

A lineage tree-based hidden Markov model to quantify cellular heterogeneity and plasticity

Farnaz Mohammadi¹, Shakthi Visagan¹, Sean M. Gross², Luka Karginov³, JC Lagarde¹,
Laura M. Heiser², and Aaron S. Meyer^{1,4,5,6}

¹Department of Bioengineering, University of California, Los Angeles, USA.

²Department of Biomedical Engineering, Oregon Health and Science University, Portland, USA.

³Department of Bioengineering, University of Illinois, Urbana Champaign, USA.

⁴Department of Bioinformatics, University of California, Los Angeles, USA.

⁵Jonsson Comprehensive Cancer Center, University of California, Los Angeles, USA.

⁶Eli and Edythe Broad Center of Regenerative Medicine and Stem Cell Research, University of California, Los Angeles, USA.

Abstract

Cell plasticity operates alongside other sources of cell-to-cell heterogeneity, such as genetic mutations and variation in signaling, together preventing most cancer therapies from being curative. The predominant methods of quantifying tumor-drug response operate on snapshot, population-level measurements and therefore lack evolutionary dynamics, which are particularly critical for dynamic processes such as plasticity. Here we apply a lineage tree-based adaptation of a hidden Markov model that employs single cell lineages as input to learn the characteristic patterns of single cell phenotypic heterogeneity and state transitions in an unsupervised fashion. To benchmark our model, we paired cell fate with either cell lifetimes or individual cell cycle phase lengths on synthetic data and demonstrated that the model successfully classifies cells within experimentally tractable dataset sizes. As an application, we analyzed experimental measurements of same measurements in cancer and non-cancer cell populations under various treatments. We find that in each case multiple phenotypically distinct states exist, with significant heterogeneity and unique drug responses. In total, this framework allows for the flexible classification of single cell heterogeneity across lineages.

Introduction

Chemotherapy and targeted therapies selectively eliminate fast-proliferating or oncogene-addicted cells and are among the primary treatments for cancer. However, long-term therapeutic efficacy is significantly limited by widespread intratumoral heterogeneity^{1,2}. Cell-to-cell variability in drug response can originate from cell-intrinsic factors—such as genomic alterations, epigenetic mechanisms like changes in chromatin state³, and variable protein levels^{4,5}—or cell-extrinsic factors such as spatial variability in the surrounding vasculature and environmental stressors^{6–8}. Moreover, cell plasticity, where cells adopt new characteristics such as those of other cell types, is observed in cancer cells, and can affect their sensitivity to therapy⁹.

37 Large-scale profiling studies can find molecular features that associate with drug response using
38 population-level samples^{10,11}. These associations, while valuable, can miss the contribution of
39 cell-to-cell heterogeneity, and especially stochastic changes in individual cell states that have
40 significant effects on overall tumor drug response^{3,12,13}. The most common methods for
41 quantifying drug response are metrics of tumor cell population expansion or contraction^{14–17}.
42 Recent research has made efforts to track phenotypic measurements of fitness at the single cell
43 level^{18,19}, however, even single cell measurements are typically performed with snapshots that
44 subsequently miss the role of individual cells in the overall population response²⁰. Though
45 population heterogeneity is usually defined through molecular measurements, studies that have
46 explicitly linked molecular and phenotypic variation have been able to identify mechanisms that
47 underly cell-to-cell variation that would otherwise remain hidden²¹, and studies starting with
48 phenotypic analysis have generally found that phenotypic variability arises from a small number
49 of molecular factors leading to the phenotypic variation^{4,22,23}.

50 Measurements accompanied by lineage relationships are uniquely valuable for studying inherited
51 phenotypes within families of people. This value is evident in linkage studies wherein relatives
52 are used to identify or refine the genetic determinants of disease^{24–26}. Notably, linkage studies
53 can identify genetic determinants with greater power than even large association studies because
54 relatives essentially serve as internal controls²⁷. Linkage studies also start with the phenotype of
55 individuals, rather than grouping based on molecular differences, ensuring discoveries are
56 phenotypically consequential. While the inherited factors are different between cells (e.g.,
57 proteins, RNA) and people (DNA), such approaches are likely to be similarly useful with
58 populations of cells. Recently, constructing phylogenetic trees of cancer cells using lineage tracing
59 and single cell sequencing has helped to characterize the directionality of metastatic seeding,
60 though these methods are limited to tracking slow processes such as mutational differences²⁸.
61 Lineage-resolved data has also demonstrated value in uncovering cell-to-cell heterogeneity due
62 to transient differences outside of cancer^{22,23}. Therefore, tools to analyze and explore these data
63 will be critical.

64 Hidden Markov models (HMMs) provide an efficient strategy to infer discrete states from
65 measurements when a series of co-dependent observations are made. An example of this is their
66 wide-spread use in time series analysis, where each measurement is dependent on those that
67 came before^{29,30}. Recognizing this co-dependence allows HMMs to make accurate inferences
68 even in the presence of extremely noisy measurements since each neighboring measurement
69 provides accumulating evidence³¹. These models derive their relative simplicity by assuming a
70 Markov process, meaning that the current behavior of a system can be assumed to be
71 independent of its earlier history should its current state be known. This assumption naturally
72 applies in many contexts. In the case of cells, this assumption aptly captures cell inheritance
73 because daughter cells inherit both molecular signals and their environment from their
74 predecessor. Indeed, several recent examples of cell-to-cell inheritance mechanisms can be
75 represented as a Markov process through linear chains or cycles of states^{12,22,23}. HMMs have
76 been adapted to lineage trees (tHMMs) so that each measurement across the tree can similarly
77 provide accumulating evidence for a prediction. Just like with time-series data, these models can
78 provide very accurate predictions despite noisy measurements and limited information by
79 recognizing the co-dependence between related measurements^{32,33}. tHMMs have been used in a
80 multitude of applications, from image classification to comparative genomics^{34,35}. In cells, these
81 models have been fit to lineages collected from stem cells and bacteria colonies, but have always

82 required tailor-made implementations^{36,37}. Improvements in cell tracking and high-throughput
83 imaging promise to make these models valuable techniques for studying the plasticity of
84 heterogeneous cell populations. However, widespread use of these models still depends on more
85 easily usable implementations, examples of successful tHMM-based discoveries, and standards
86 for experimental application.

87 Here, we develop an extensible implementation of tHMMs with a defined interface for
88 integrating diverse types of measurements on cell lineage trees. This model allows us to quantify
89 the dynamics and phenotypic features of drug response heterogeneity. We used the relationship
90 between the cells to analyze how populations of breast cancer cells respond to therapy with a cell
91 cycle reporter, and how normal breast cells respond to growth factor treatment. Single cell
92 measurements of the cell cycle revealed extensive variation not captured by population-level
93 measurement. Using the tHMM model, we inferred the number of phenotypically distinct
94 subpopulations, the characteristics of those subpopulations, the transition probabilities from one
95 state to another, and each cell's expected state. We also confirmed that the tHMM model could
96 use patterns of inheritance to predict cell behavior. This work, therefore, provides a flexible
97 phenotype-driven route to discovering cell-to-cell variation in drug response, demonstrates an
98 overall strategy for quantifying the dynamics of cell heterogeneity, and implements a very
99 general software tool for the widespread use of tHMM models.

100 **Results**

101 **Lineage information provides unique information about the source and** 102 **dynamics of intratumoral heterogeneity**

103 Single cells grow and then divide into two daughter cells, eventually forming a binary
104 genealogical tree, also known as a lineage tree. We collected single-cell measurements in the
105 form of lineage trees to track these relationships. The life cycle of each cell before division
106 includes phases that must pass one after another. To illustrate the unique value of lineage
107 measurements in analyzing intra-tumoral and drug response heterogeneity, we collected cell fates
108 alongside either cell lifetimes (MCF10A) or individual cycle phase durations (AU565). Two
109 random subsets of the tracked lineages of the breast cancer cell line AU565 are plotted in Fig. 1a.
110 The single cell lineages reveal striking variation in cell cycle phase durations and cell division
111 dynamics despite coming from the same sample. Population-level measurements would be
112 unable to identify this difference as the starting and ending cell numbers are the same.
113 Measurements that record or reflect the history of cells (e.g., CFSE staining, Luria-Delbruck
114 experiment) can help to identify these variations within cell populations but must make
115 assumptions about the dynamics of heterogeneity^{13,23}. Lineage measurements, by contrast,
116 provide sufficiently rich temporal information to quantify the specific structure of the phenotypic
117 heterogeneity.

118 As a further exploration of the cell tracking data, we randomly sampled lineages from
119 gemcitabine treated AU565 cells (Fig. 1b). Gemcitabine is a chemotherapy agent that disrupts
120 DNA replication and results in extension of and apoptosis in S phase³⁸. We found that the S/G2
121 phase lengths in treated cells were noticeably extended compared to untreated cells, slowing
122 population growth. There was generally striking variation between lineages of a single condition,
123 including anywhere from zero to three cell divisions, but tightly shared behavior among cells and

124 their relatives in each lineage. These observations demonstrate some of the unique advantages of
125 collecting lineage-based measurements.

126 **Figure 1: Total cell number is insufficient to distinguish the structure of heterogeneous**
127 **populations. (a)** Randomly sampled lineages of untreated AU565 cells from the same replicate
128 and experiment. **(b)** Randomly sampled lineages of AU565 cells treated with 5 nM gemcitabine
129 from a single replicate and experiment. Each line indicates the lifetime of one cell. A line
130 branching into two lines indicates cell division. The G1 and S/G2 phase durations are indicated
131 by solid thick and thin lines, respectively.

132 **A lineage tree-based hidden Markov model infers the state of cells given** 133 **measurements on lineage trees**

134 Given the unique insights that single cell measurements on lineage trees can provide, we
135 implemented a strategy for classifying cells based on their phenotype and lineage relationships.
136 We used a tree-based hidden Markov model (tHMM) to fit a set of measurements made across a
137 lineage tree (Fig. 2a). Like a typical hidden Markov model, a tHMM can infer the hidden
138 discrete “states” of cells given a series of measurements where a state is defined by specific
139 phenotype distributions. The inference of these states takes place using an iterative strategy
140 wherein the states of each cell are predicted by the phenotype of both the cell and its relatives in
141 a lineage (“expectation” step), and then each distribution of phenotypes is fit to match the cells
142 within that state (“maximization” step) (Fig. 2b). This expectation-maximization (EM) process
143 repeats until convergence.

144 After fitting, the model can provide a variety of information (Fig. 2c). First, it infers the starting
145 and transition probabilities of each state. Second, the distribution of cells’ phenotypes in each
146 state are estimated and can be compared to distinguish how cells of each state behave. For
147 instance, if we use the growth rates of cells as their phenotype, we may observe a subpopulation
148 of cells with shorter times to division, and another with longer times. Moreover, the state of each
149 individual cell can be predicted from the fit data or new measurements. Finally, the model
150 provides a likelihood of each cell’s observations and therefore the data overall. This last quantity
151 can be used, for example, to estimate the number of distinguishable cell states. When
152 implementing these processes, we ensured that a cell’s measurements are defined through a
153 modular interface, allowing many other forms of data to be easily integrated, such as cell
154 morphology or molecular measurements.

155 **Figure 2: The tHMM interface. (a)** Input data takes the form of single cell measurements
156 across time, where the lineage relationship between cells is known. **(b)** The fitting process
157 includes expectation and maximization steps, where model parameters are iteratively updated
158 until convergence. **(c)** Output predictions of the model after fitting including the tree of inferred
159 cell states, probabilities of transition between each state, starting abundance of each cell state,
160 and distributions that describe the behavior of cells within each state. The model likelihood can
161 be used to estimate the number of distinguishable cell states.

162 **Experiments of finite time necessitate corrections for experimental censorship**

163 Modeling the duration of each cell’s lifetime is complicated by the influence of experimental
164 parameters. Specifically, cells at the beginning or end of an experiment persist beyond the

165 experiment's duration and so, while we observe these cells, we do not know their exact lifetime.
166 Data censorship occurs when a measurement is systematically affected by an undesired
167 influence. For instance, in our case, phase durations are censored because the experiment started
168 after cells had already begun their initial cell cycle phase or the experiment ended before they
169 had completed their last phase. Previously, this has been addressed by removing incompletely
170 observed cells²². However, doing so results in a systematic bias, where longer-lived cells are
171 preferentially eliminated. On the other hand, ignoring the truncation of these values also creates
172 bias by creating an upper bound on the cells' lifetimes (Fig. 3b,c).

173 To correct for this effect in our model, we marked cells that encountered the start or end bounds
174 of the experiment. When estimating the properties of these cells' lifetime we instead used a
175 censored estimator or the survival function of the distribution³⁹. Using synthetic data, we verified
176 that this correction resulted in accurate phenotype estimations (Figs. 3d, S3, S10). Thus,
177 accounting for cells that outlive the bounds of the experiment through a censored estimator
178 removes the contribution of this experimental confounder.

179 **Figure 3: Experiments of finite time necessitate data censorship corrections.** (a) An example
180 synthetic, uncensored two-state lineage. (b) An example synthetic, censored two-state lineage.
181 Cells in state 0 and 1 are shown in green and blue. (c) State assignment accuracy with censored
182 lineages using an uncorrected model. (d) State assignment accuracy with censored lineages using
183 the corrected model. Each scatter point represents the state assignment accuracy of the model
184 when fit to a lineage with the indicated number of cells. The solid lines show the Lowess
185 trendline of the individual run accuracies. 100 trials are plotted.

186 **Synthetic lineage benchmarks show a tHMM can accurately infer population** 187 **behavior**

188 To evaluate how accurately a tHMM model could infer the behavior of multi-state cell
189 populations, we used synthetic populations of cells in a wide variety of configurations, such as
190 various populations sizes, numbers of states, and abundance of the states. In each case, we
191 determined that the tHMM model could accurately infer the hidden states and parameters of a
192 population given at least 100 cells. This synthetic data included uncensored (Figs. S1, S2, S8, S9,
193 Table 1-2) or censored (Figs. 4, S10, S3, S15; Table 1–3) situations. Synthetic data were created
194 by lengthening the simulated experiment time, in effect creating “deeper” lineages, or by
195 increasing the number of initial cells to have a greater number of lineages, increasing the
196 experiment's “breadth”. In addition to varying the number of cells in a population, we
197 benchmarked populations with varied cell state percentages (Figs. S4, S5) and varied degrees of
198 phenotypic differences (Figs. S6, S7, 5). This benchmarking consistently showed that the tHMM
199 model would provide accurate results across a range of circumstances, and generally provided
200 accurate results with datasets consisting of at least 10 lineages, 100 cells overall, and 10 cells
201 from each state.

202 More specifically, one of the benchmarking studies we performed was with data matching our
203 measurements of AU565, where G1 and S/G2 phase durations represented by a gamma
204 distribution, and their corresponding cell fate represented by a Bernoulli distribution, were
205 quantified (Fig. 4). The choice of the gamma distribution for cell cycle phase was inspired by a
206 previous study⁴⁰ and verified by evaluating a variety of distributions; the gamma distribution fit
207 the cell lifetime data best. Although the tHMM model was fit with no information about the true

208 underlying parameters of the simulated cells, it distinguished the pre-assigned two underlying
209 cell states' phenotypes (Fig. 4b–d) and member cells with >95% accuracy (Fig. 4e). The
210 Wasserstein distance metric was used to quantify the difference between the true and estimated
211 cell cycle phase duration distributions to show the accuracy of parameter estimation (Fig. 4d).
212 On the population level, the difference between the true and estimated transition probabilities, as
213 calculated by the sum of squared difference, was less than 0.1 for 100 cells or more. Starting
214 probabilities were compared to their corresponding true values using the Euclidean distance and
215 showed less than a 0.2 error for populations with 10 lineages or more (Fig. 4f–g). Thus, we are
216 confident that with similar experimental data, we should derive accurate results.

217 **Figure 4: Model performance on censored lineages of two states with increasing breadth**
218 **and depth.** (a) Synthetic two-state populations of increasing breadth (increasing number of
219 initial cells and therefore lineages) and of increasing depth (increasing experiment time and
220 therefore more cells in each lineage) are analyzed. The states are shown as green and blue colors.
221 Red indicates cell death. (b–c) The accuracy of estimating the Bernoulli parameters for G1 and
222 S/G2 phase, respectively. Each point in the scatter plots represents the inferred value for a model
223 evaluation trial with the number of cells shown in the x-axis. The dark solid lines are the Lowess
224 trendline across the individual trials. The light green and light blue lines show the true value of
225 the parameters. (d) The distance between the true and estimated gamma distributions associated
226 with phase lengths for the two states. (e) The state assignment accuracy. (f) The errors in the
227 estimated and transition rate matrices. (g) The initial probability vector. Note that the
228 Wasserstein distance between the true and estimated distributions for each state is much lower
229 than the distance between two distributions that are quite similar (Fig. 5b). 100 trials are plotted.

230 **Lineage information improves cell state identification with heritable** 231 **phenotypes**

232 Cells of even very distinct molecular states can have partly overlapping phenotypes due to non-
233 heritable variation. Therefore, we sought to evaluate how different two states need to be for us to
234 accurately identify them as distinct (Fig. 5a). We varied the G1 phase duration of two states from
235 identical to very distinct (Fig. 5b) and quantified the state assignment accuracy of our model
236 (Fig. 5c). While the phenotypic observation of a given state had to be different for our model to
237 accurately assign cells, even moderately overlapping phenotypes (Wasserstein distance of ~20)
238 could be distinguished by using the lineage relationships of cells. As a baseline comparison, we
239 analytically identified the optimal classifier in the absence of lineage information (see Methods).
240 The tHMM consistently outperformed this approach (Fig. 5c). The model performance in
241 censored and uncensored populations were similar (Fig. S6, S7). This shows that lineage
242 relationships can be used to identify cell states more accurately with partially overlapping
243 phenotypes.

244 **Figure 5: Model performance versus the difference between states.** (a) Cartoon of how two
245 states can vary in their phenotypic similarity, in a synthetic population of two states. On the top,
246 cells might be virtually indistinguishable (here based on shape). On the bottom, they might be so
247 different that looking at one cell is sufficient to identify its state. (b) The distribution of G1
248 duration is varied in state 1 (blue) while the other state is kept constant. (c) State assignment
249 accuracy versus the Wasserstein distance between state phenotypes. Each point represents the
250 accuracy of state assignment for a lineage created by a set of parameters that yield the shown

251 Wasserstein distance between the two state distributions. 100 trials are plotted. Either the tHMM
252 model (blue) or an optimal classifier without lineage information (orange) was used. The solid
253 lines show a Lowess trendline of the model accuracy.

254 **Likelihood-based model selection can effectively identify the number of** 255 **distinct states**

256 One does not usually know the number of distinct cell states within a population. Further, the
257 number of distinct states may depend on the environmental context of the cells, particularly
258 because we use phenotypic measurements^{41,42}. To test whether we could infer the number of
259 phenotypically distinct states, we performed model selection using the Bayesian information
260 criterion (BIC) while varying the number of states in synthetic data (Fig. 6). We normalized the
261 BIC values such that zero corresponds to the state with the highest likelihood. The synthetic
262 populations included approximately 250 to 650 cells with known cell phase fate and phase
263 lengths (Table 3). The inferred number of cell states was consistently correct, and the few
264 incorrect predictions still centered around the true answer for both uncensored and censored
265 lineages (Fig. 6). This indicated that model selection can help to identify the appropriate number
266 of cell states for a set of measurements.

267 **Figure 6: Model selection effectively identifies the number of distinct states. (a-d)** Model
268 BIC for synthetic uncensored lineages with 1–4 states. **(e-h)** Model BIC for synthetic censored
269 lineages with 1–4 states. BIC values are normalized such that the optimum is equal to 0. The
270 minimum BIC value corresponds to the predicted number of states in each repetition. 5 trials
271 plotted.

272 **tHMM infers multiple distinct subpopulations in experimental drug response** 273 **data**

274 As an application of our model, we used phenotypic measurements from two cell lines. With the
275 first, AU565, we measured of the G1 and S/G2 phase durations and terminal cell fates of cells in
276 a control condition and when treated with 3 concentrations of gemcitabine or lapatinib. For the
277 second, MCF10A, we measured the overall cell lifetimes and terminal fates of cells treated with
278 PBS or single concentrations of the growth factors EGF, HGF, or OSM. Cells were imaged every
279 30 minutes and then tracked over time to assemble lineage relationships. The lapatinib and
280 gemcitabine treated AU565 populations (including control) contained a total of 5290 and 4537
281 cells, respectively. The MCF10A population contained 1306 cells. Lineages included 1–5
282 generations of cells. The model was then fit to each experiment's data across all conditions,
283 enforcing that the initial and transition probabilities are shared across concentrations but
284 allowing the phenotype distributions to vary. We enforced a unidirectional phenotypic shift with
285 drug concentration in AU565 cells, reflecting the expectation of a dose-response effect on cell
286 phenotype within each state. The cell fate parameters were estimated without constraints. We
287 assumed the number of states is shared across drug concentrations in AU565 cells and across
288 growth factor treatments in MCF10A cells. To determine the number of cell states, we compared
289 models of 1–7 states using the BIC, where the lowest BIC value across numbers of states
290 indicates the most optimal model correcting for complexity (Fig. 7a–c). The data for each
291 compound indicated the presence of multiple inherited states.

292 To verify the model's predictive ability, we additionally implemented a cross-validation scheme
293 for the lineage data. Briefly, roughly 20% of the cells were chosen at random and then masked
294 from the fitting process. The model parameters were estimated using only the unmasked cells,
295 though all cells received state assignments through use of their relatives. At the end, the log-
296 likelihood of the masked cells' observations were evaluated using the fit model. We tested this
297 cross-validation approach by creating synthetic cell populations of 2–5 true states with
298 conditions matching the experimental data. For each scenario, we were able to identify the
299 correct number of states based on which gave the highest log-likelihood (Fig. S16a–d, Table 3).
300 Cross-validating the experimental data again confirmed the 4 and 5 phenotypic states within the
301 lapatinib and gemcitabine data, respectively (Fig. S16e,f). It also directly demonstrated that the
302 inclusion of multiple states enables the tHMM model to predict unseen data, and that this
303 prediction is dependent on inheritance; a no-inheritance model in which all transitions were
304 equally likely performed relatively poorly.

305 **Figure 7: BIC-based model selection infers the number of phenotypically distinct states.**
306 Normalized BIC values for (a) AU565 cells in control and treated with 5 nM, 25 nM, and 250
307 nM of lapatinib; (b) AU565 cells in control and treated with 5 nM, 10 nM, and 30 nM of
308 gemcitabine; and (c) MCF10A cells treated with PBS, 10 ng/ml EGF, 40 ng/ml HGF, and 10
309 ng/ml OSM. The BIC values for all conditions were normalized such that the minimum value
310 was zero.

311 **Lapatinib response is defined by both stable and cyclical states**

312 We fit the lapatinib-treated data to the model with 4 states based on our BIC-based model
313 selection, confirmed by cross-validation (Fig. 7a, S16f). Fitting revealed states of widely varying
314 persistence over generations, from less than a 0.01 probability of remaining in state 2 to a 0.94
315 probability of remaining in states 1 and 3 (Fig. 8a). Interestingly, states 2 and 4 formed a cycle
316 wherein the most probable transition was between these two states (Fig. 8a, S11).

317 Examining the phenotypes of each state revealed distinct drug responses. Lapatinib is an
318 EGFR/HER2 inhibitor that induces cell cycle arrest in G1 phase⁴³. Every state displayed a dose-
319 dependent increase in G1 phase lifetime with lapatinib treatment, and G1 effects were more
320 pronounced as compared to those involving S/G2 (Figs. 8b–i, S11). While the probability of
321 survival at the end of the cell cycle phase decreased at higher concentrations, very few cell death
322 events were observed (Figs. 8h–i, S11). Consequently, the chances of cell death likely have high
323 uncertainty at higher concentrations of lapatinib. States 2 and 4 were highly arrested in both G1
324 and S/G2 phase; in contrast, states 1 and 3 experienced little arrest in G1 and no arrest in S/G2
325 phase (Fig. 8f–g). Thus, cell states seemed to be primarily distinguished based on the degree of
326 lapatinib response. The cycle between states 2 and 4 seems to reflect the observation that cells
327 more highly arrested in G1 than G2/S give rise to cells that spend longer in G2/S than G1, and
328 vice versa (Fig. 8, S11).

329 **Figure 8: Lapatinib response is defined by phenotypically distinct stable and cyclical states.**
330 (a) State transition graph showing the probability of state transitions among the predicted states.
331 Transitions with less than a 0.03 probability have been removed. (b–e) A sample of lineage trees
332 after fitting the model and state assignment (control, 25 nM, 50 nM, and 250 nM). (f–g) The
333 log₁₀ of fit mean time of G1 and S/G2 phase durations for different concentrations. (h–i) The

334 Bernoulli parameter, indicating the probability of G1-to-S phase transition versus cell death (**h**),
335 and the probability of division versus cell death (**i**) for each concentration.

336 **Gemcitabine-treated populations are clustered into phase-specific responses**

337 Gemcitabine, as mentioned previously, is a chemotherapy agent that induces cell cycle arrest and
338 apoptosis in S phase by disrupting DNA repair. The AU565 cells were treated with 5, 10, and 30
339 nM of gemcitabine; model selection, confirmed by cross-validation, inferred 5 states in the
340 population (Figs. 7b, S16a/g). Examining the 5-state fit revealed relatively stable states 1, 3, and
341 4 (Fig. 9a–e, S12). States 2 and 5 formed a cycle with high rates of interconversion between one
342 another.

343 The phenotypic effects of drug treatment were evident in both cell cycle phases, though the most
344 affected phase differed between states (Fig. 9f–i). State 5 showed S/G2-specific arrest and
345 always resulted in cell death at the highest concentration (Fig. 9g, i, S12). Meanwhile, cells in
346 state 4 grew almost normally, with some cell death in G1 at the highest concentration (Fig. 9f–i).
347 At the highest concentration, state 3 represents the cells arrested at S/G2 that have not divided
348 even once, and state 5 is the representative of almost all of those undergoing cell death at S/G2
349 (Fig. 8i, S12).

350 Lastly, we wished to explore whether the phenotypic heterogeneity we observed was limited to
351 cancer cells or cytotoxic drug treatment. To determine this, we tracked non-tumorigenic
352 MCF10A breast cells. These cells are normally grown in the presence of epidermal growth factor
353 (EGF); we compared this condition to growth factor withdrawal (PBS) or rescue with hepatocyte
354 growth factor (HGF) or oncostatin M (OSM)⁴⁴. Each growth factor consistently promoted
355 proliferation on a population level compared to the PBS control, though with considerable inter-
356 and intra-lineage variation (Fig. S13). BIC-based model selection inferred the presence of 3
357 distinct states (Fig. 7c). Inspecting the model revealed generally more dynamic transitions
358 between states as compared to the AU565 experiments (Fig. S14). Due to the lack of growth
359 factors, most cells arrested in the PBS condition; the lack of cells completing their lifetime is the
360 reason for the division probability being 0.5 (Fig. S14g). State 1 was distinct in being relatively
361 less responsive to HGF and OSM treatments (Fig. S14a/f), while state 1 displayed higher rates of
362 cell death overall (Fig. S14a/g). In total, the tHMM model was effective in identifying subsets of
363 cells with divergent phenotypic responses to drug treatment alongside the relationships between
364 cells in the population.

365 **Figure 9: State-specific emissions of the gemcitabine-treated data.** (a) State transition graph
366 showing the probability of state transitions among the predicted states. The transitions with less
367 than a 0.03 probability have been removed. (b–e) A sample of lineage trees after fitting the
368 model and state assignment (control, 5 nM, 10 nM, and 30 nM). (f–g) The log₁₀ fit mean time of
369 G1 and S/G2 phase durations for different concentrations. (h–i) The Bernoulli parameter,
370 indicating the probability of G1-to-S phase transition versus cell death (**h**), and the probability of
371 division versus cell death (**i**) for each concentration.

372 **Discussion**

373 Heterogeneity and plasticity in cancer cells enables them to adapt in response to therapy. Even in
374 the absence of genetic mutations, other heritable variation serves as a substrate for selection^{45,46}.

375 In this paper, we introduced a tree-based hidden Markov model that clusters single cells from a
376 heterogeneous population based solely on their phenotypic traits and relationships. Model
377 benchmarking showed that it can provide accurate results using feasible experimental designs.
378 Of particular importance, the tHMM showed good sensitivity to subpopulations at lower
379 frequencies (Figs. S6, S7). Comparing the model to more standard clustering, the tHMM showed
380 that lineage information helps to identify cell states more accurately (Fig. 5). Using cross-
381 validation, we were able to show that accounting for cell inheritance allowed the model to
382 accurately predict unseen observations (Fig. S16). Several critical advancements in the current
383 work are a modular interface for using tHMM models with various phenotypes (Fig. 2), proper
384 censorship handling (Fig. 3), strategies for model evaluation (Figs. 7 & S16), and demonstrating
385 that such a model can be applied to study cancer heterogeneity at baseline and in response to
386 perturbation.

387 We used single cell lineage tracking data of AU565 cancer cells treated with lapatinib and
388 gemcitabine as a demonstration of the model. G1 and S/G2 cell cycle phase durations and cell
389 fate measurements were used as relevant cell phenotypes to quantify the anti-cancer effects of
390 these drugs. We were able to identify 4 and 5 distinct subpopulations within the lapatinib and
391 gemcitabine-treated data, respectively (Fig. 7). The phenotypic features of each state were
392 quantified in parallel (Figs. 8, 9). Lapatinib is known to inhibit cell proliferation by inhibiting
393 Akt/mTOR pathway activity, which is a key regulator of G1 phase progression⁴⁷. Similarly, our
394 analysis in the lapatinib-treated population indicated that cells, regardless of their state,
395 experienced a prolonged G1 phase, but individual states varied in their susceptibility. In
396 gemcitabine treated cells, we observed that most states were highly heritable, with more varied
397 phenotypic effects. This included cells that became arrested in S/G2 and underwent apoptosis
398 (state 5), cells that were selectively arrested in G1 (state 1), and cells that hardly responded to
399 drug treatment at all (state 4; Fig. 9). While gemcitabine canonically works by inducing cell
400 arrest in G2/S, previous work has characterized its effects on G1 phase by separating the effects
401 on both cell cycle phases⁴⁹. They similarly identified that G1 arrest was associated with cell
402 death, which is also evident in cells of state 3 where G1 arrest is seen alongside cell death in both
403 phases (Fig. 9f–i). Our results would further suggest that those cells with G1 effects are
404 molecularly and heritably distinct from those that are arrested in S/G2. MCF10A cells with
405 growth-factor-induced proliferation showed a very distinct pattern of variation, suggesting that
406 the phenotypic cell states identified by the model reflect a confluence of cell features and
407 treatment conditions (Fig. S13 & S14).

408 We present several lines of evidence supporting the accuracy of the model and existence of
409 heritable cell states. First, across a diverse array of benchmarking experiments, we show that the
410 model can derive accurate conclusions from synthetic data with properties like those we find in
411 the experimental measurements (e.g., Fig. 4). Through an informatic model selection scheme, we
412 find statistical evidence for the existence of multiple states (Figs. 6 & 7). Examining these cell
413 states, we find patterns consistent with the biological mechanisms of the compounds we used to
414 alter cell proliferation (Figs. 8 & 9). Reassuringly, we were able to confirm that the abundance of
415 cell states was consistent across experimental replicates, ruling out day-to-day variation between
416 experiments. Finally, and most critically, we showed that the model could more effectively
417 predict the behavior of masked cells with the inclusion of multiple cell states, and that this
418 prediction is dependent on allowing inheritance between cell generations (Fig. S16). While we
419 have considered the use of experimental control conditions, it is important to keep in mind that

420 the variation observed here arises both through external perturbation and natural variation within
421 the population. Consequently, we have not been able to identify a context in which one might
422 expect to *not* observe such variation. Also, while experiments in which distinct cell lines are
423 mixed can help to validate methods in which cell relationships are inferred, such as pseudotime
424 methods⁵⁰, the cell relationships are not modeled here because they are explicitly known through
425 the measured lineage relationships. Ultimately, experiments uncovering molecular markers and
426 mechanisms of these cell states will provide the best independent validation for their biological
427 significance.

428 Modeling advancements will further improve on our approach. Cells may express a continuum
429 of, rather than discrete, phenotypic states²². If this is the case, a continuous latent variable model
430 would lead to a refined view of the population-level heterogeneity. A discrete model like the one
431 used here should, however, still provide an accurate estimate by breaking up the continuous
432 state-space into discrete steps. Continuous latent variable models also have additional challenges
433 in implementation and interpretation^{51,52}. Careful handling of each states' phenotypic
434 distributions might also improve the model's accuracy and power to identify distinct states. For
435 example, the eventual fate of cells and their cell cycle durations are likely correlated which could
436 be handled through a multivariate distribution accounting for this covariance⁵³. This becomes
437 even more important with the incorporation of other phenotypic information such as migration,
438 cell shape, or other features, all of which are likely to be correlated to some extent.

439 Experimental advancements will improve the utility and accuracy of single cell analysis using
440 lineage information. Currently our experimental data is limited to 96 hours, covering up to
441 roughly five generations of cells. However, traits such as resistance may develop over more
442 generations and longer timescales^{13,54,55}. Longer data collection becomes challenging due to
443 factors such as phototoxicity and cell stress⁵⁶. Improved imaging modalities and experimental
444 platforms might allow for longer tracking experiments, with reduced phototoxicity, in more
445 physiologically representative environments such as engineered 3D extracellular matrix^{57,58}.
446 Currently, the model is agnostic as to whether the heterogeneity it identifies is pre-existing or
447 induced by drug treatment. Collecting data in which cells are tracked before and after drug
448 treatment, and after a wash-out, would help to link pre- and post-exposure phenotypes of cells⁵⁹.

449 While we have identified states that are phenotypically distinct subpopulations of cells, we
450 currently cannot comment on the molecular factors leading to these phenotypes. Molecular
451 barcoding has been a popular approach for identifying subpopulations of cells with genetic
452 predispositions toward unique phenotypes, but we do not expect it would identify the same
453 subpopulations as we do here⁵⁵, unlike in barcoding experiments, we do not see a bottleneck in
454 the clonality of cells that survive treatment, and rapid interconversion between states should
455 corrupt the relationship between ancestor phenotype and descendent molecular state^{3,12}.
456 However, we expect that single cell molecular analyses such as single cell tracking tied with
457 transcriptional profiling of the same cells at the end of the experiment, should allow us to align
458 molecular and phenotypic states in the same populations of cells⁶⁰. Such experiments would also
459 provide a common baseline by which to link lineage-based phenotypic analysis and various
460 snapshot measurements of the same cell population. In this way it should be able to pinpoint the
461 underlying molecular mechanisms driving distinct phenotypic responses.

462 In total, the pipeline developed here provides a unique approach for understanding the structure
463 of dynamic, heterogeneous tumor populations. By capturing the dynamics of state transitions, it

464 links single cell phenotypes to overall population behavior. Incorporating molecular
465 measurements, and a broader set of drug interventions, will then also help to identify means of
466 modulating state and overall population behavior. Ultimately, we expect this integrative view
467 will help to identify treatments alone and in combination that allow for population-level control
468 by affecting the growth of and transitions between individual cell subpopulations.

469 **Materials and Methods**

470 **Experimental cell lineage data**

471 Stable cell line creation, drug treatments, and tracking of AU565 and MCF10A cells were
472 performed as previously described in Gross et al⁶¹ and Gross et al⁴⁴, respectively. Briefly, AU565
473 cells were co-transfected with a transposase plasmid (Addgene #34879) and a donor plasmid that
474 drove expression of a nuclear localized mCherry, puromycin resistance, and a fragment of
475 HDHB fused to the clover fluorescent protein, which was used to track progression through the
476 cell cycle⁶². Cells stably expressing the nuclear and cell cycle reporter were selected for 7 days
477 with 0.75 µg/ml puromycin. The phase of the cells is determined based on whether the amount of
478 fluorescence is greater within nucleus or the cytoplasm⁶². As a result, the reporter signal is
479 invariant to changes in exposure and background. To track drug responses AU565 reporter cells
480 were plated into 24-well plates with fluorobrite media containing 10% FBS, glutamine, and
481 penicillin-streptomycin. 24 hours later fresh media containing escalating doses of lapatinib and
482 gemcitabine was added. MCF10A cells were cultured in growth media (DMEM/F12, 5% horse
483 serum, 20 ng/ml cholera toxin, 10 µg/ml insulin, and 1% Pen/Strep), grown to 50–80%
484 confluency, and detached with 0.05% trypsin-EDTA. 7 hours after seeding 75000 cells, they
485 were washed with PBS and the experiment media (DMEM/F12, 5% horse serum, 0.5 µg/ml
486 hydrocortisone, 100 ng/ml cholera toxin, and 1% Pen/Strep) was added to the 8 well-plates
487 which was followed by 18 hours of incubation. Afterward, cells were treated with growth factors
488 10 ng/ml EGF, 40 ng/ml HGF, and 10 ng/ml OSM in fresh experiment media. After drug
489 addition, plates were placed in the IncuCyte S3 and four image locations per treatment were
490 imaged every 30 minutes. AU565 were imaged for 96 hours and MCF10A cells for 48 hours.
491 After half the experiment times, fresh media and drugs/growth factors were added. Cell lineages
492 from the IncuCyte images were manually tracked in FIJI⁶³ to record cell division, death, and the
493 transition from G1 to S/G2 phase (in AU565). AU565 cells are non-motile and fewer than 4% of
494 cells were within one cell length of the image boundary, ensuring minimal sampling bias from
495 the microscopy field of view. Three biological replicates were collected and combined in the
496 final data set. To verify that results did not reflect batch effects, we checked that state
497 assignments were not enriched or depleted within a replicate.

498 **Lineage tree-based hidden Markov model**

499 The core assumption of a Markov chain is that the next state and current observations are only
500 dependent on the current state. Proof of the expressions below involving cell state assignment
501 (expectation step), including the upward recursion, downward recursion, and Viterbi algorithms,
502 can be found in Durand³³. All other model elements, including the emissions distribution fitting,
503 model evaluation strategies, and censorship corrections, were developed in this study.

504 **Basic model structure**

505 The initial probabilities of a cell being in state k are represented by the vector π that sums to 1:
506

$$507 \quad \pi_k = P(z_1 = k), \quad k \in \{1, \dots, K\} \quad (1)$$

508

509 where z indicates the state and K is the total number of states. The probability of state i
510 transitioning to state j is represented by the $K \times K$ matrix, T , in which each row sums to 1:
511

$$512 \quad T_{i,j} = T(z_i \rightarrow z_j) = P(z_j | z_i), \quad i, j \in \{1, \dots, K\} \quad (2)$$

513

514 The emission likelihood matrix, EL , is based on the cell observations. It is defined as the
515 probability of an observation conditioned on the cell being in a specific state:
516

$$517 \quad EL(n, k) = P(x_n = x | z_n = k) \quad (3)$$

518

519 where x_n is the observation for cell number n , with a total of N cells in a lineage. Separate
520 observations were assumed to be independent; for instance, cell fate is assumed to be
521 independent from the duration of each cell phase. This facilitates calculating the likelihood of
522 observations, because we can multiply the likelihood of all observations together for the total
523 likelihood.

524 **Assigning cell states (expectation step)**

525 *Upward recursion*

526 An *upward-downward* algorithm for calculating the probabilities in hidden Markov chains was
527 previously proposed by Erphaim and Merhav⁶⁴ which suffered from underflow. This problem
528 was originally solved by Levinson⁶⁵, where they adopted a heuristic-based scaling, and then was
529 improved by Devijver⁶⁶ where they introduced smooth probabilities. Durand³³, however, revised
530 this approach for hidden Markov trees to avoid underflow when calculating $P(Z|X)$ probability
531 matrices. To explain we need the following definitions:

- 532 • $p(n)$ is the parent cell of cell n , and $c(n)$ is the children of cell n .
- 533 • \bar{X} is the observation of the whole tree and \bar{X}_a is a subtree of \bar{X} which is rooted at cell a .
- 534 • \bar{Z} is the complete hidden state tree.
- 535 • $\bar{X}_{a/b}$ is the subtree rooted at a except for the subtree rooted at cell b , if \bar{X}_b is a subtree of
536 \bar{X}_a .

537 For the state prediction we start by calculating the marginal state distribution (MSD) matrix.
538 MSD is an $N \times K$ matrix that for each cell is marginalizing the transition probability over all

539 possible current states by traversing from root to leaf cells:
540

$$541 \quad MSD(n, k) = P(z_n = k) = \sum_i P(z_n = k | z_{n-1} = i) \times P(z_{n-1} = i) \quad (4)$$

542 During upward recursion, the flow of upward probabilities is calculated from leaf cells to the
543 root cells generation by generation. First, for leaf cells, the probabilities (β) are calculated by:
544

$$545 \quad \beta_n(k) = P(z_n = k | X_n = x_n) = \frac{EL(n, k) \times MSD(n, k)}{NF_l(n)} \quad (5)$$

546
547 in which X_n is the leaf cell's observation, and NF (Normalizing Factor) is an $N \times 1$ matrix that is
548 the marginal observation distribution. Since $\sum_k \beta_n(k) = 1$, we find the NF for leaf cells using:
549

$$550 \quad NF_l(n) = \sum_k EL(n, k) \times MSD(n, k) = P(X_n = x_n) \quad (6)$$

551
552 For non-leaf cells the values are given by:
553

$$554 \quad \beta_n(k) = P(z_n = k | \bar{X}_n = \bar{x}_n) = \frac{EL(n, k) \times MSD(n, k) \times \prod_{v \in c(n)} \beta_{n,v}(k)}{NF_{nl}(n)} \quad (7)$$

555
556 where we calculate the non-leaf NF using:

$$557 \quad NF_{nl}(n) = \sum_k [EL(n, k) \times MSD(n, k) \prod_{v \in c(n)} \beta_{n,v}(k)] \quad (8)$$

558 and linking β between parent-daughter cells is given by:
559

$$560 \quad \beta_{p(n),n}(k) = P(\bar{X}_n = \bar{x}_n | z_{p(n)} = k) = \sum_j \frac{\beta_n(j) \times T_{k,j}}{MSD(n, j)} \quad (9)$$

561
562 By recursing from leaf to root cells, the β and NF matrices are calculated as upward recursion.
563 The NF matrix gives a convenient expression for the observation log-likelihoods. For each root
564 cell we have:

565
$$P(\bar{X} = \bar{x}) = \prod_n \frac{P(\bar{X}_n = \bar{x}_n)}{\prod_{v \in c(n)} P(\bar{X}_v = \bar{x}_v)} = \sum_n NF(n) \quad n \in \{1, \dots, N\} \quad (10)$$

566

567 The overall model log-likelihood is given by the sum over root cells:

568
$$\log P(\bar{X} = \bar{x}) = \sum_n \log NF(n) \quad (11)$$

569

570 *Downward recursion*

571 For computing *downward recursion*, we need the following definition for each root cells:

572
$$\gamma_1(k) = P(z_1 = k | \bar{X}_1 = \bar{x}_1) = \beta_1(k) \quad (12)$$

573

574 The other cells follow in an $N \times K$ matrix by writing the conditional probabilities as the
575 summation over the joint probabilities of parent-daughter cells:

576
$$\gamma_n(k) = P(z_n = k | \bar{X}_1 = \bar{x}_1) = \frac{\beta_n(k)}{MSD(n, k)} \sum_i \frac{T_{i,k} \gamma_{p(n)}(i)}{\beta_{p(n),n}(i)} \quad (13)$$

577

578 *Viterbi algorithm*

579 Given a sequence of observations in a hidden Markov chain, the Viterbi algorithm is commonly
580 used to find the most likely sequence of states. Equivalently, here it returns the most likely
581 sequence of states of the cells in a lineage tree using upward and downward recursion³³.

582 The algorithm follows an upward recursion from leaf to root cells. We define δ , an $N \times K$
583 matrix:

584
$$\delta_n(k) = \max_{\bar{Z}_{c(n)}} \{P(\bar{X}_n = \bar{x}_n, \bar{Z}_{c(n)} = \bar{z}_{c(n)} | z_n = k)\} \quad (14)$$

585

586 and the links between parent-daughter cells as:

587
$$\delta_{p(n),n}(k) = \max_{\bar{Z}_n} \{P(\bar{X}_n = \bar{x}_n, \bar{Z}_n = \bar{z}_n | z_{p(n)} = k)\} = \max_k \{\delta_n(k) T_{k,k}\} \quad (15)$$

588

589 We initialize from the leaf cells as:

590
$$\delta_n(k) = P(X_n = x_n | z_n = k) = EL(n, k) \quad (16)$$

591

592 and for non-leaf cells use:

593
$$\delta_n(k) = \left[\prod_{v \in c(n)} \delta_{n,v}(k) \right] \times EL(n, k) \quad (17)$$

594

595 The probability of the optimal state tree corresponding to the observations tree, assuming root
596 cell is noted as cell 1, is then given by:

597
$$Z^* = \max_k \{ \delta_1(k) \pi_k \} \quad (18)$$

598

599 which arises from maximization over the conditional emission likelihood (EL) probabilities by
600 factoring out the root cells as the outer maximizing step over all possible states.

601 **Fitting the cell phenotypes (maximization step)**

602 In the maximization step, we find the maximum likelihood of the hidden Markov model
603 distribution parameters. We estimate the initial probabilities, the transition probability matrix,
604 and the parameters of the observation distributions. The maximum likelihood estimation of the
605 initial probabilities can be found from each state's representation among the root cells:
606

607
$$\pi_k^* = \gamma_1(k) \quad (19)$$

608

609 Similarly, the transition probability matrix is estimated by calculating the prevalence of each
610 transition across the lineage trees:
611

612
$$T_{i,j}^* = \frac{\sum_{n=1}^{N-1} \xi_n(i, j)}{\sum_{n=1}^{N-1} \gamma_n(i)} \quad (20)$$

613

614 where

615
$$\xi_n(i, j) = \left(\frac{\gamma_{p(n)}(i)}{\frac{\beta_n(i)}{MSD(n, i)} T(i, j)} \right)^T \times \frac{\beta_n(j)}{MSD(n, j)} \quad (21)$$

616

617 **Estimating emissions distribution parameters**

618 In the current study, we used two emissions distributions; first, a Bernoulli distribution for the
619 probability of each cell fate, either at the end of each phase or at the end of cell's lifetime;
620 second, a gamma distribution for the durations of each cell cycle phase or overall cell lifetime.
621 To estimate the distribution parameters after finding the cell state assignments, we calculated
622 their maximum likelihood estimation weighted by their proportional assignment to that state. The
623 initial and transition probabilities were shared across drug concentrations.

624 For estimating the Bernoulli distribution parameter for cell fate, we simply found the state
625 assignment-weighted sample mean of the observations. To estimate the gamma distribution
626 parameters, we fit all concentrations of each drug simultaneously and assumed that increasing
627 drug concentration had a unidirectional effect on the observed phenotype within each state. This
628 was implemented, using sequential least-squares programming (SLSQP)⁶⁷, through a linear
629 constraint on the scaling parameter of the gamma distributions between concentrations so that
630 higher concentrations had equal or greater average durations. The gamma distribution likelihood
631 fitting is a convex optimization problem, indicating that local optimization can arrive at the
632 globally optimal solution. Linear constraints do not change this property, and we confirmed
633 fitting with different starting points arrived at the same solution. We used censored estimators to
634 handle the effect of time censorship (explained below) in the duration distribution fitting. This
635 was done by fitting uncensored and censored observations to the complete and survival
636 distributions, respectively, and using the accumulated log-likelihood to estimate the distribution
637 parameters.

638 **Baum-Welch**

639 Since both the hidden states and model parameters are unknown, we applied expectation-
640 maximization (EM), known as the Baum-Welch algorithm for HMMs, to find both the model
641 parameters and cell states.

642 The EM algorithm consists of two steps: expectation and maximization. During expectation, the
643 probabilities of all cells being in specific states are calculated, such that for every cell and every
644 state we have $P(z_n = k | X_n)$ and $P(z_n = k, z_{n+1} = l | X_n)$. The expectation step is calculated
645 by the upward and downward recursion algorithms described above. In the maximization step,
646 described above, the distribution parameters of each state, the initial (π) probabilities, and the
647 transition probability (T) matrices are estimated, given the state assignments of each cell.

648 The EM algorithm is initialized by randomly assigning the cells to states using a Dirichlet
649 distribution. During fitting we iteratively switch between the expectation and maximization steps
650 and then calculate the likelihood. If the likelihood improves less than a set threshold, we take that
651 to indicate that to indicate convergence.

652 **Model evaluation**

653 To find the most likely number of states corresponding to the observations, the Bayesian
654 Information Criterion (BIC) was used⁶⁸. The BIC requires the number of cells and the degrees of
655 freedom, which we calculate using the number of independent parameters. Our model estimates
656 a k element initial probability vector, a $k \times k$ transition matrix, and a $k \times m$ matrix of state-wise
657 parameters where k is the number of states and m is the number of parameters associated with

658 observation distributions. For the phase-specific observation distributions we have a total of 6
659 parameters including 2 Bernoulli parameters and 2 pairs of shape and scale parameters for
660 gamma distribution. Since the row-sums for transition and initial probability matrices must be 1,
661 these values are not entirely independent. From distribution analysis of the phase lengths, we
662 realized the shape parameter of the gamma distribution remains constant over different
663 conditions, while the scale parameter changes. Therefore, the shape parameter was shared
664 between the populations treated with 4 different concentrations of the same compound. Each
665 condition therefore introduced 2 free parameters (1 Bernoulli parameter and 1 scale parameter).
666 For the MCF10A experiments, terminal fates and cell cycle durations were also assumed to be
667 Bernoulli- and gamma-distributed, respectively. The shape of cell lifetime was similarly shared
668 among the four conditions (PBS, EGF, HGF, and OSM).

669 The Wasserstein or Kantorovich–Rubinstein metric is a measure of distance between two
670 distributions. This metric was used to determine the difference between state emissions⁶⁹. An
671 analytical solution, the absolute value of the difference in distribution means, was used for the
672 gamma distribution.

673 **Model benchmarking**

674 We used emission distributions to represent the phenotypic characteristics of the cells within the
675 lineages. To create our synthetic data, we considered two possible options as our set of
676 observations throughout an experiment. In one case, we modeled the overall cell fate and cell
677 lifetime; in the second, we modeled the phase-specific fate and duration. In both, we used a
678 Bernoulli distribution for the fate outcomes and a gamma distribution for durations. The state
679 assignment accuracy was calculated using the Rand Index⁷⁰. The difference between true and
680 estimated probability matrices was assessed using the Frobenius norm, or the sum of each
681 element squared.

682 **Synthetic lineage data generation**

683 We generated synthetic lineage trees with K discrete states and N total number of cells for
684 benchmarking. Lineages were composed of two primary data structures: the state and emissions
685 trees. The state tree was randomly seeded with a root cell determined by the starting
686 probabilities, then expanded by randomly sampling transitions based on the transition probability
687 matrix. The lineages were extended by either increasing the number of initial cells, resulting in a
688 greater number of lineages (breadth), or by lengthening the experiment time resulting in each
689 lineage containing more cells (depth). After creating the tree of states with the desired number of
690 cells, the emission tree is built upon it. Emissions were randomly sampled from the distributions
691 for each cell's state. Finally, the effects of the emissions were applied to the tree when necessary.
692 If any cells died, their progeny were marked as unobserved by making their emissions equal to
693 NaN (Not a Number). If applicable, the effects of finite-duration experiments were also applied.
694 Cells existing outside of the experiment duration were marked as unobserved, and those crossing
695 the bounds of an experiment were marked as censored with duration clipped by the experiment.

696 **Time censorship**

697 Our phenotypic measurements include the cell fate (progression or cell death) and duration.
698 These measurements are made for each cell cycle phase (G1 or S/G2) in the case of AU565 cells
699 and the entire lifetime for MCF10A cells. These measurements can contain incomplete

700 information due to the bounds of an experiment. For instance, it is unknown when initial cells
701 present at the start of the experiment began their cell cycle. The same is true of the cells present
702 at the end of the experiment because we do not observe their end. Hence, a cell's lifetime and/or
703 fate may be partially observed. To ensure our synthetic data is a close reflection of experimental
704 data, we incorporated this effect in our synthetic data. Cells with lifetimes that extend beyond the
705 end of the experiment were marked as censored for the lifetime estimation.

706 Cell overall lifetime observations

707 The parameters are reflective of the cell phenotypes we observed with 5 nM lapatinib treatment.
708 Figures S1-S5 are based on these parameters. Each figure includes 100 trials.

709 Transition probability matrix: $T = \begin{bmatrix} 0.9 & 0.1 \\ 0.1 & 0.9 \end{bmatrix}$

710 The initial probability vector is then calculated as the stationary distribution of states from
711 transition probability matrix, satisfying $\pi = \pi * T$.

712 In this case, we have: $\pi = \begin{bmatrix} 0.5 \\ 0.5 \end{bmatrix}$

713 The same T and π were used for phase-specific emissions.

714 In the following table, “Bern p” refers to the Bernoulli parameter, the cell division probability at
715 the end of its lifetime which is equal to 1 - the probability of cell death at the end of its lifetime.
716 “Shape” and “Scale” refer to the gamma distribution parameters. The cells' lifetimes were fit to
717 gamma distributions.

718 Table 1. State distribution parameters.

State	Bern p	Shape	Scale
State 1	0.99	8	6
State 2	0.75	8	1

719 Cell cycle phase-specific observations

720 The synthetic data used in Figures 3, 4, S8–S10 were created based on the following parameters.
721 These parameters are based on estimations from AU565 cells treated with 5 nM lapatinib. Each
722 figure includes 100 trials.

723 In the following table, “G1 bern” and “S/G2 bern” are the cell division probabilities at the end of
724 G1 and S/G2 phase, respectively. The “G1 shape” and “G1 scale” are the gamma distribution
725 parameters of G1 phase lengths. “S/G2 shape” and “S/G2 scale” are the gamma distribution
726 parameters of S/G2 phase lengths.

727 Table 2. State distribution parameters.

State	G1 bern	S/G2 bern	G1 shape	G1 scale	S/G2 shape	S/G2 scale
State 1	0.99	0.95	8	7	4	2
State 2	0.95	0.9	6	4	3	5

728

729 To benchmark the model with 5 states, we simulated 25–500 lineages, each with up to 31 cells,
 730 to create a population with 5 states. Like with the experimental data, we assumed the experiment
 731 ends after 96 hours and censored the cells’ observations accordingly. The model parameters,
 732 including the distribution parameters, transition probabilities, and initial probabilities, are listed
 733 below. The analysis results are shown in Figure S14.

734 Table 3. State distribution parameters.

State	G1 bern	S/G2 bern	G1 shape	G1 scale	S/G2 shape	S/G2 scale
State 1	0.7	0.99	250	0.2	50	0.1
State 2	0.95	0.9	200	0.2	100	0.1
State 3	0.9	0.85	150	0.2	150	0.1
State 4	0.99	0.75	100	0.2	200	0.1
State 5	0.99	0.75	50	0.2	250	0.1

735

$$T = \begin{bmatrix} \mathbf{0.6} & \mathbf{0.1} & \mathbf{0.1} & \mathbf{0.1} & \mathbf{0.1} \\ \mathbf{0.05} & \mathbf{0.8} & \mathbf{0.05} & \mathbf{0.05} & \mathbf{0.05} \\ \mathbf{0.01} & \mathbf{0.1} & \mathbf{0.7} & \mathbf{0.09} & \mathbf{0.1} \\ \mathbf{0.1} & \mathbf{0.1} & \mathbf{0.05} & \mathbf{0.7} & \mathbf{0.05} \\ \mathbf{0.1} & \mathbf{0.1} & \mathbf{0.05} & \mathbf{0.05} & \mathbf{0.7} \end{bmatrix} \quad (22)$$

736

$$\pi = \begin{bmatrix} \mathbf{0.13} \\ \mathbf{0.33} \\ \mathbf{0.16} \\ \mathbf{0.18} \\ \mathbf{0.18} \end{bmatrix} \quad (23)$$

737 Figure 6 uses the first 4 states of table 3 as the parameter set for the emissions matrix to simulate
 738 varying state numbers in the BIC calculation.

739 Varying emission differences

740 To create synthetic data with subpopulations of varying dissimilarity (Fig. 5), we use the phase-
 741 specific parameters, with the values for the G1 phase gamma scale parameter for state 1 varying
 742 over [4, 20]. This results in an increase in the Wasserstein distance between the two cell states,
 743 allowing us to measure state assignment accuracy for different dissimilarity amounts between the
 744 two states. Likewise, for Figures S6 and S7, we simulated the overall cell lifetime and varied the
 745 gamma distribution scale parameter from 1 to 8 for state 1.

746 Optimal baseline classifier

747 To compare the tHMM with a classifier that ignores heritability, we manually calculated the
 748 optimal classification boundary between the gamma distributions for state 1 and state 2. The best
 749 choice of classification boundary between two gamma distributions is the point at which the
 750 likelihood of the random variable, x , is equal between the two distributions:

$$751 \quad p(x | G(k_1, \theta_1)) = p(x | G(k_2, \theta_2)) \quad (24)$$

752 where k_1 , θ_1 , k_2 , and θ_2 are the shape and scale parameters of the gamma distribution
753 corresponding to state 1 and 2, respectively. The shape parameter was shared between the two
754 distributions. Consequently, this can be simplified to:

$$755 \quad x = \frac{k \ln \frac{\theta_2}{\theta_1}}{\frac{1}{\theta_1} - \frac{1}{\theta_2}} \quad (25)$$

756 We assigned the classification labels to the observations using this classification boundary,
757 which formed the baseline accuracy shown in Fig. 5c. As states 1 and 2 are identical at the very
758 first point, we used the distribution mean ($k \times \theta$) as the threshold.

759 **Cross-validation**

760 To split the lineage data into train and test sets, we randomly selected 20% of cells from each
761 condition and masked their observations such that they would not contribute to the fitting
762 process. This was performed by setting the log-likelihood of the masked cells' observations to be
763 uniformly zero for all the states. During the Baum-Welch fitting, the algorithm estimates the
764 parameters using only the training cells. However, during the expectation step, the state of
765 masked cells is still inferred via information about their relatives. After the fitting converges, we
766 calculate the log-likelihood of the test cells' observations given their state assignments. This is
767 accumulated into an overall likelihood of the held-out observations given the tHMM state
768 assignments and fit.

769 To test this cross-validation scheme's ability to determine the optimum number of states for a
770 cell population, we created synthetic populations with 2–5 true states. States 1– n were used,
771 where n is the number of true states, to generate data that is like the experimental data. The state
772 observation distributions shown in Table 2. The transition probabilities were generated by adding
773 0.1 elementwise to the identity matrix and then normalizing it. The initial probabilities for all
774 states were equal. Fitting was performed with models including 1–7 states. The optimum number
775 of states was taken to be the smallest number of states at which the log-likelihood plateaus.

776 **Lowess trendline**

777 Locally Weighted Scatterplot Smoothing (Lowess) was used to provide the trendlines in the
778 figures with repeated model runs.

779 **Code availability**

780 All analysis were implemented in Python v3.9 and can be found at <https://github.com/meyer-lab/tHMM>.
781

782 **Data availability**

783 The experimental lineage data for AU565 and MCF10A cell lines can be found at
784 <https://github.com/meyer-lab/tHMM>

785 The synthetic data from which we plotted Fig. 3c–d, Fig. 4b–g, Fig. 5c, Fig. 6, and Fig. 7 uses
786 the code in the file named after the corresponding figure number. Data used in Figs. 7a–b, 8, and

787 9 uses the AU565 cell line experimental lineage data, and figures 7c and 10 use the MCF10A
788 lineage data.

789 The cell lines used in this study (AU565, MCF10A) can be made available upon request.

790 **Statistics and Reproducibility**

791 The experiments were repeated in three independent biological replicates and yielded similar
792 results.

793 **Acknowledgements**

794 This work was supported by the Jayne Koskinas Ted Giovanis Foundation for Health and Policy,
795 NIH U01-CA215709 (A.S.M.), NIH U54-CA209988 (L.M.H.), NIH U54-HG008100 (L.M.H.).

796 The authors thank Scott Taylor for his critical feedback that helped to improve the manuscript.

797 The authors thank Ali Farhat, Adam Weiner, and Nikan Namiri for early exploratory work.

798 **Competing financial interests:** The authors declare no competing financial interests.

799 **Author contributions statement**

800 A.S.M. and L.M.H. conceived of the study; A.S.M. conceived of the model; A.S.M, F.M., S.V.
801 designed model; A.S.M., F.M., J.L., L.K., S.V. performed computational experiments; S.M.G.
802 performed the experiments; F.M., J.L., L.K., S.M.G. conducted data analysis; A.S.M. and
803 L.M.H. supervised the research; all authors wrote the paper.

804 **References**




- 805 1. Di Maio, M. *et al.* Chemotherapy-induced neutropenia and treatment efficacy in advanced
806 non-small-cell lung cancer: a pooled analysis of three randomised trials. *Lancet Oncol.* **6**,
807 669–677 (2005).
- 808 2. De Roock, W. *et al.* Effects of KRAS, BRAF, NRAS, and PIK3CA mutations on the
809 efficacy of cetuximab plus chemotherapy in chemotherapy-refractory metastatic colorectal
810 cancer: a retrospective consortium analysis. *Lancet Oncol.* **11**, 753–762 (2010).
- 811 3. Sharma, S. V. *et al.* A Chromatin-Mediated Reversible Drug-Tolerant State in Cancer Cell
812 Subpopulations. *Cell* **141**, 69–80 (2010).
- 813 4. Spencer, S. L., Gaudet, S., Albeck, J. G., Burke, J. M. & Sorger, P. K. Non-genetic origins
814 of cell-to-cell variability in TRAIL-induced apoptosis. *Nature* **459**, 428–432 (2009).
- 815 5. Sigal, A. *et al.* Variability and memory of protein levels in human cells. *Nature* **444**, 643–
816 646 (2006).
- 817 6. Feinberg, A. P., Koldobskiy, M. A. & Göndör, A. Epigenetic modulators, modifiers and
818 mediators in cancer aetiology and progression. *Nat. Rev. Genet.* **17**, 284–299 (2016).
- 819 7. Falkenberg, K. J. & Johnstone, R. W. Histone deacetylases and their inhibitors in cancer,
820 neurological diseases and immune disorders. *Nat. Rev. Drug Discov.* **13**, 673–691 (2014).
- 821 8. Inde, Z. & Dixon, S. J. The impact of non-genetic heterogeneity on cancer cell death. *Crit.*

- 822 *Rev. Biochem. Mol. Biol.* **53**, 99–114 (2018).
- 823 9. Pisco, A. O. & Huang, S. Non-genetic cancer cell plasticity and therapy-induced stemness
824 in tumour relapse: ‘What does not kill me strengthens me’. *Br. J. Cancer* **112**, 1725–1732
825 (2015).
- 826 10. Ghandi, M. *et al.* Next-generation characterization of the Cancer Cell Line Encyclopedia.
827 *Nat. 2019 5697757* **569**, 503–508 (2019).
- 828 11. Barretina, J. *et al.* The Cancer Cell Line Encyclopedia enables predictive modelling of
829 anticancer drug sensitivity. *Nat. 2012 4837391* **483**, 603–607 (2012).
- 830 12. Gupta, P. B. *et al.* Stochastic State Transitions Give Rise to Phenotypic Equilibrium in
831 Populations of Cancer Cells. *Cell* **146**, 633–644 (2011).
- 832 13. Shaffer, S. M. *et al.* Rare cell variability and drug-induced reprogramming as a mode of
833 cancer drug resistance. *Nature* **546**, 431–435 (2017).
- 834 14. Gett, A. V., Sallusto, F., Lanzavecchia, A. & Geginat, J. T cell fitness determined by signal
835 strength. *Nat. Immunol.* **4**, 355–360 (2003).
- 836 15. Arai, T. *et al.* Tumor Doubling Time and Prognosis in Lung Cancer Patients: Evaluation
837 from Chest Films and Clinical Follow-up Study. *Jpn. J. Clin. Oncol.* **24**, 199–204 (1994).
- 838 16. Bourhis, J. *et al.* Potential doubling time and clinical outcome in head and neck squamous
839 cell carcinoma treated with 70 GY in 7 weeks. *Int. J. Radiat. Oncol.* **35**, 471–476 (1996).
- 840 17. Yachida, S. *et al.* Distant metastasis occurs late during the genetic evolution of pancreatic
841 cancer. *Nature* **467**, 1114–1117 (2010).
- 842 18. Huang, D. *et al.* High-Speed Live-Cell Interferometry: A New Method for Quantifying
843 Tumor Drug Resistance and Heterogeneity. *Anal. Chem.* **90**, 3299–3306 (2018).
- 844 19. Tyson, D. R., Garbett, S. P., Frick, P. L. & Quaranta, V. Fractional proliferation: A
845 method to deconvolve cell population dynamics from single-cell data. *Nat. Methods*
846 (2012) doi:10.1038/nmeth.2138.
- 847 20. O’Connor, J. P. B. *et al.* Imaging intratumor heterogeneity: role in therapy response,
848 resistance, and clinical outcome. *Clin. Cancer Res.* **21**, 249–257 (2015).
- 849 21. Chen, K. *et al.* Phenotypically supervised single-cell sequencing parses within-cell-type
850 heterogeneity. *iScience* **24**, 101991 (2020).
- 851 22. Kuchen, E. E., Becker, N. B., Claudino, N. & Höfer, T. Hidden long-range memories of
852 growth and cycle speed correlate cell cycles in lineage trees. *Elife* **9**, e51002 (2020).
- 853 23. Mitchell, S., Roy, K., Zangle, T. A. & Hoffmann, A. Nongenetic origins of cell-to-cell
854 variability in B lymphocyte proliferation. *Proc. Natl. Acad. Sci. U. S. A.* **115**, E2888–
855 E2897 (2018).
- 856 24. Young, A. I., Benonisdottir, S., Przeworski, M. & Kong, A. Deconstructing the sources of

- 857 genotype-phenotype associations in humans. *Science* **365**, 1396–1400 (2019).
- 858 25. Brumpton, B. *et al.* Within-family studies for Mendelian randomization: avoiding
859 dynastic, assortative mating, and population stratification biases. (2019)
860 doi:10.1101/602516.
- 861 26. Concannon, P. *et al.* Genome-wide scan for linkage to type 1 diabetes in 2,496 multiplex
862 families from the Type 1 Diabetes Genetics Consortium. *Diabetes* **58**, 1018–1022 (2009).
- 863 27. Concannon, P., Rich, S. S. & Nepom, G. T. Genetics of Type 1A Diabetes. *N. Engl. J.*
864 *Med.* **360**, 1646–1654 (2009).
- 865 28. Quinn, J. J. *et al.* Single-cell lineages reveal the rates, routes, and drivers of metastasis in
866 cancer xenografts. *Science* **371**, eabc1944 (2021).
- 867 29. Choo, K. H., Tong, J. C. & Zhang, L. Recent applications of Hidden Markov Models in
868 computational biology. *Genomics. Proteomics Bioinformatics* **2**, 84–96 (2004).
- 869 30. Rabiner, L. R. A tutorial on hidden Markov models and selected applications in speech
870 recognition. *Proc. IEEE* **77**, 257–286 (1989).
- 871 31. Yanagawa, M. *et al.* Single-molecule diffusion-based estimation of ligand effects on G
872 protein-coupled receptors. *Sci. Signal.* **11**, (2018).
- 873 32. Crouse, M. S., Nowak, R. D. & Baraniuk, R. G. Wavelet-based statistical signal
874 processing using hidden Markov models. *IEEE Trans. Signal Process.* **46**, 886–902
875 (1998).
- 876 33. Durand, J. B., Gonçalves, P. & Guédon, Y. Computational methods for hidden Markov
877 tree models-An application to wavelet trees. *IEEE Trans. Signal Process.* **52**, 2551–2560
878 (2004).
- 879 34. Choi, H. & Baraniuk, R. G. Multiscale image segmentation using wavelet-domain hidden
880 Markov models. *IEEE Trans. Image Process.* **10**, 1309–1321 (2001).
- 881 35. Bykova, N. A., Favorov, A. V & Mironov, A. A. Hidden Markov models for evolution
882 and comparative genomics analysis. *PLoS One* **8**, e65012–e65012 (2013).
- 883 36. Olariu, V. *et al.* Modified variational Bayes EM estimation of hidden Markov tree model
884 of cell lineages. *Bioinformatics* **25**, 2824–2830 (2009).
- 885 37. Nakashima, S., Sughiyama, Y. & Kobayashi, T. J. Lineage EM algorithm for inferring
886 latent states from cellular lineage trees. *Bioinformatics* **36**, 2829–2838 (2020).
- 887 38. Lund, B., Kristjansen, P. E. G. & Hansen, H. H. Clinical and preclinical activity of 2',2'-
888 difluorodeoxycytidine (gemcitabine). *Cancer Treat. Rev.* (1993) doi:10.1016/0305-
889 7372(93)90026-N.
- 890 39. Bolstad, B. M. Comparing some iterative methods of parameter estimation for censored
891 gamma data. (1998).


- 892 40. Chao, H. X. *et al.* Evidence that the human cell cycle is a series of uncoupled, memoryless
893 phases. *Mol. Syst. Biol.* **15**, e8604–e8604 (2019).
- 894 41. Lee, J. A. *et al.* Microbial phenotypic heterogeneity in response to a metabolic toxin:
895 Continuous, dynamically shifting distribution of formaldehyde tolerance in
896 *Methylobacterium extorquens* populations. *PLoS Genet.* **15**, e1008458–e1008458 (2019).
- 897 42. van Boxtel, C., van Heerden, J. H., Nordholt, N., Schmidt, P. & Bruggeman, F. J. Taking
898 chances and making mistakes: non-genetic phenotypic heterogeneity and its consequences
899 for surviving in dynamic environments. *J. R. Soc. Interface* **14**, 20170141 (2017).
- 900 43. Tang, L., Wang, Y., Strom, A., Gustafsson, J. A. & Guan, X. Lapatinib induces p27Kip1-
901 dependent G1 arrest through both transcriptional and post-translational mechanisms. *Cell*
902 *Cycle* (2013) doi:10.4161/cc.25728.
- 903 44. Gross, S. M. *et al.* A LINCS microenvironment perturbation resource for integrative
904 assessment of ligand-mediated molecular and phenotypic responses. (2021)
905 doi:10.1101/2021.08.06.455429.
- 906 45. Brock, A., Chang, H. & Huang, S. Non-genetic heterogeneity — a mutation-independent
907 driving force for the somatic evolution of tumours. *Nat. Rev. Genet.* **10**, 336–342 (2009).
- 908 46. Fan, Y. & Meyer, T. Molecular control of cell density-mediated exit to quiescence. *Cell*
909 *Rep.* **36**, 109436 (2021).
- 910 47. Zhu, X. *et al.* Autophagy stimulates apoptosis in HER2-overexpressing breast cancers
911 treated by lapatinib. *J. Cell. Biochem.* **114**, 2643–2653 (2013).
- 912 48. Richards, R. *et al.* Drug antagonism and single-agent dominance result from differences in
913 death kinetics. *Nat. Chem. Biol.* **16**, 791–800 (2020).
- 914 49. Johnson, T. I. *et al.* Quantifying cell cycle-dependent drug sensitivities in cancer using a
915 high throughput synchronisation and screening approach. *EBioMedicine* **68**, 103396
916 (2021).
- 917 50. Campbell, K. R. & Yau, C. Uncovering pseudotemporal trajectories with covariates from
918 single cell and bulk expression data. *Nat. Commun.* **2018 91 9**, 1–12 (2018).
- 919 51. Lartillot, N. A phylogenetic Kalman filter for ancestral trait reconstruction using
920 molecular data. *Bioinformatics* **30**, 488–496 (2013).
- 921 52. Ding, J. *et al.* Reconstructing differentiation networks and their regulation from time
922 series single-cell expression data. *Genome Res.* **28**, 383–395 (2018).
- 923 53. Tentner, A. R. *et al.* Combined experimental and computational analysis of DNA damage
924 signaling reveals context-dependent roles for Erk in apoptosis and G1/S arrest after
925 genotoxic stress. *Mol. Syst. Biol.* **8**, 568 (2012).
- 926 54. Hata, A. N. *et al.* Tumor cells can follow distinct evolutionary paths to become resistant to
927 epidermal growth factor receptor inhibition. *Nat. Med.* **22**, 262–269 (2016).

- 928 55. Bhang, H. C. *et al.* Studying clonal dynamics in response to cancer therapy using high-
929 complexity barcoding. *Nat. Med.* **21**, 440–448 (2015).
- 930 56. Burke, R. T. & Orth, J. D. Through the Looking Glass: Time-lapse Microscopy and
931 Longitudinal Tracking of Single Cells to Study Anti-cancer Therapeutics. *J. Vis. Exp.*
932 53994 (2016) doi:10.3791/53994.
- 933 57. Han, K. *et al.* CRISPR screens in cancer spheroids identify 3D growth-specific
934 vulnerabilities. *Nature* **580**, 136–141 (2020).
- 935 58. Schwartz, A. D. *et al.* A biomaterial screening approach reveals microenvironmental
936 mechanisms of drug resistance. *Integr. Biol. (Camb)*. **9**, 912–924 (2017).
- 937 59. Emert, B. L. *et al.* Variability within rare cell states enables multiple paths toward drug
938 resistance. *Nat. Biotechnol.* **39**, 865–876 (2021).
- 939 60. Papalexli, E. & Satija, R. Single-cell RNA sequencing to explore immune cell
940 heterogeneity. *Nat. Rev. Immunol.* **18**, 35–45 (2017).
- 941 61. Gross, S. M. *et al.* Analysis and modeling of cancer drug responses using cell cycle phase-
942 specific rate effects. *bioRxiv* 2020.07.24.219907 (2021) doi:10.1101/2020.07.24.219907.
- 943 62. SL, S. *et al.* The proliferation-quiescence decision is controlled by a bifurcation in CDK2
944 activity at mitotic exit. *Cell* **155**, 369 (2013).
- 945 63. Meijering, E., Dzyubachyk, O. & Smal, I. Methods for cell and particle tracking. in
946 *Methods in Enzymology* (2012). doi:10.1016/B978-0-12-391857-4.00009-4.
- 947 64. Ephraim, Y. & Merhav, N. Hidden Markov processes. *IEEE Trans. Inf. Theory* **48**, 1518–
948 1569 (2002).
- 949 65. Levinson, S. E., Rabiner, L. R. & Sondhi, M. M. An Introduction to the Application of the
950 Theory of Probabilistic Functions of a Markov Process to Automatic Speech Recognition.
951 *Bell Syst. Tech. J.* **62**, 1035–1074 (1983).
- 952 66. Devijver, P. A. Baum’s forward-backward algorithm revisited. *Pattern Recognit. Lett.* **3**,
953 369–373 (1985).
- 954 67. Virtanen, P. *et al.* SciPy 1.0: fundamental algorithms for scientific computing in Python.
955 *Nat. Methods* 2020 173 **17**, 261–272 (2020).
- 956 68. Neath, A. A. & Cavanaugh, J. E. The Bayesian information criterion: background,
957 derivation, and applications. *WIREs Comput. Stat.* **4**, 199–203 (2011).
- 958 69. Vallender, S. S. Calculation of the Wasserstein Distance Between Probability
959 Distributions on the Line. *Theory Probab. & Its Appl.* **18**, 784–786 (1974).
- 960 70. Rand, W. M. Objective Criteria for the Evaluation of Clustering Methods. *J. Am. Stat.*
961 *Assoc.* **66**, 846–850 (1971).



a

A lineage tree object stores information on each cell

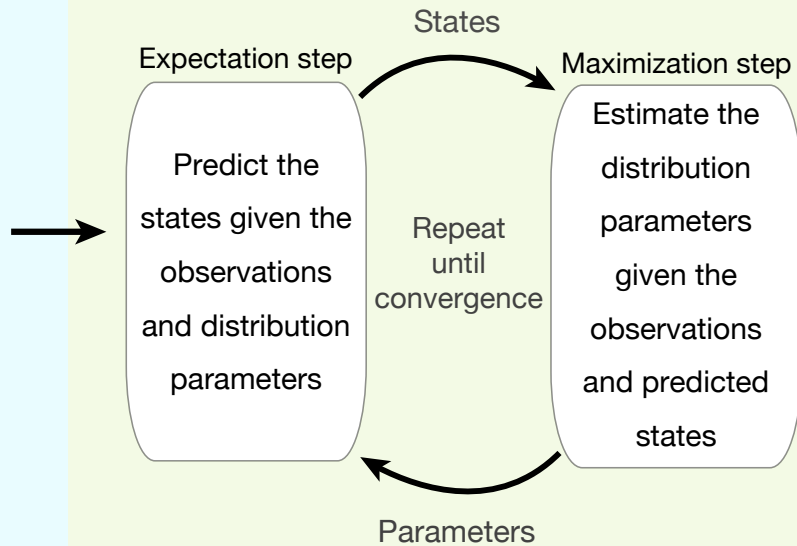


Observations

- Survived or died in G1
- Survived or died in S/G2
- Duration of G1 phase
- Duration of S/G2 phase
- Any new observed phenotype

b


tHMM

**c**


State transition probabilities




State distribution parameters


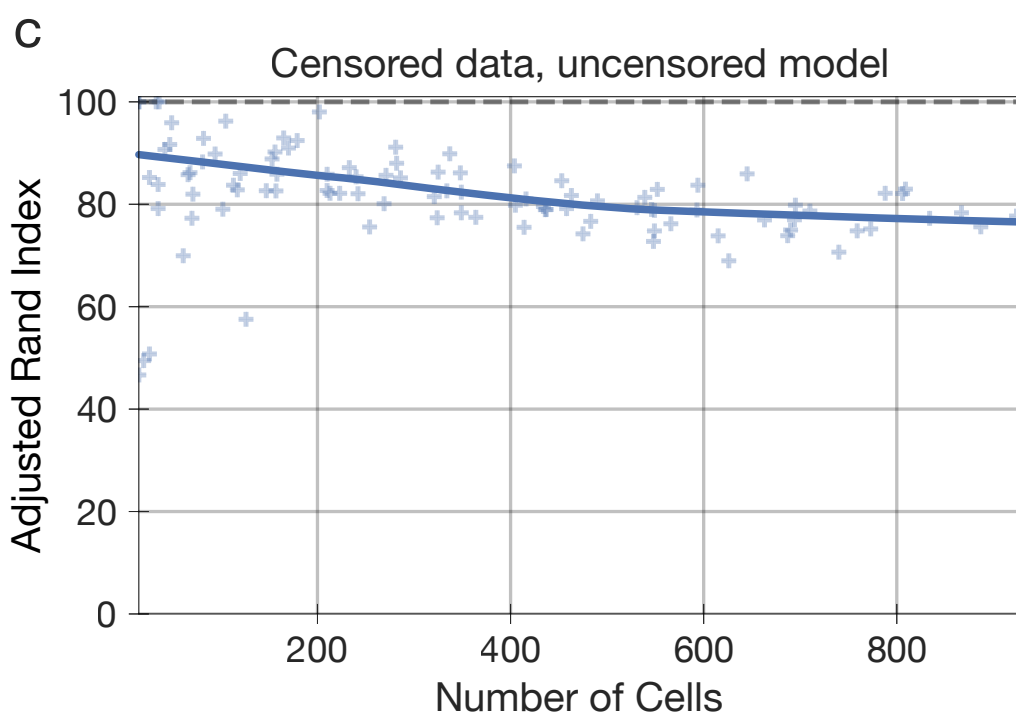




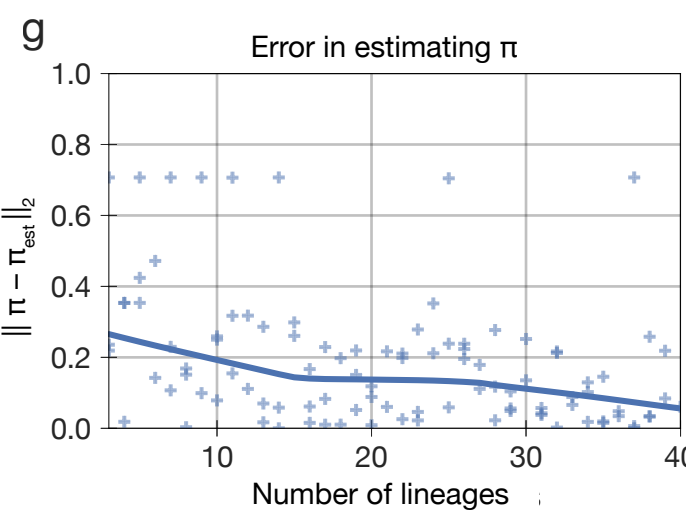
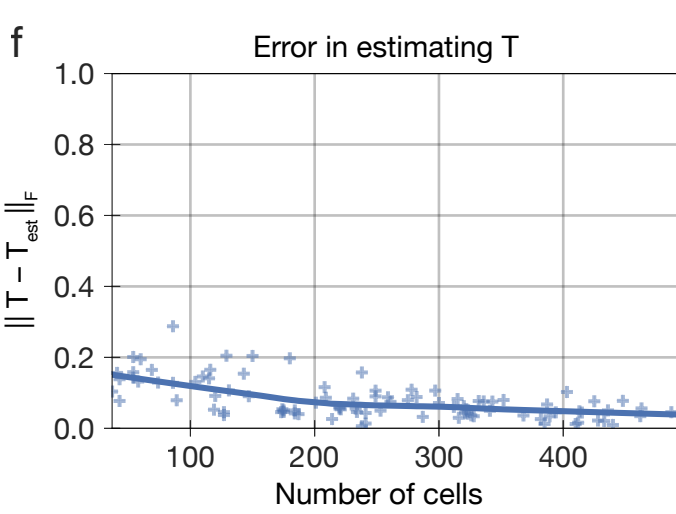
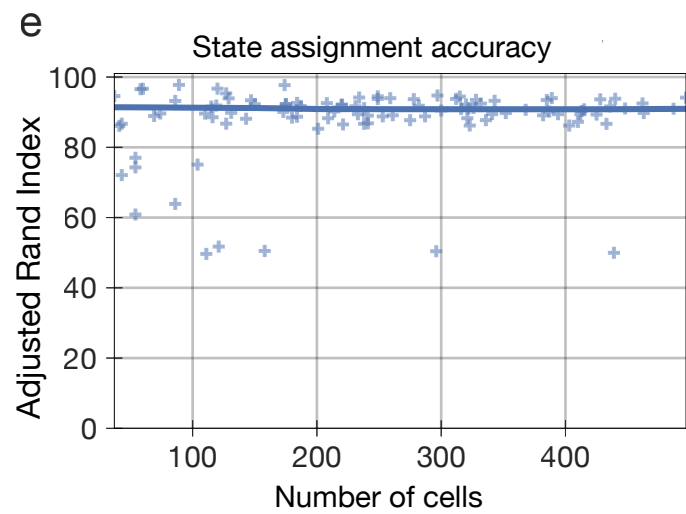

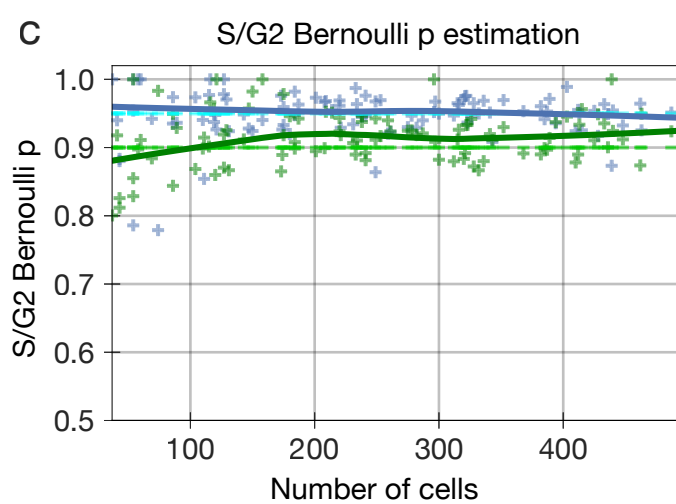
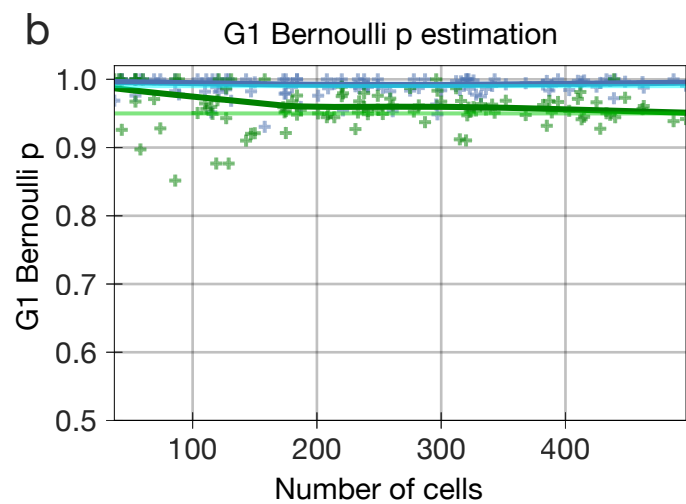
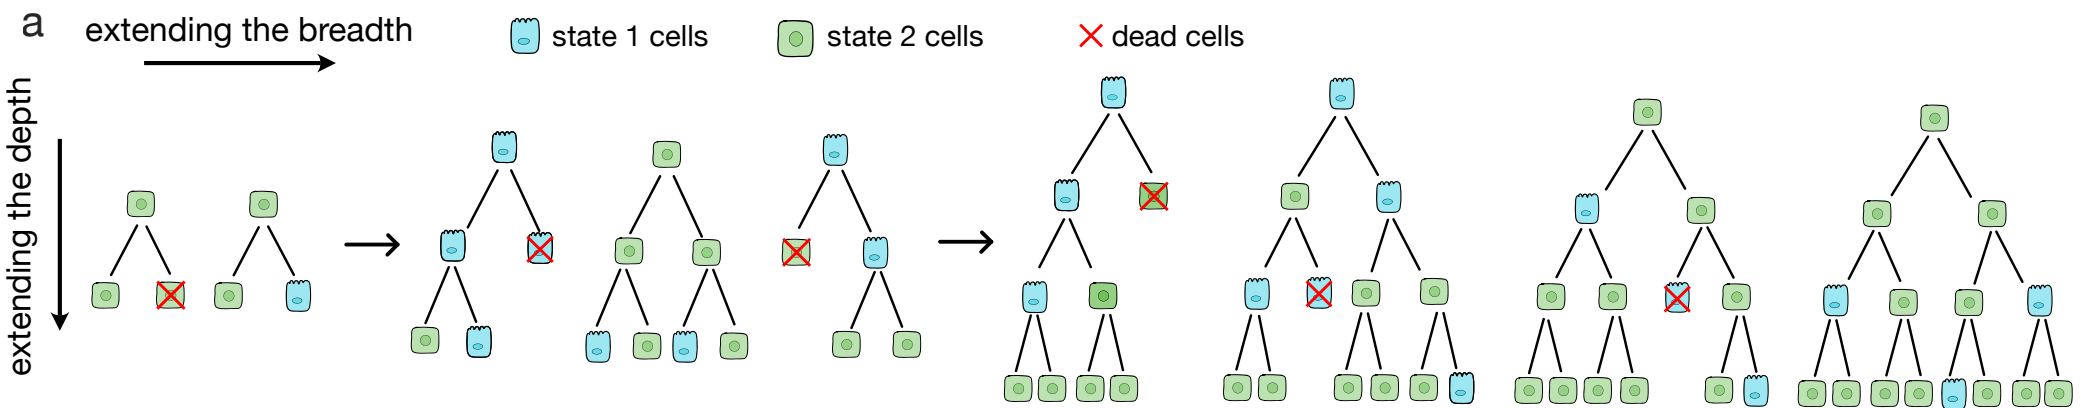
Cell state assignments


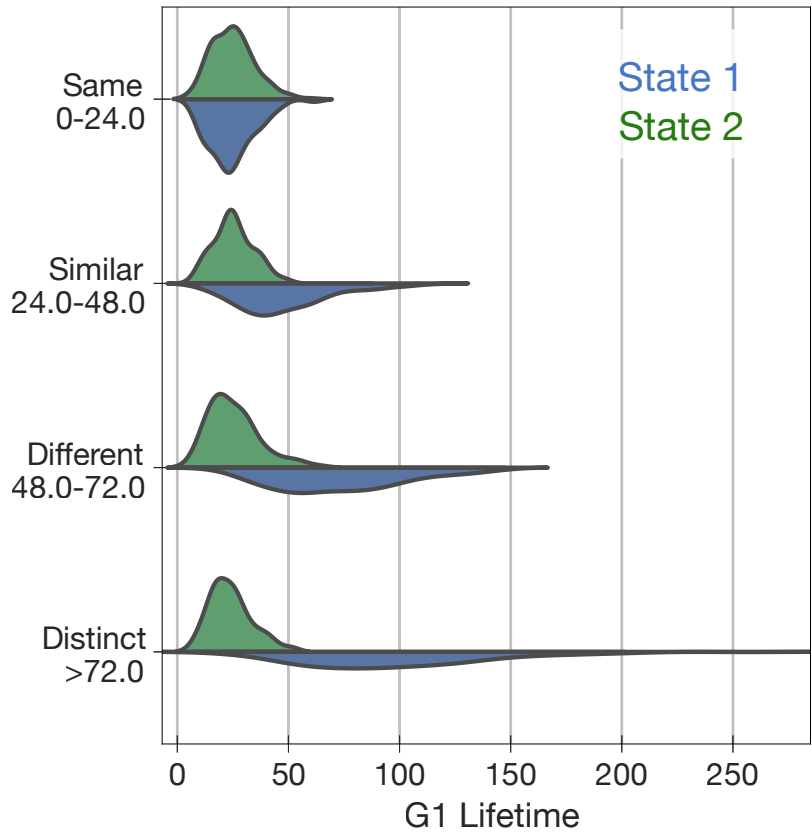
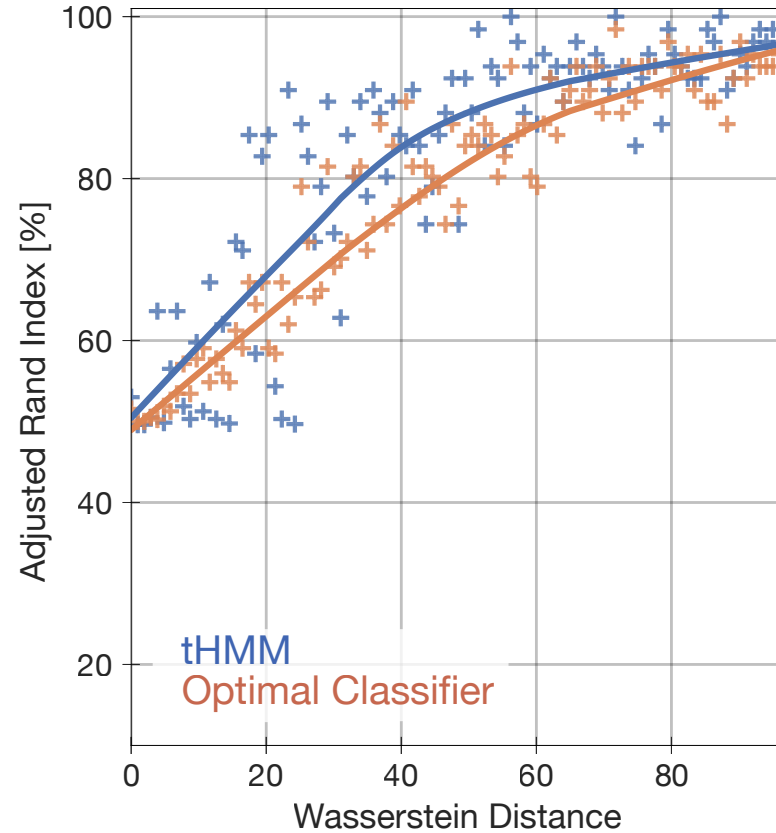


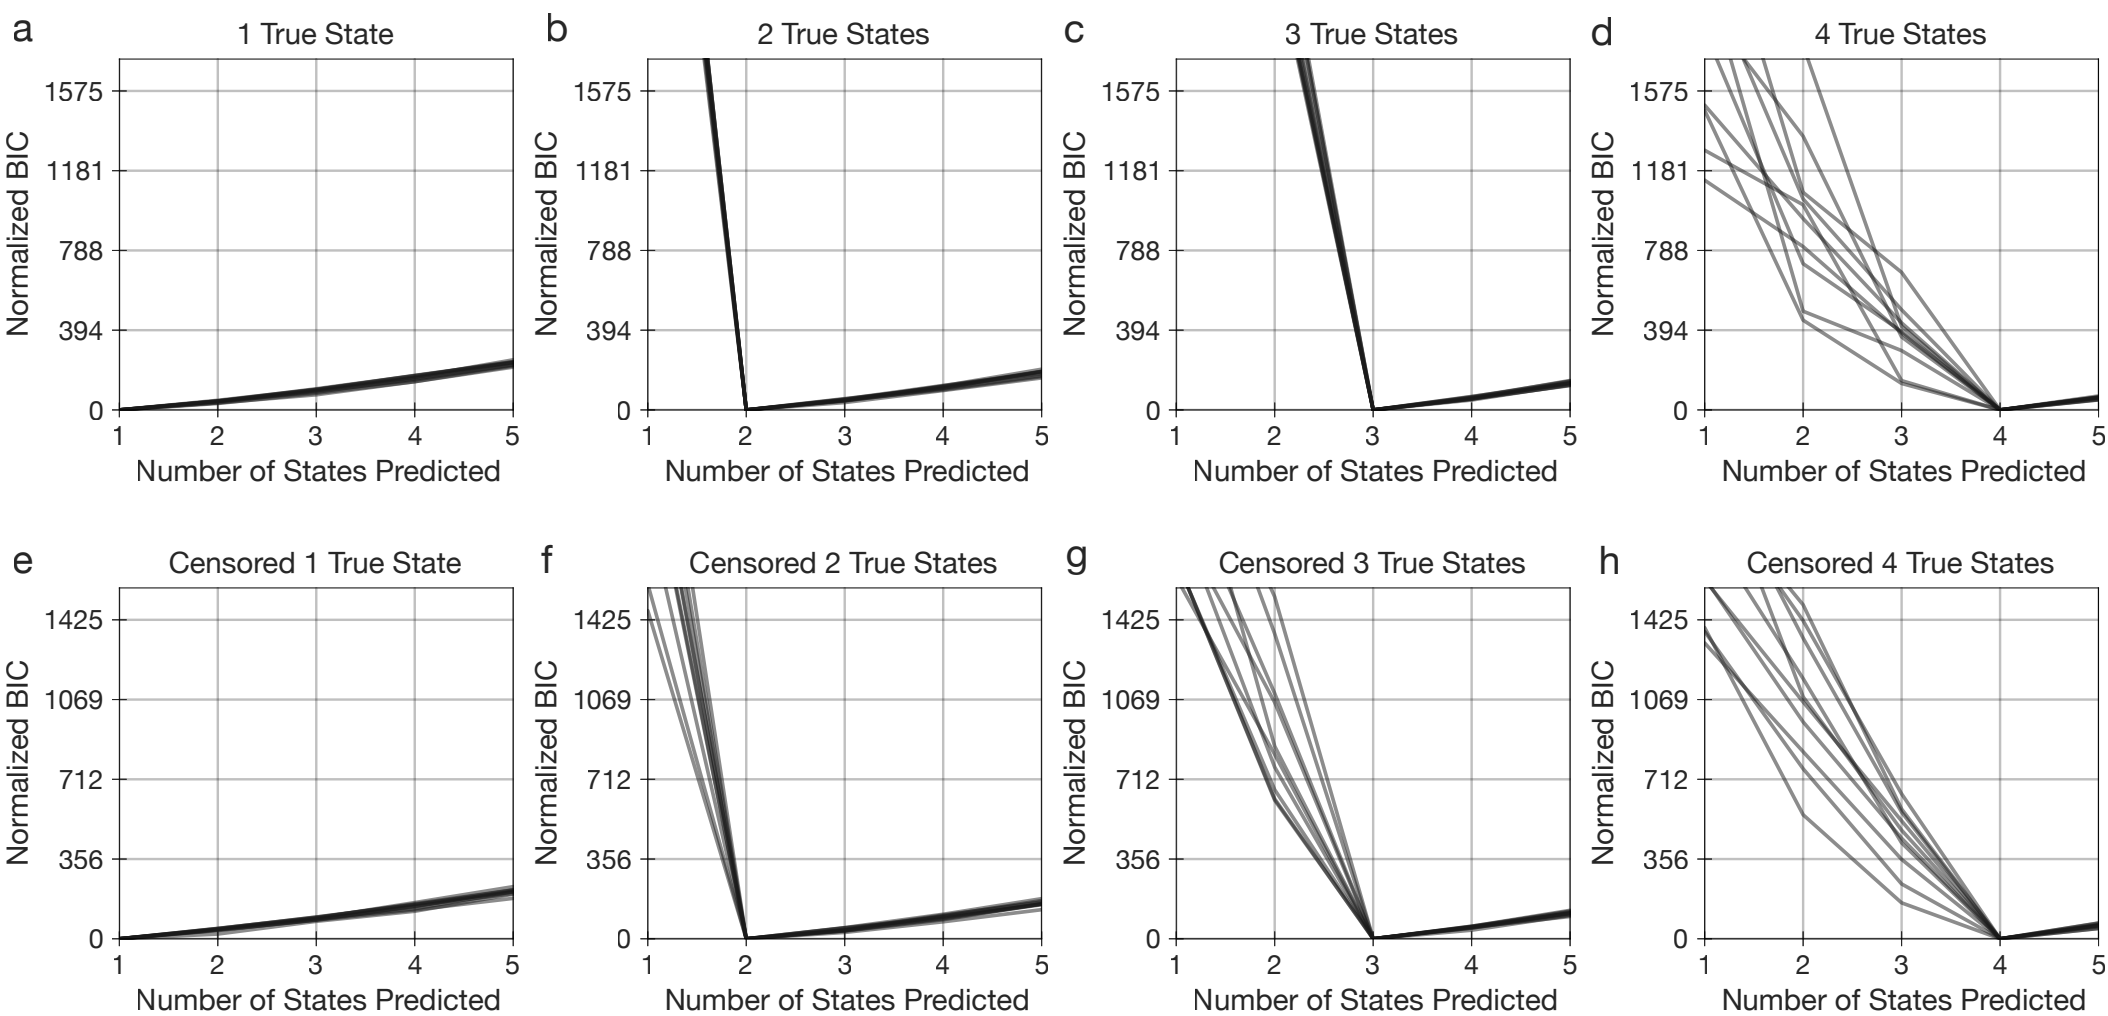
Number of states








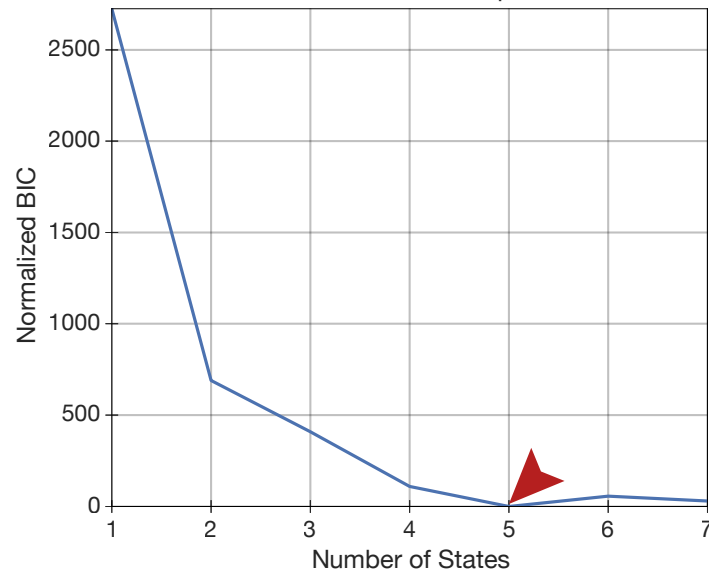
a State difference**b** Distribution similarity**c** State assignment accuracy



Lapatinib-Treated Populations



Gemcitabine-Treated Populations



Growth Factor-Treated Populations

