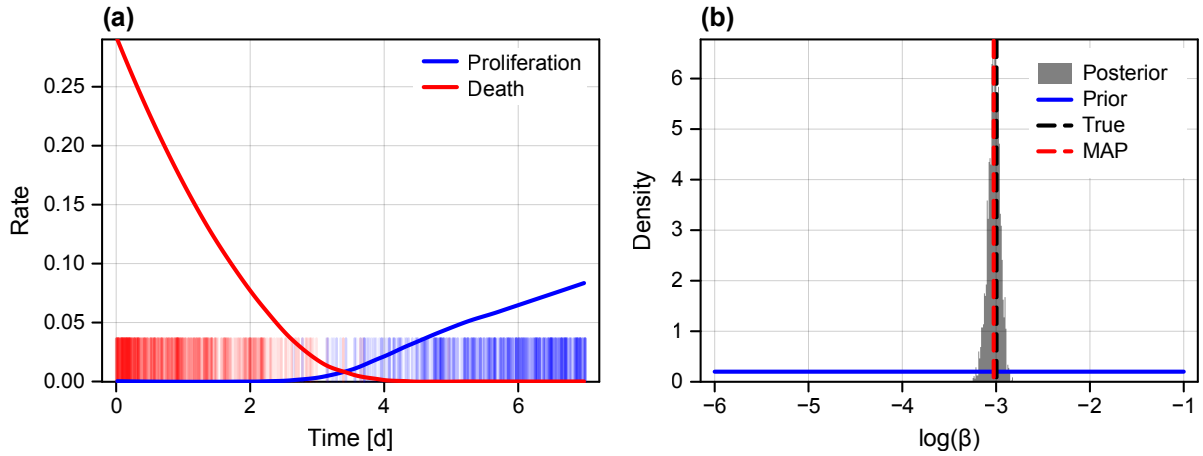

Figure 5. Identifiability of heterogeneity from event timing data. We generate a synthetic data set from an experiment (or set of experiments) that are initiated with a total of 10,000 cells that are under continuous treatment. The exact event timings (i.e., time of cell proliferation, and time of cell death) are recorded and used for inference. (a) Synthetic event timing data. Shown is a rug plot of a sample of 500 each of proliferation and death events, and a local regression (LOESS) of the observed proliferation and death rate. (b) Posterior distribution for $\log(\beta)$, previously non-identifiable, constructed using a Poisson likelihood for the exact timing data. Shown also is the uniform prior (blue), the true value (black dashed), and the MAP (red dashed).

the transition from primarily cell proliferation to primarily cell death to the results shown previously in Fig. 5a. We proceed to perform inference on this synthetic data set using the Poisson likelihood, with the posterior shown for β in Fig. 5b (all other relevant parameters remain identifiable). Clearly, heterogeneity is now identifiable; estimates of β can be drawn precisely.

3.3 Phenotype heterogeneity is not identifiable from proliferation marker data

Our study is in part motivated by Kavran et al. [6] who provide compelling evidence for a continuous transition from a sensitive to resistant state through the cell-adhesion marker L1CAM. Such data are difficult to interpret directly due to uncertainty in the precise link between the net growth rate and the expected marker expression and the resultant flow cyometry measurement. Challenges aside, we now consider identifiability of β in the case that the measured marker expression correlates linearly with the proliferation rate (and effectively, since the link between the net growth rate and phenotype index is also linear, the phenotype index).

We assume that the observed marker expression for a cell, M_i, is given by

$$M_i \sim x_i + \varepsilon \tag{15}$$

where ε is independent of x_i. We consider both that ε is normally distributed with zero mean and unknown standard deviation σ, and a scenario where the shape of ε is additionally unknown such that ε is given by a translated Gamma distribution with zero mean, unknown standard deviation, and unknown skewness ω (this distribution becomes normal as $\omega \to 0$) [42]. By convoluting the distribution of ε with that for x, we can construct an exact log-likelihood for

a set of marker data. In Fig. 6a, we show the resultant (weak) linear correlation between phenotype index and marker measurement.

We fix all other mechanistic parameters, which we previously established to be identifiable from cell count data, at the corresponding true values. We then consider a synthetic data set in which marker data is taken from a set of proliferation assays terminated at $t = \{1\,\mathrm{d}, 3\,\mathrm{d}, 5\,\mathrm{d}, 7\,\mathrm{d}\}$. Results in Fig. 6b show samples from the joint posterior distribution for $\log(\sigma)$ and $\log(\beta)$ in the case that $\omega = 0$. In both the case where the marker error shape is known ($\omega = 0$) and unknown, we are unable to place a lower bound on β. Furthermore, the shape of the posterior in Fig. 6b indicates that, even if we had knowledge of σ, β would remain non-identifiable. We conclude that, from marker data, cellular heterogeneity is indistinguishable from marker noise.

3.4 Model selection and misspecification for cell-count data

We have made two significant observations thus far: first, that the $\beta > 0$ regime is indistinguishable from the $\beta = 0$ regime from cell count data; and second, that all other model parameters are identifiable given a correctly specified model. As a consequence of the first observation, we perform all remaining analysis using what we term the "homogeneous model": an ordinary differential equation (ODE) model given by Eq. (1) with $\beta = 0$ in place of the PDE model. In light of the second observation, we assess whether we can not only identify model parameters, but also the functional form of the adaptation velocity $v(x, T_\mathrm{x})$ (without loss of generality in the case that $T_\mathrm{x} = 1$ such that $v(x, 1) := v(x)$).

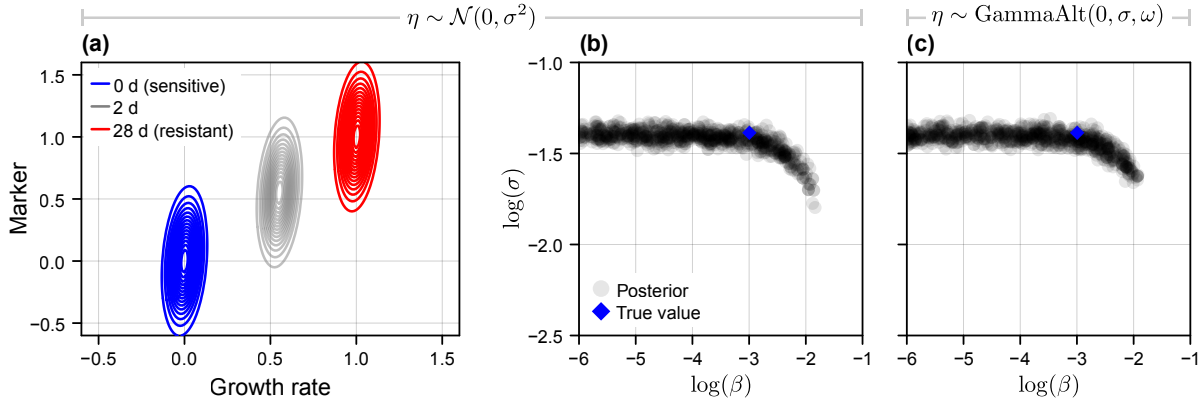

Figure 6. Identifiability of heterogeneity from noisy marker data. We generate a synthetic data set comprising noisy measurements of the phenotype state of each cell using a hypothetical marker for cell proliferation (i.e., L1CAM). (a) Measurements are normalised such that the mean of fully sensitive cells is approximately zero, and that of fully resistant cells (which arise in the limit as $t \to \infty$) is approximately unity. The marker is assumed to weakly linearly correlate with growth rate (and hence, the phenotype index); we model this by a measurement noise process that is normally distributed with variance σ^2. (b) We perform Bayesian inference on a dataset generated from cell proliferation assays with fully sensitive cells, exposed to drug continuously, with independent measurements taken at $t = \{1\,\mathrm{d}, 3\,\mathrm{d}, 5\,\mathrm{d}, 7\,\mathrm{d}\}$ ($M = 48$ replicates per measurement time). All other parameters, identifiable from cell count data, are fixed, and σ and β are estimated, with the joint posterior (grey discs) shown alongside the true value (blue diamond). (c) We repeat the analysis in the case that the shape (skewness, quantified by ω) of the measurement noise distribution is additionally unknown.

We consider that $v(x)$ is given by

$$v(x) = a\text{sgn}(x) + (1-x)(b + cx + dx^2), \tag{16}$$

and consider cell proliferation assays conducted with drug sensitive cells under continuous treatment and terminated at $t = \{1\,\text{d}, 3\,\text{d}, 5\,\text{d}, 7\,\text{d}\}$ ($M = 48$ per condition). The true model is given by setting $a = c = d = 0$. We can recover a variety of velocity models using the functional form given by Eq. (16), including for $b = c = d = 0$ the constant adaptation presented in our previous work [32]. As the growth rate parameters for the drug-on experiment, γ_2 and γ_4, were found to be identifiable (and can be established by conducting drug-off and drug-on experiments with sensitive and resistant cells, respectively) we fix each to their corresponding true value.

We perform model selection using the frequentist likelihood ratio test (equivalent to profile likelihood). For example, to test whether $a = 0$, we compare the likelihood at the MLE (equivalently, the MAP) for $a = 0$, denoted by $\hat{\boldsymbol{\theta}}_{[a]}$ to that for the model where all parameters in Eq. (16) are non-zero, denoted by $\hat{\boldsymbol{\theta}}$. Figure 7a shows the resultant set of log-likelihoods, translated such that $\ell(\hat{\boldsymbol{\theta}}) = 0$. From the likelihood ratio test [39], we can construct a threshold based on a 95% confidence interval outside of which we reject a null hypothesis that the parameter set, i.e., $[a]$ is equal to zero.

Results in Fig. 7a show that any individual parameter can be set to zero. Furthermore, any pair of parameters can be set to zero *except* a and b simultaneously. Finally, only the parameter triples that do not contain *both* a and b can be set to zero. If the goal was to identify a single model, one would use an information criterion [39] (or similar) to penalise differences in log-likelihood by the dimensionality of the non-zero parameter set; in our case, we expect a model where only one of a or b is non-zero as the most parsimonious.

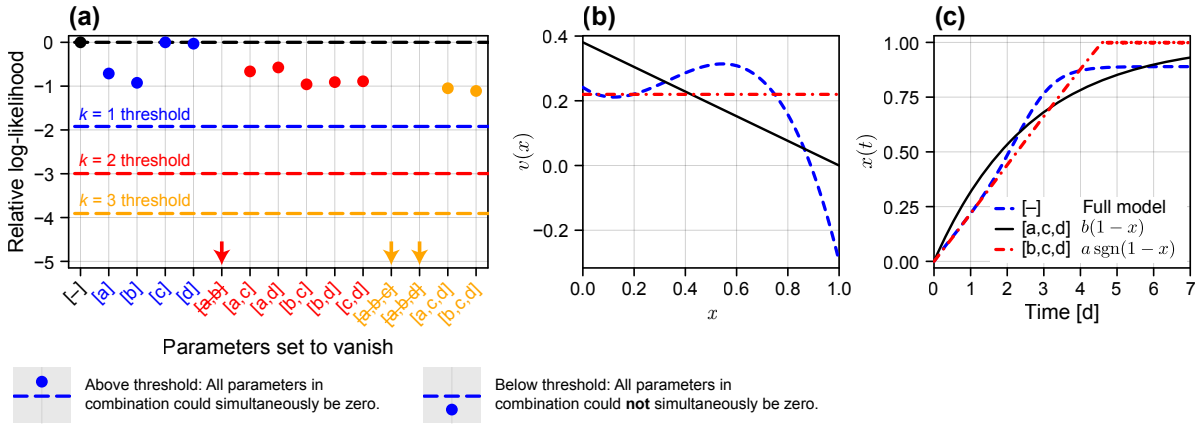

Figure 7. Model selection and misspecification. We perform inference and model selection on a general adaptation velocity function of form given by Eq. (16). The true model corresponds to $b = 0.4$ and $a = c = d = 0$ (i.e., the combination $[a, c, d]$). (a) Results from a likelihood ratio test where the null hypothesis in each column is that the stated parameter combination $[\cdot]$ is zero. Relative log-likelihood values below the relevant threshold (colours correspond to different dimensionalities) indicate that the null hypothesis can be rejected at the level of a 95% confidence interval. Arrows indicate that observed statistics are below the plotted region. (b,c) Identified possible adaptation velocities and phenotype transitions respectively.

Our analysis has identified a family of possible adaptation velocity functions, given by the MLE for each combination for which the relative log-likelihood in Fig. 7a is above the corresponding threshold. In Fig. 7b we compare the identified adaptation velocities for the true model (b non-zero) to the full model and a model where only a is non-zero. Clearly, there remains large uncertainty as to the functional form of $v(x)$ throughout the phenotype space. Results in Fig. 7c, however, demonstrate why these differences do not manifest in statistically different cell count observations: while $v(x)$ varies significantly, the possible paths for $x(t)$ are similar.

3.5 Continuous and discrete-binary heterogeneity may be indistinguishable

Arguably the standard model of plasticity describes a drug-dependent switch between two discrete phenotypes: sensitive and resistant. Such an analogue of our model is

$$X_0 \underset{r_{10}}{\overset{r_{01}}{\rightleftharpoons}} X_1, \tag{17}$$

where sensitive cells, X_0, have net growth rate $\tilde{\lambda}_0$, and resistant cells, X_1, have net growth rate $\tilde{\lambda}_1$ (Fig. 1g). We assume that r_{10} and r_{01} are drug-dependent, and focus our analysis on adaptation from the sensitive to the resistant state.

As Eq. (17) is linear, the mean cell count in each subpopulation, denoted by $n_0(t)$ and $n_1(t)$, is given by

$$\begin{aligned} \frac{\mathrm{d}n_0}{\mathrm{d}t} &= \tilde{\lambda}_0 n_0 - r_{01} n_0 + r_{10} n_1, \\ \frac{\mathrm{d}n_1}{\mathrm{d}t} &= \tilde{\lambda}_1 n_1 + r_{01} n_0 - r_{10} n_1. \end{aligned} \tag{18}$$

To draw a correspondence to the continuous model, we consider now the mean $\tilde{x}(t) := n_1(t)/(n_0(t)+n_1(t))$, which we expect to correspond with $x(t)$ in the continuous model (although not exactly, as in general at equilibrium $\tilde{x}(t) \neq 1$). The dynamics of $\tilde{x}(t)$ are governed by

$$\frac{\mathrm{d}\tilde{x}}{\mathrm{d}t} = r_{01} + \big(\tilde{\lambda}_1 - \tilde{\lambda}_0 - r_{01} - r_{10}\big)\tilde{x} + \big(\tilde{\lambda}_0 - \tilde{\lambda}_1\big)\tilde{x}^2 = A + B\tilde{x} + C\tilde{x}^2. \tag{19}$$

Thus, we expect the average cell count in the discrete model to correspond exactly to the average cell count in a continuous model with a quadratic adaptation velocity. We cannot make an equivalent statement for higher order moments, however we can define an exact CME for the evolution of the joint density $\tilde{q}(n_1, n_2, t) := \mathbb{P}(N_1(t) = n_1, N_2(t) = n_2)$ and hence the probability mass $\tilde{q}(n, t) := \mathbb{P}(N_1(t) + N_2(t) = n)$ in the discrete model (supplementary material).

For a given set of discrete model parameters, we compute a rescaled velocity function and set of continuous model net growth rates such that both models have equivalent initial and fully adapted net growth rates. In Fig. 8a, we demonstrate that the mean cell counts are identical between models; from average cell count data, and by extension large-cell-count proliferation assays, we cannot distinguish a discrete model from a continuous model with quadratic adaptation velocity. Results in Fig. 8b–d demonstrate (subtle) differences in higher-order moments and the mass function for each model. We conclude, therefore, that within our modelling framework

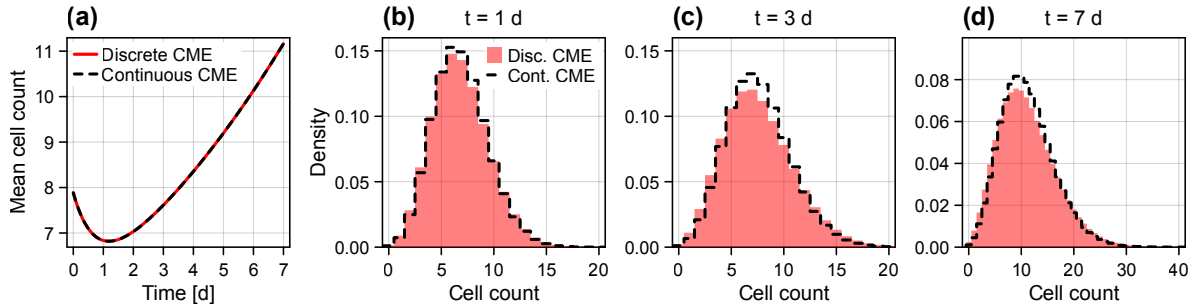

Figure 8. Continuous and discrete-binary models are only distinguishable from higher order moments. We compare the solutions of the CME for a discrete-binary model and a continuous model with quadratic adaptation velocity. (a) Exact correspondence in the mean cell count for each model; the regimes are structurally non-identifiable. (b–d) Solution to the complete CME at various time points, showing higher-order differences between the models. Discrete model parameters are give by $r_{01} = 1$, $r_{10} = 0.01$, $\hat{\lambda}_0 = -0.3$, $\hat{\lambda}_1 = 0.1$.

it may be possible to distinguish between the discrete and continuous models using higher order moments in low-cell-count proliferation assays; however this is unlikely to be the case if only imprecise cell count observations are available.

4 Discussion and Conclusion

Phenotype plasticity and the rapid adaptation of cells upon the application of treatment are widely recognised as a significant factor in the failure of many anti-cancer treatments [43]. Complicating a comprehensive characterisation of adaptive resistance is a lack of consensus as to whether adaptation occurs between a set of well-defined discrete cell states or across a continuous spectrum of phenotypes. While both hypotheses are associated with mature subsets of the mathematical modelling literature, there remains—particularly for the latter—a dearth of statistical methodology to parameterise such models. Indeed, key questions relating to the identifiability of adaptation mechanisms and within-population heterogeneity, the ability of practitioners to distinguish between discrete and continuous adaptation, and the experimental design requirements to parameterise models, remain unanswered.

Our most significant result is that we are unable to identify heterogeneity from population-level (i.e., cell count or proliferation marker) data. While we find that the relevant model parameter is *in theory* identifiable, this identifiability is lost for imprecise cell-counts. Indeed, the difficulty in distinguishing between the functional form of the adaptation velocity (Fig. 7), combined with the narrow time window in which the proliferation and death rates are distinguishable (Fig. 4) suggests that heterogeneity may be indistinguishable from misspecification of other model terms. It is only if cell-level information (i.e., timings of proliferation and death events in the population) are available that we are, in theory, able to establish heterogeneity; although the timescale of adaptation compared to the cell doubling time (less than 7 d compared to ∼2–7 d for melanoma [6]) may arise as a practical limitation that prevents heterogeneity from becoming practically identifiable.

A consequence of the non-identifiability of heterogeneity is that we cannot distinguish be-

tween a heterogeneous and a homogeneous model. Mathematically, this offers a practical benefit as it implies that population-level behaviours are well characterised by a simple, and in many cases analytically tractable, ODE model. All other model parameters: the on- and off-drug growth rates in each fully adapted state, and the adaptation velocity, are practically identifiable. In the supplementary material, we explore a number of experimental designs in which various combinations of termination times are considered for a fixed total number of proliferation assays. Even if all experiments are terminated after 3 d, all relevant parameters remain identifiable; albeit estimates are drawn with reduced precision. The indistinguishability of the heterogeneous and homogeneous models motivates us to explore the model selection question using an ODE-based homogeneous model. We are unable to distinguish the functionally correct adaptation velocity, however we do identify a class of models that manifest similar trajectories through phenotype space (Fig. 7c).

The theoretical identifiability of heterogeneity from event-timing data using our stochastic formulation highlights two potential (and rarely considered) sources of model misspecification. First, that proliferation and death events are mutually exclusive: often it is only in a stochastic modelling framework that the two can be distinguished [44]. A more realistic (and correspondingly, further parameterised) model would consider individual and phenotype-dependent proliferation and death rates. Depending on the action of the drug and the metabolic cost of resistance, it may be appropriate for one of these rates to be phenotype-independent. For large-cell-count experiments, it may also be appropriate to consider a stochastic analogue of logistic, rather than exponential, growth. A second source of misspecification, the effects of which are, to the best of our knowledge, largely unknown in the context of population-level behaviour in IBMs, is that proliferation occurs according to a Markov process. Clearly, this is a strong assumption that, while routine in the mathematical literature, may be inappropriate. Given that adaptation occurs on a similar timescale as proliferation, what an appropriate time-to-proliferation distribution would be is entirely unclear. Furthermore, any move away from a Markovian formulation in the IBM would render intractable the ODE, PDE, and CME formulations that we rely upon for inference.

We follow the vast majority of the PDE literature in assuming that heterogeneity arises solely according to a (biased) random walk in phenotype space [27]. A consequence of this choice is that the phenotypic state of a cell is constantly evolving, even within a system that appears static at the population-level. It is only this formulation of heterogeneity that we find to be non-identifiable. A potentially more realistic model of heterogeneity is one that considers inherent heterogeneity between individual cells; for example, where individual cells have a unique growth rate when both fully sensitive and fully adapted or where the adaptation speed varies. The question of identifiability of these population-level distributions from population-level statistics, such as cell count, remains open, although there is a fast-growing set of statistical tools that could be adapted to answer these questions [42, 45].

Our final result is to demonstrate that the discrete-binary model of heterogeneous phenotypic adaptation is indistinguishable at the population level from an appropriately formulated continuous model. Our results do show very minor differences in high-order behaviours (cell

count variance), although we still expect both models to remain indistinguishable upon consideration of potential model misspecification and measurement noise. We do expect the binary strategy to be identifiable from marker data, as, mid-adaptation, a binary strategy may manifest a bimodal distribution that contrasts with the continuous transition we see both experimentally (Fig. 1b), and in our model (Fig. 6a). The equivalence we derive in Section 3.5 also suggests at a hybrid discrete-continuous model that could be studied in future. Namely, a model in which drug-sensitive cells switch to an intermediate transition state with some drug-dependent propensity, in which the phenotypic state varies continuously until the cell reaches the drug-resistant state.

We establish the identifiability of potentially heterogeneous phenotype adaptation from commonly reported low-cell-count proliferation assay experiments. To achieve this, we develop a computationally efficient inference framework for heterogeneous phenotype adaptation that captures potential information arising as intrinsic noise, without resorting to the study of a mean-field model subject to an additive Gaussian measurement process. That we find heterogeneity non-identifiable is significant to the mathematical modelling community, and implies that population-level behaviours (including, importantly, the response of systems to drugs and the design of adaptive therapies) are well characterised by homogeneous ODE models. For the experimental community, our methodology can be used to design and characterise experiments that probe continuous adaptive resistance.

Data availability

Code used to produce the results is available at `https://github.com/ap-browning/phenotypic_heterogeneity_ibm`.

Acknowledgements

The authors thank Adriana Zanca for helpful discussions. APB thanks the Mathematical Institute, University of Oxford, for a Hooke Research Fellowship. SH was funded by Wenner-Gren Stiftelserna/the Wenner-Gren Foundations (WGF2022-0044) and the Kjell och Märta Beijer Foundation. ALJ thanks the London Mathematical Society. CV is a Fellow of the Paris Region Fellowship Programme. This work was partially supported by a Heilbronn Institute for Mathematical Research Small Maths Grant to TC.

References

[1] McGranahan N, Swanton C. 2017 Clonal heterogeneity and tumor evolution: past, present, and the future. *Cell* **168**, 613–628.

[2] Arozarena I, Wellbrock C. 2019 Phenotype plasticity as enabler of melanoma progression and therapy resistance. *Nature Reviews Cancer* **19**, 377–391.

[3] Boumahdi S, de Sauvage FJ. 2020 The great escape: tumour cell plasticity in resistance to targeted therapy. *Nature Reviews Drug Discovery* **19**, 39–56.

[4] Labrie M, Brugge JS, Mills GB, Zervantonakis IK. 2022 Therapy resistance: opportunities created by adaptive responses to targeted therapies in cancer. *Nature Reviews Cancer* **22**, 323–339.

[5] Shaffer SM, Dunagin MC, Torborg SR, Torre EA, Emert B, Krepler C, Beqiri M, Sproesser K, Brafford PA, Xiao M, Eggan E, Anastopoulos IN, Vargas-Garcia CA, Singh A, Nathanson KL, Herlyn M, Raj A. 2017 Rare cell variability and drug-induced reprogramming as a mode of cancer drug resistance. *Nature* **546**, 431–435.

[6] Kavran AJ, Stuart SA, Hayashi KR, Basken JM, Brandhuber BJ, Ahn NG. 2022 Intermittent treatment of BRAFV600E melanoma cells delays resistance by adaptive resensitization to drug rechallenge. *Proceedings of the National Academy of Sciences* **119**, e2113535119.

[7] Sharma SV, Lee DY, Li B, Quinlan MP, Takahashi F, Maheswaran S, McDermott U, Azizian N, Zou L, Fischbach MA, Wong KK, Brandstetter K, Wittner B, Ramaswamy S, Classon M, Settleman J. 2010 A chromatin-mediated reversible drug-tolerant state in cancer cell subpopulations. *Cell* **141**, 69–80.

[8] Altrock PM, Liu LL, Michor F. 2015 The mathematics of cancer: integrating quantitative models. *Nature Reviews Cancer* **15**, 730–745.

[9] Rockne RC, Hawkins-Daarud A, Swanson KR, Sluka JP, Glazier JA, Macklin P, II DAH, Jarrett AM, Lima EABF, Oden JT, Biros G, Yankeelov TE, Curtius K, Bakir IA, Wodarz D, Komarova N, Aparicio L, Bordyuh M, Rabadan R, Finley SD, Enderling H, Caudell J, Moros EG, Anderson ARA, Gatenby RA, Kaznatcheev A, Jeavons P, Krishnan N, Pelesko J, Wadhwa RR, Yoon N, Nichol D, Marusyk A, Hinczewski M, Scott JG. 2019 The 2019 mathematical oncology roadmap. *Physical Biology* **16**, 041005.

[10] Anderson AR, Weaver AM, Cummings PT, Quaranta V. 2006 Tumor morphology and phenotypic evolution driven by selective pressure from the microenvironment. *Cell* **127**, 905–915.

[11] Clairambault J, Pouchol C. 2019 A survey of adaptive cell population dynamics models of emergence of drug resistance in cancer, and open questions about evolution and cancer. *BIOMATH* **8**, 1905147.

[12] Marusyk A, Janiszewska M, Polyak K. 2020 Intratumor heterogeneity: the Rosetta Stone of therapy resistance. *Cancer Cell* **37**, 471–484.

[13] Gunnarsson EB, De S, Leder K, Foo J. 2020 Understanding the role of phenotypic switching in cancer drug resistance. *Journal of Theoretical Biology* **490**, 110162.

[14] Stace RE, Stiehl T, Chaplain MA, Marciniak-Czochra A, Lorenzi T. 2020 Discrete and continuum phenotype-structured models for the evolution of cancer cell populations under chemotherapy. *Mathematical Modelling of Natural Phenomena* **15**, 14.

[15] Cassidy T, Nichol D, Robertson-Tessi M, Craig M, Anderson ARA. 2021 The role of memory in non-genetic inheritance and its impact on cancer treatment resistance. *PLoS Computational Biology* **17**, e1009348.

[16] Martínez VA, Salvioli M, Laleh NG, Thuijsman F, Brown JS, Cavill R, Kather JN, Staňková K. 2022 Improving mathematical models of cancer through game-theoretic modelling: a study in non-small cell lung cancer. *bioRxiv* p. 2021.10.29.466444.

[17] West J, Ma Y, Newton PK. 2018 Capitalizing on competition: An evolutionary model of competitive release in metastatic castration resistant prostate cancer treatment. *Journal of Theoretical Biology* **455**, 249–260.

[18] Craig M, Kaveh K, Woosley A, Brown AS, Goldman D, Eton E, Mehta RM, Dhawan A, Arai K, Rahman MM, Chen S, Nowak MA, Goldman A. 2019 Cooperative adaptation to therapy (CAT) confers resistance in heterogeneous non-small cell lung cancer. *PLoS Computational Biology* **15**, e1007278.

[19] Kaznatcheev A, Peacock J, Basanta D, Marusyk A, Scott JG. 2019 Fibroblasts and Alectinib switch the evolutionary games played by non-small cell lung cancer. *Nature Ecology & Evolution* **3**, 450–456.

[20] Kareva I. 2022 Different costs of therapeutic resistance in cancer: Short- and long-term impact of population heterogeneity. *Mathematical Biosciences* **352**, 108891.

[21] Lorz A, Lorenzi T, Hochberg ME, Clairambault J, Perthame B. 2013 Populational adaptive evolution, chemotherapeutic resistance and multiple anti-cancer therapies. *ESAIM: Mathematical Modelling and Numerical Analysis* **47**, 377–399.

[22] França GS, Baron M, King BR, Bossowski JP, Bjornberg A, Pour M, Rao A, Patel AS, Misirlioglu S, Barkley D, Tang KH, Dolgalev I, Liberman DA, Avital G, Kuperwaser F, Chiodin M, Levine DA, Papagiannakopoulos T, Marusyk A, Lionnet T, Yanai I. 2024 Cellular adaptation to cancer therapy along a resistance continuum. *Nature* **631**, 876–883.

[23] Smalley I, Kim E, Li J, Spence P, Wyatt CJ, Eroglu Z, Sondak VK, Messina JL, Babacan NA, Maria-Engler SS, Armas LD, Williams SL, Gatenby RA, Chen YA, Anderson AR, Smalley KS. 2019 Leveraging transcriptional dynamics to improve BRAF inhibitor responses in melanoma. *EBioMedicine* **48**, 178–190.

[24] Cassidy T, Craig M. 2019 Determinants of combination GM-CSF immunotherapy and oncolytic virotherapy success identified through in silico treatment personalization. *PLoS Computational Biology* **15**, e1007495.

[25] Strobl MA, West J, Viossat Y, Damaghi M, Robertson-Tessi M, Brown JS, Gatenby RA, Maini PK, Anderson AR. 2021 Turnover modulates the need for a cost of resistance in adaptive therapy. *Cancer Research* **81**, 1135–1147.

[26] Dieckmann U, Law R. 1996 The dynamical theory of coevolution: a derivation from stochastic ecological processes. *Journal of Mathematical Biology* **34**, 579–612.

[27] Chisholm RH, Lorenzi T, Lorz A, Larsen AK, de Almeida LN, Escargueil A, Clairambault J. 2015 Emergence of drug tolerance in cancer cell populations: an evolutionary outcome of selection, nongenetic instability, and stress-induced adaptation. *Cancer Research* **75**, 930–939.

[28] Celora GL, Byrne HM, Zois CE, Kevrekidis P. 2021 Phenotypic variation modulates the growth dynamics and response to radiotherapy of solid tumours under normoxia and hypoxia. *Journal of Theoretical Biology* **527**, 110792.

[29] Lorenzi T, Chisholm RH, Desvillettes L, Hughes BD. 2015 Dissecting the dynamics of epigenetic changes in phenotype-structured populations exposed to fluctuating environments. *Journal of Theoretical Biology* **386**, 166–176.

[30] Raue A, Kreutz C, Maiwald T, Bachmann J, Schilling M, Klingmüller U, Timmer J. 2009 Structural and practical identifiability analysis of partially observed dynamical models by exploiting the profile likelihood. *Bioinformatics* **25**, 1923–1929.

[31] Elsasser WM. 1984 Outline of a theory of cellular heterogeneity. *Proceedings of the National Academy of Sciences* **81**, 5126–5129.

[32] Hamis S, Browning AP, Jenner AL, Villa C, Maini P, Cassidy T. 2024 Growth rate-driven modelling reveals how phenotypic adaptation drives drug resistance in BRAFV600E-mutant melanoma. *bioRxiv*. `https://doi.org/10.1101/2024.08.14.607616`.

[33] Kiefel H, Bondong S, Pfeifer M, Schirmer U, Erbe-Hoffmann N, Schäfer H, Sebens S, Altevogt P. 2012 EMT-associated up-regulation of L1CAM provides insights into L1CAM-mediated integrin signalling and NF-κB activation. *Carcinogenesis* **33**, 1919–1929.

[34] Markham DC, Simpson MJ, Baker RE. 2014 Choosing an appropriate modelling framework for analysing multispecies co-culture cell biology experiments. *bioRxiv*. `https://doi.org/10.1101/008318`.

[35] Markham DC, Simpson MJ, Baker RE. 2015 Choosing an appropriate modelling framework for analysing multispecies co-culture cell biology experiments. *Bulletin of Mathematical Biology* **77**, 713–734.

[36] Villa C, Maini PK, Browning AP, Jenner AL, Hamis S, Cassidy T. 2024 Reducing phenotype-structured PDE models of cancer evolution to systems of ODEs: a generalised moment dynamics approach. *arXiv*.

[37] Warne DJ, Baker RE, Simpson MJ. 2019 Simulation and inference algorithms for stochastic biochemical reaction networks: from basic concepts to state-of-the-art. *Journal of The Royal Society Interface* **16**, 20180943.

[38] Risken H, Frank T. 1996 *The Fokker-Planck Equation Methods of Solution and Applications* . Berlin: Springer 2nd edition.

[39] Pawitan Y. 2013 *In all likelihood: statistical modelling and inference using likelihood*. Oxford: Oxford University Press.

[40] Vihola M. 2020 Ergonomic and reliable Bayesian inference with adaptive Markov chain Monte Carlo. *Wiley StatsRef: Statistics Reference Online* pp. 1–12.

[41] Johnson SG. 2021 The NLopt module for Julia. `https://github.com/JuliaOpt/NLopt.jl`.

[42] Browning AP, Drovandi C, Turner IW, Jenner AL, Simpson MJ. 2022 Efficient inference and identifiability analysis for differential equation models with random parameters. *PLOS Computational Biology* **18**, e1010734.

[43] Hanahan D. 2022 Hallmarks of cancer: new dimensions. *Cancer Discovery* **12**, 31–46.

[44] Huynh L, Scott JG, Thomas PJ. 2023 Inferring density-dependent population dynamics mechanisms through rate disambiguation for logistic birth-death processes. *Journal of Mathematical Biology* **86**, 50.

[45] Hasenauer J, Waldherr S, Doszczak M, Radde N, Scheurich P, Allgwer F. 2011 Identification of models of heterogeneous cell populations from population snapshot data. *BMC Bioinformatics* **12**, 125.

Supplementary Material

S1 IBM/CME comparison

In Fig. S1 we compare the probability mass functions arising from the solution of the CME to the empirical distribution arising from 1000 cell proliferation assays simulated using the IBM.

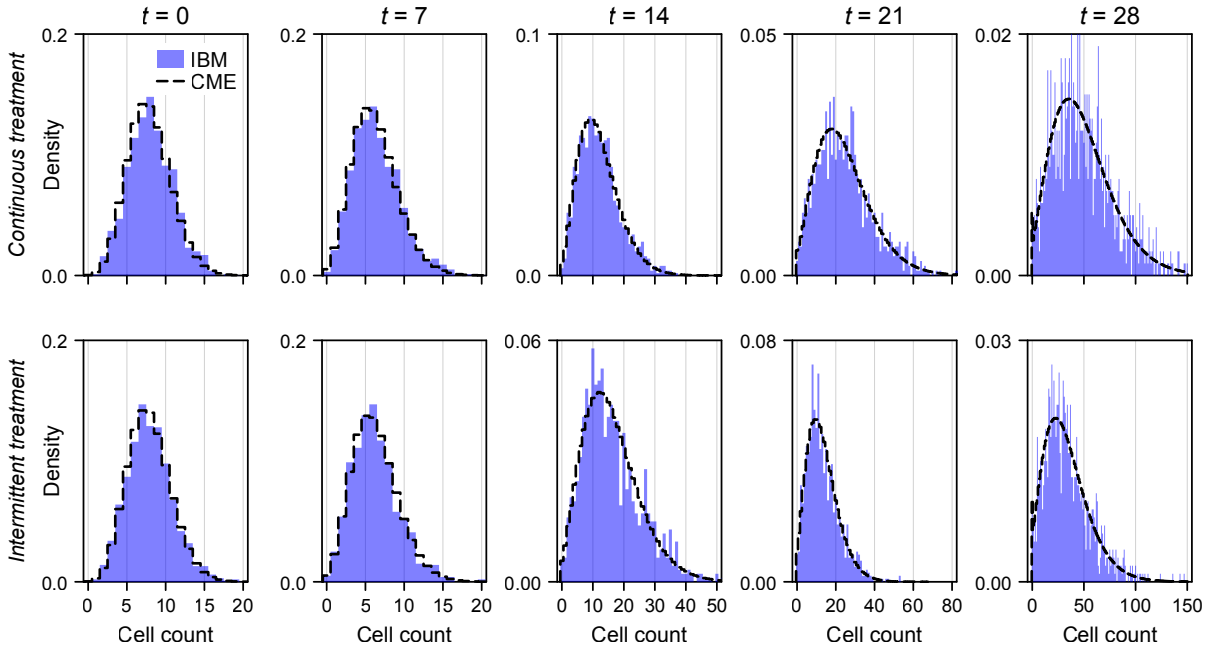

Figure S1. Chemical master equation comparison. Comparison between $n = 1000$ realisations of the IBM (blue) and the solution to the CME (black dashed) under continuous treatment (top row) and intermittent treatment (bottom row). All parameters are consistent with those in Fig. 2 of the main text.

S2 MCMC priors and results

The priors for each parameter are given in Table S1. In Fig. S2, we show marginal posterior distributions for all parameters corresponding to the analysis in Fig. 3 of the main text. In the main text, we present results for experiment termination times $t = \{1\,\mathrm{d}, 3\,\mathrm{d}, 5\,\mathrm{d}, 7\,\mathrm{d}\}$. In Fig. S2, we additionally consider various other termination time sets, as indicated.

Table S1. Parameters, parameter descriptions, and parameter priors used for analysis in the main text.

Parameter	Description	Prior
γ_1	Sensitive ($x = 0$) growth rate off drug	Uniform$(-1, 1)$
γ_2	Sensitive ($x = 0$) growth rate on drug	Uniform$(-1, 1)$
γ_3	Resistant ($x = 1$) growth rate off drug	Uniform$(-1, 1)$
γ_4	Resistant ($x = 1$) growth rate on drug	Uniform$(-1, 1)$
$\log \nu$	Adaptation speed	Uniform$(-6, 1)$
$\log \beta$	Heterogeneity/diffusivity parameter	Uniform$(-6, -1)$

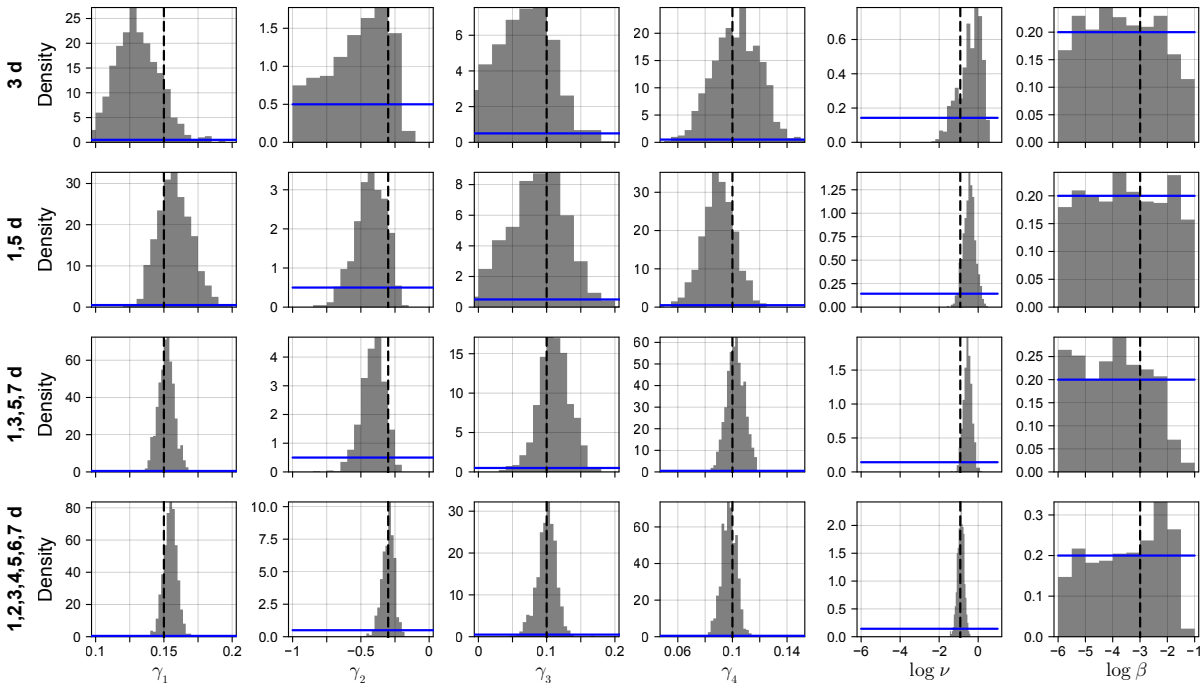

Figure S2. Marginal posterior distributions. We show marginal posterior distributions for all parameters, for various sets of observation termination times. Shown also are the marginal prior distributions (blue) and the true values (black dashed).

S3 CME for discrete model

Let $q(n_1, n_2, t)$ be the density for n_1 sensitive cells and n_2 resistant cells. Let λ_i denote the proliferation rate of each cell population, and δ_i denote the death rate of each cell population (one of which will be zero). We have that

$$\begin{aligned}\frac{\mathrm{d}q(n_1, n_2, t)}{\mathrm{d}t} = &(n_1 - 1)\lambda_1 q(n_1 - 1, n_2, t) + (n_1 + 1)\delta_1 q(n_1 + 1, n_2, t) \\ &+ (n_2 - 1)\lambda_2 q(n_1, n_2 - 1, t) + (n_2 + 1)\delta_2 q(n_1, n_2 + 1, t) \\ &+ n_1 r_{12} q(n_1 + 1, n_2 - 1, t) + n_2 r_{21} q(n_1 - 1, n_2 + 1, t) \\ &- (n_1\lambda_1 + n_1\delta_1 + n_2\lambda_2 + n_2\delta_2) q(n_1, n_2, t).\end{aligned} \tag{20}$$

To obtain the mass function for the total cell count, $q(n, t)$, we consider that

$$q(n, t) = \sum_{n_1=0}^{\infty} q(n_1, n - n_1, t). \tag{21}$$

In practice, we consider a partial sum truncated at $n_1 = 50$, which we find to be sufficient given the maximum cell counts observed in Fig. 8 of the main text.

S4 Structural identifiability of heterogeneity

In the main text, we observe that the heterogeneity parameter, β, is one-sided practically identifiable. In this section, we reproduce the results in Fig. 3 of the main text (and equivalently, Fig. S2 of this supplementary material document) to investigate a scenario where a large data set is generated that comprises 768 cell proliferation assays (i.e., eight plates), for each condition, at a set of termination times $t = \{0.5\,\mathrm{d}, 1\,\mathrm{d}, 1.5\,\mathrm{d}, \cdots, 6.5\,\mathrm{d}, 7\,\mathrm{d}\}$. Results in Fig. S3 show that, in this large data-set regime (5,376 proliferation assays equivalent to a total of 56 plates), the diffusivity parameter is practically identifiable, demonstrating that it is a structurally identifiable parameter.

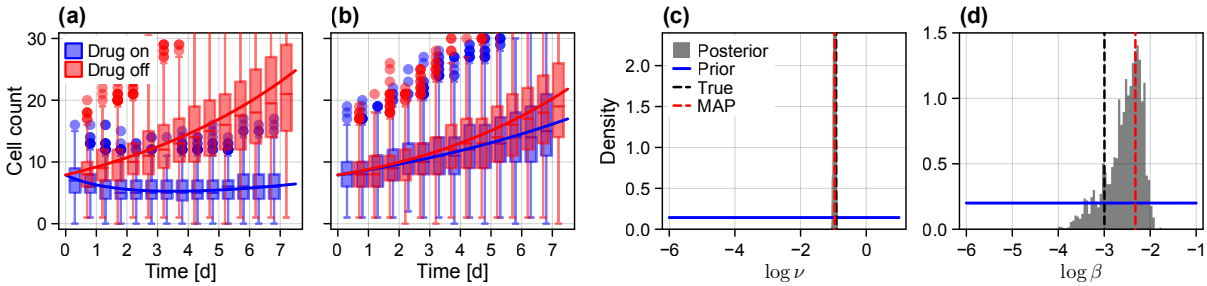

Figure S3. Large data set proliferation assay inference. We show marginal posterior distributions for all parameters, for various sets of observation termination times. We perform Bayesian inference on synthetic cell proliferation assay data using the CME as a likelihood. Independent cell count observations ($M = 768$ replicates per condition) are made from experiments conducted with fully sensitive or fully resistant cells, with and without drug, and terminated at $t = \{0.5\,\mathrm{d}, 1\,\mathrm{d}, 1.5\,\mathrm{d}, \cdots, 6.5\,\mathrm{d}, 7\,\mathrm{d}\}$. (a–b) Synthetic proliferation assay cell count data (box plots with outliers shown as discs) and the model predicted mean cell count at the MAP (solid lines). (c–d) Posterior distributions for the logarithm of v, the adaptation speed, and β, the diffusivity. Shown also is the uniform prior (blue), the true value (black dashed), and the MAP (red dashed). In this regime, both the adaptation speed and diffusivity are identifiable.

S5 Large data set inference with noisy data

In the main text, and in all supplementary results hitherto, we have studied a scenario where exact cell counts are made. We find in Section S4 that, in this effectively noise-free regime, the heterogenity parameter is identifiable provided that a sufficiently large data set is available. We now revisit this assumption by assuming that cell count observations, denoted y, are subject to binomial noise such that

$$y \,|\, n \sim \text{Truncated}\left(\text{Binomial}(M(n), 0.5) - \frac{M(n)}{2} + n, 0, \infty\right), \tag{22}$$

where n is the exact cell count, and $M(n)$ is chosen such that the standard deviation of the noise term scales with the cell count (note that for the right hand side of Eq. (22) to be valid, we require that $M(n)$ is always even, such that the noise term $\text{Binomial}(M(n), 0.5) - M(n)/2$ is symmetric about zero. In this section, we set

$$M(n) = 2\text{round}\left(\frac{4\alpha^2 n^2 + n_0}{2}\right), \tag{23}$$

such that the noise comprises a count independent term, n_0 (i.e., noise present even in very low cell count observations arising from, e.g., cellular debris), and a count dependent term of magnitude α which scales such that the standard deviation of the noise term is approximately αn for large n. For the results that follow, we set $\alpha = 0.1$ and $n_0 = 5$; for these parameter values we demonstrate the noise distribution and resultant observed cell count distribution in Fig. S4.

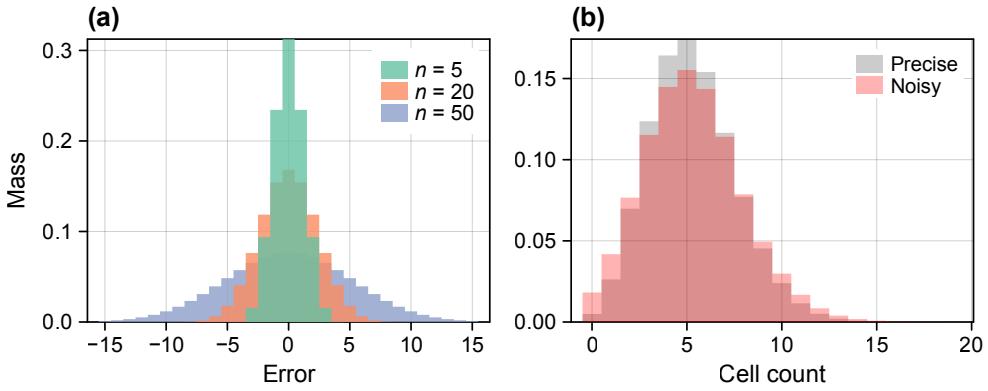

Figure S4. Observation noise model. (a) We consider cell count observations subject to additive Binomial error that scales with the cell count, n, according to Eq. (23). (b) Comparison between precise (black) and noisy (red) cell count distributions from the CME.

Following the construction of the statistical observation noise model, we reproduce the large data set results of Fig. S3 in the case that only noisy observations are available. Priors are given in Table S2 and both fits and marginal posterior distributions in Fig. S4. Results in Fig. S5 demonstrate that the diffusivity parameter is no longer identifiable; even from a large data set, we cannot distinguish heterogeneity from observation noise.

Table S2. Prior distributions for the noise distribution parameters used in the inference of noisy data.

Parameter	**Prior**
α	Uniform$(0, 1)$
n_0	Uniform$(0, 10)$

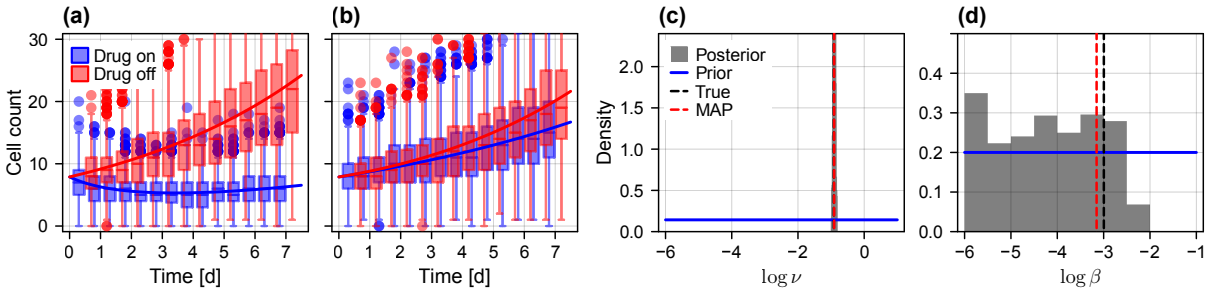

Figure S5. Large data set proliferation assay inference with noisy data. We reproduce the results in Fig. S3 in the case that cell count observations are subject to noise of the form given in Eq. (22). The noise parameters α and n_0 are assumed to be unknown, with priors given in Table S2.