RESEARCH ARTICLE 7 NOVEMBER 2026

A uniform sampling pool underestimates melanoma lineage overlap between BRAF and MEK inhibitor treatments

New York General Group, Inc.

A melanoma lineage that is abundant after two different treatments may share a response to both drugs, or it may simply have been abundant before the culture was divided. We used published barcode counts to test how far shared sampling opportunity can explain recurrence after separate exposures to dabrafenib and trametinib.[1] Of 682 lineages in the top decile after dabrafenib and 585 after trametinib, 331 appeared in both lists; only 11 appeared in the top decile under all six conditions examined. Treatment-exclusive counts establish that at least 2,208 distinct lineages ranked highly in at least one condition. Any common pool large enough to contain those lineages, with equal hit probabilities among eligible lineages and independent sampling across conditions, predicts at most 180.69 inhibitor-pair hits. Allowing unequal lineage weights does not reconcile the two expected overlaps if those weights affect all six treatments proportionally: matching the inhibitor-pair count then predicts at least 91.29 six-condition hits. By contrast, a shared size advantage whose magnitude differs across treatments can match both counts. The overlap therefore distinguishes between specific sampling explanations without, by itself, identifying a mechanism of drug resistance.

Lineage recurrence across separate treatment arms

Schaff *et al.* divided an expanded population of barcoded WM989 melanoma cells into duplicate cultures for each of six conditions: dabrafenib, trametinib, cobalt chloride, acidic medium, cisplatin and doxorubicin. They ranked lineages within each condition by endpoint genomic DNA barcode abundance averaged across split replicates.[1] We define a *hit* as a lineage in the published top decile for that condition. Its denominator was the number of lineages detected after that treatment, not the 17,593 lineages in the overall comparison. A hit measures relative endpoint abundance; it does not distinguish survival of the original cells from proliferation of their descendants.[1,2]

Table 1 | Published lineage counts by treatment.

The denominator and hit count are from Fig. 2d of Schaff *et al.*; exclusive counts are from Fig. 2g.[1] An exclusive hit ranked in the top decile in that condition and in none of the other five. It need not have been absent from the other cultures.

Condition	Lineages detected at endpoint	Top-decile hits	Hits exclusive to condition
Dabrafenib	6,819	682	258
Trametinib	5,851	585	219
Cobalt chloride	4,848	485	181
Acidic medium	10,230	1,023	193
Cisplatin	9,965	996	179
Doxorubicin	11,785	1,178	239

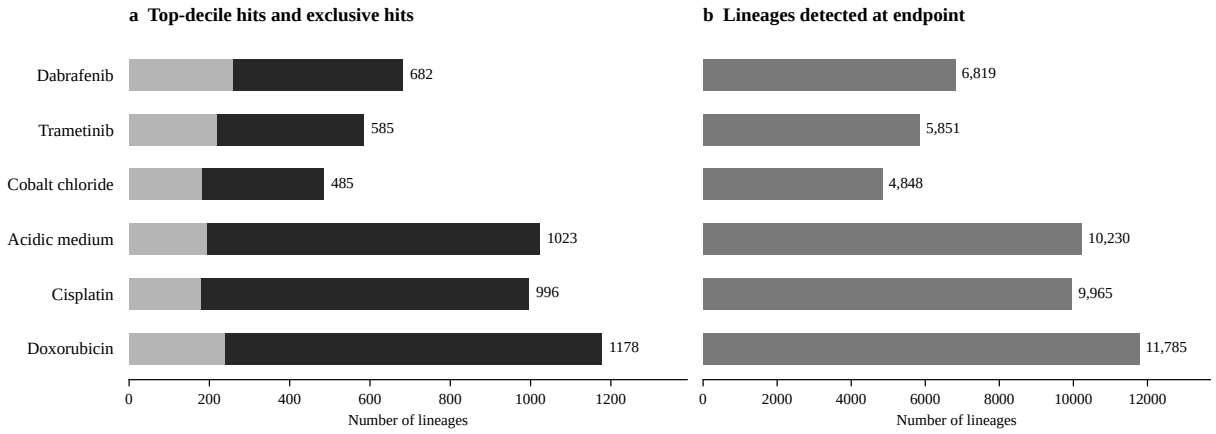


Figure 1 | Published lineage counts by treatment. Dark bars indicate all top-decile hits; light bars overlaid on the dark bars indicate hits exclusive to the condition. The right panel shows endpoint lineage detection counts. Values correspond to Table 1.

The inhibitor-pair overlap was 331 lineages (Fig. 2e of ref. 1). This count includes the 11 lineages that ranked in the top decile in all six conditions (Fig. 2g). The original study established that the overlap between the inhibitors was greater than expected when sampling from all 17,593 lineages.[1] We ask a narrower question: **could a smaller pool, common to all six conditions, account for that overlap?**

Treatment-exclusive hits constrain the size of the pool

Let U be the number of distinct lineages that were hits in at least one condition. Doxorubicin alone contributed 1,178. The lineages exclusive to each of the other five conditions must be distinct from these 1,178 and from one another. Thus,

$$U \;\geq\; 1{,}178 + 258 + 219 + 181 + 193 + 179 \;=\; 2{,}208.$$

This is a lower bound, not an estimate of U. The exact union could in principle be obtained by summing every disjoint region in Fig. 2g of ref. 1. The lower bound requires only one marginal total and five exclusive counts, avoids reliance on the smallest intersection labels and is already sufficient for the test. Any common eligible pool must contain every observed hit, so its size M must satisfy $M \geq U$.

Write n_j for the published hit count in condition j. If each condition independently selects exactly n_j lineages uniformly from the same pool, the expected overlap Y between the inhibitors is

$$\mathbb{E}_M[Y] = \frac{682 \times 585}{M} \leq \frac{398{,}970}{2{,}208} = 180.69.$$

The observed overlap of 331 is therefore **1.83 times the largest expected overlap** compatible with the union bound. At $M = 2{,}208$, the exact hypergeometric probability of at least 331 shared hits is 7.39×10^{-53}. Increasing M can only reduce that probability: it adds non-dabrafenib lineages to the pool from which the 585 trametinib hits are drawn. The six-condition count of 11 is not needed for this result.

The pairwise and six-condition counts imply different pool sizes

The six-condition count provides a second check. Let X denote the number of lineages that are hits in every condition. Under independent sampling across conditions,

$$\mathbb{E}_M[X] = \frac{\prod_{j=1}^{6} n_j}{M^5}, \qquad \mathbb{E}_M[Y] = \frac{n_{\mathrm{D}} n_{\mathrm{T}}}{M},$$

where D and T denote dabrafenib and trametinib. These expectations hold both when list sizes are fixed and when individual hit indicators are independent Bernoulli variables with probabilities n_j/M. The distributions, however, are not the same under the two schemes.

Matching the six-condition expectation to 11 gives $M = 1{,}840.41$; at the nearest integer, only 216.83 inhibitor-pair hits are expected. Matching the inhibitor-pair expectation to 331 instead gives $M = 1{,}205.35$, for which the six-condition expectation is 91.29. Neither fitted pool is large enough to contain the 2,208 distinct lineages established directly from the data.

Table 2 | Expected overlaps and exact six-condition tail probabilities.

Expectations apply to both sampling schemes. The probability column applies only to the independent Bernoulli scheme, under which X is binomial. Values are rounded from the exact distribution, not estimated by simulation. The last two rows fit one overlap in isolation and are incompatible with the observed union bound.

Eligible lineages M	Basis for pool size	Expected inhibitor-pair hits	Expected six-condition hits	$\Pr_M(X \geq 11)$
17,593	Full comparison pool	22.68	0.000138	8.50×10^{-51}
2,208	Lower bound on distinct hits	180.69	4.43	0.005874
1,840	Fitted to six-condition count, rounded	216.83	11.01	0.541926
1,205	Fitted to inhibitor-pair count, rounded	331.10	91.42	$1 - 1.36 \times 10^{-28}$

At $M = 1{,}840$, obtaining at least 11 six-condition hits has probability 0.542. Thus, the six-condition count alone is readily explained by concentrating the opportunity to rank highly in a smaller pool. It is the combination of the inhibitor-pair overlap and the minimum number of distinct hits that rules out this explanation.

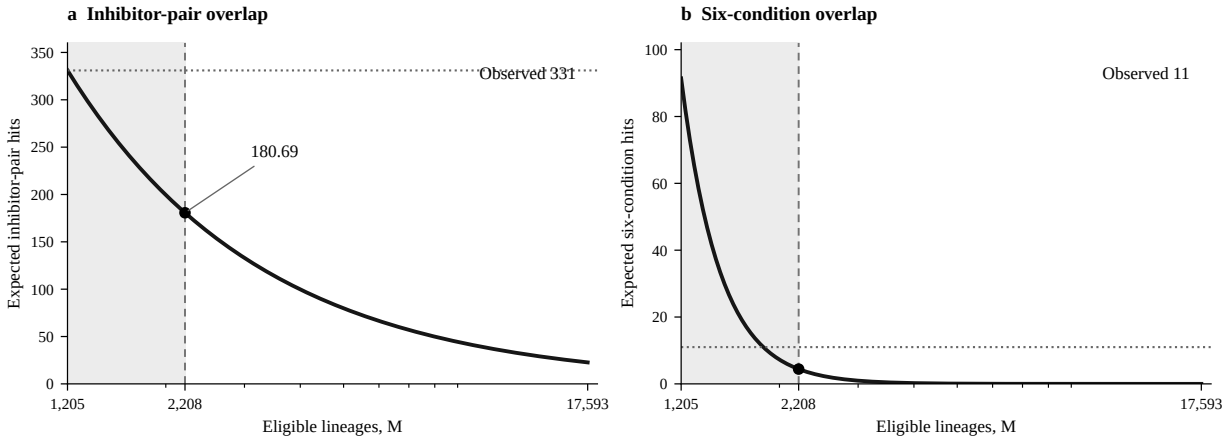


Figure 2 | Expected overlaps under uniform sampling. The horizontal lines mark the observed inhibitor-pair count (331) and six-condition count (11); the shaded region indicates pool sizes below the directly established minimum of 2,208 distinct lineages.

What unequal clone sizes can explain

A common pool treats eligibility as binary. A more biological alternative allows each lineage i to have a different non-negative sampling weight w_i, with $\sum_i w_i = 1$. Suppose first that the same weight acts **proportionally** in every treatment: the hit probability is $p_{ij} = n_j w_i$, with all probabilities at most one. Conditional on the weights, treatments sample independently. Then

$$\mathbb{E}[Y] = n_{\mathrm{D}} n_{\mathrm{T}} \sum_i w_i^2, \qquad \mathbb{E}[X] = \left(\prod_j n_j\right) \sum_i w_i^6.$$

These expectations obey a restriction that does not depend on the shape of the clone-size distribution. Choose a lineage I with probability w_i. Because x^5 is convex for $x \geq 0$,

$$\sum_i w_i^6 = \mathbb{E}[w_I^5] \geq \left(\mathbb{E}[w_I]\right)^5 = \left(\sum_i w_i^2\right)^5.$$

Consequently, calibrating the proportional-weight explanation to $\mathbb{E}[Y] = 331$ forces

$$\mathbb{E}[X] \;\geq\; \left(\prod_j n_j\right) \left(\frac{331}{682 \times 585}\right)^5 = 91.2853.$$

Equality in the moment inequality requires all positive weights to be equal. Thus, concentrating sampling opportunity in a few very large clones cannot lower the six-condition expectation towards 11 while retaining a pairwise expectation of 331. This is a statement about **expected counts** under proportional weights, not a claim that observing 11 is impossible.

Proportionality is stronger than the biological premise that larger clones are more likely to rank highly. To see the difference, consider 1,000 larger lineages among the 17,593. Give each larger lineage hit probabilities 0.682 under dabrafenib, $331/682$ under trametinib and $q = (11/331)^{1/4} \approx 0.427$ under each of the other four conditions. In each condition, distribute the remaining expected hits uniformly across the 16,593 smaller lineages. The smaller lineages then have a lower probability than the larger ones in **every** treatment and zero probability under dabrafenib. With independent sampling across arms,

$$\mathbb{E}[Y] = 1{,}000(0.682)(331/682) = 331, \qquad \mathbb{E}[X] = 331q^4 = 11.$$

This example matches all six marginal counts and both overlaps **in expectation**. It is not a fit to the full intersection diagram or a proposed distribution of clone sizes. It establishes the limit of the inference: a single ranking of lineages by starting size remains compatible with both overlaps if the effect of size on hit probability differs among treatments.

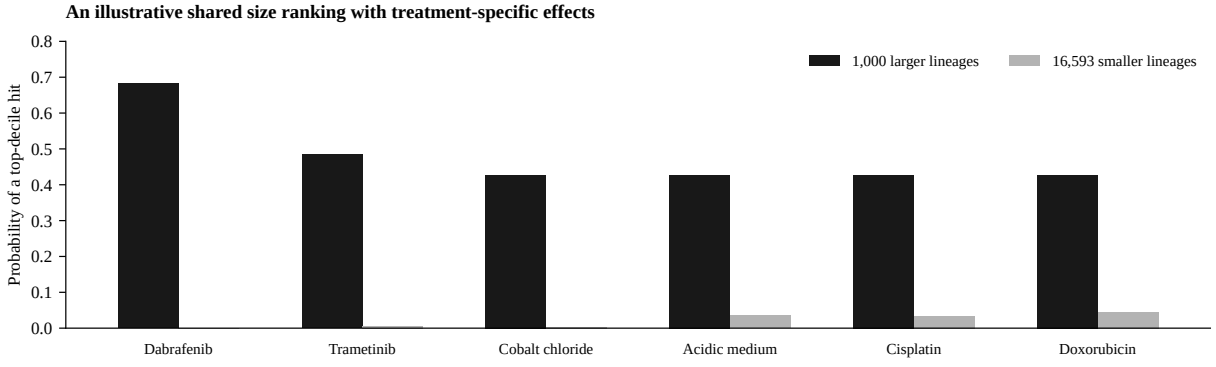


Figure 3 | Treatment-dependent size advantage. A two-class probability construction matches the six treatment hit counts and both overlap counts in expectation while retaining the same ranking of larger and smaller lineages in every treatment.

Discussion

The strongest evidence against a uniform pool comes from lineages that did **not** rank highly under both inhibitors. The treatment-exclusive hits establish how large the pool must be; that minimum size places the observed inhibitor-pair overlap far into the tail of the sampling distribution. The six-condition count answers a different question. It shows why unequal clone sizes do not resolve the discrepancy when every treatment inherits the same proportional sampling weights.

A shared response to dabrafenib and trametinib is biologically plausible because the drugs inhibit different components of the same MAP kinase pathway.[1] Independent measurements in the source study are informative here: in a separate single-cell experiment, Schaff *et al.* reported weak correlations between pretreatment and endpoint clone size for dabrafenib and trametinib ($r = 0.11$ and 0.07), compared with $r = 0.71$ for acidic medium.[1] Earlier lineage-tracing work also showed that pretreatment cell states can predict differences in the proliferative capacity of resistant melanoma clones.[2] Neither observation converts endpoint barcode rank into a direct measurement of cell survival.

The discriminating comparison is therefore between lineages of **similar measured starting abundance**, not simply between recurrent and non-recurrent barcodes. Pretreatment counts from the population used to seed the treatment arms would allow inhibitor-pair recurrence to be assessed within abundance strata. Measuring early survival separately from subsequent regrowth would then distinguish two explanations that the endpoint lists leave open: a lineage may be unusually likely to survive both inhibitors, or its surviving descendants may be unusually likely to expand under both.

Methods

Data extraction. The unit of analysis was a barcode-defined lineage. Condition-specific denominators and hit counts were taken from Fig. 2d, the inhibitor-pair overlap and comparison-pool size from Fig. 2e, and the exclusive and six-condition counts from Fig. 2g of Schaff *et al.*[1] Split cultures contributed to the published within-condition rankings; they were not treated as independent observations of the overlap. No sequencing reads were reprocessed.

Fixed-count calculation. Under independent, uniform selection of lists of sizes 682 and 585 from M eligible lineages, the overlap has probability mass

$$\Pr_M(Y = y) = \frac{\binom{682}{y}\binom{M-682}{585-y}}{\binom{M}{585}}.$$

For $M < 2{,}208$, the event $U \geq 2{,}208$ is impossible. For $M \geq 2{,}208$, the upper tail for Y is largest at $M = 2{,}208$. The reported hypergeometric tail therefore also bounds the probability of observing **both** $U \geq 2{,}208$ and $Y \geq 331$ under the fixed-count scheme. It is not a probability calculated after conditioning on the observed union size.

Independent-hit calculation. With independent hit indicators across lineages and conditions, $X \sim \text{Binomial}(M, p_6)$, where $p_6 = \prod_j (n_j/M)$. The exact upper tail is $\Pr(X \geq 11) = \Pr(X > 10)$. The inhibitor-pair count is also binomial, with mean $\mu = 682 \times 585/M$. For $t > \mu$, applying the exponential probability bound and minimizing over the positive exponent gives

$$\Pr_M(Y \geq t) \leq \exp[t - \mu - t\log(t/\mu)]\,.$$

The same bound holds for fixed-count sampling because the exponential moment of a sum sampled without replacement does not exceed its counterpart with replacement.[3] At $M = 2{,}208$ and $t = 331$, the bound is 1.8281×10^{-22}; it decreases for larger pools. This looser bound applies to both schemes, whereas 7.39×10^{-53} is the exact fixed-count tail at the boundary. All probabilities concern the specified retrospective comparison, not an independently selected endpoint or a causal hypothesis.

Data and code availability. Table 1 and the two published intersection counts contain all biological inputs to this analysis.[1] The following base R code reproduces Table 2, the pairwise probability bounds and the unequal-weight calculations.[4] Both tails are evaluated directly so that very small probabilities are not lost by subtraction from one.

References

1. Schaff, D. L. *et al.* Pre-existing cell states predict resistance to multiple treatments. *Cell Genom.* **6**, 101191 (2026). (pubmed.ncbi.nlm.nih.gov)

2. Emert, B. L. *et al.* Variability within rare cell states enables multiple paths toward drug resistance. *Nat. Biotechnol.* **39**, 865–876 (2021). (nature.com)

3. Hoeffding, W. Probability inequalities for sums of bounded random variables. *J. Am. Stat. Assoc.* **58**, 13–30 (1963). (cs.rpi.edu)

4. R Core Team. *R: A Language and Environment for Statistical Computing* (R Foundation for Statistical Computing, Vienna, 2026). (cloud.r-project.org)

Reproducibility code

```
hits <- c(dabrafenib = 682, trametinib = 585, cobalt = 485,
          acid = 1023, cisplatin = 996, doxorubicin = 1178)
exclusive <- c(dabrafenib = 258, trametinib = 219, cobalt = 181,
               acid = 193, cisplatin = 179, doxorubicin = 239)

N <- 17593
Y <- 331
X <- 11
pair_product <- prod(hits[1:2])

# One full list plus the exclusive hits in each of the other five lists.
L <- unname(hits["doxorubicin"] + sum(exclusive[names(exclusive) != "doxorubicin"]))
stopifnot(L == 2208)

fit_six  <- (prod(hits) / X)^(1/5)
fit_pair <- pair_product / Y

pool_summary <- function(M) {
  stopifnot(M >= max(hits))
  p6 <- prod(hits / M)
  data.frame(
    pool = M,
    expected_pair = pair_product / M,
    expected_six = M * p6,
    prob_X_ge_11 = pbinom(X - 1, M, p6, lower.tail = FALSE),
    prob_X_le_10 = pbinom(X - 1, M, p6, lower.tail = TRUE)
  )
}

pools <- c(N, L, round(fit_six), round(fit_pair))
print(do.call(rbind, lapply(pools, pool_summary)), digits = 8,
      row.names = FALSE)

# Exact fixed-count tail and a bound valid under both sampling schemes.
mu <- pair_product / L
fixed_tail <- phyper(Y - 1, hits[1], L - hits[1], hits[2],
                     lower.tail = FALSE)
common_bound <- exp(Y - mu - Y * log(Y / mu))
weighted_min <- prod(hits) * (Y / pair_product)^5

print(c(fit_six = fit_six, fit_pair = fit_pair,
        fixed_count_tail = unname(fixed_tail),
        exponential_bound = common_bound,
        proportional_weight_lower_bound = weighted_min), digits = 9)

# A shared size ranking with different size effects across treatments.
n_large <- 1000
n_small <- N - n_large
p_large <- c(hits[1] / n_large, Y / hits[1], rep((X / Y)^(1/4), 4))
p_small <- (unname(hits) - n_large * unname(p_large)) / n_small
p_small[1] <- 0  # Exactly zero by construction.

expected_hits <- n_large * p_large + n_small * p_small
expected_pair <- n_large * prod(p_large[1:2]) + n_small * prod(p_small[1:2])
expected_six  <- n_large * prod(p_large) + n_small * prod(p_small)

stopifnot(all(p_small >= 0), all(p_small < p_large), all(p_large <= 1),
          max(abs(expected_hits - hits)) < 1e-9,
          abs(expected_pair - Y) < 1e-9,
          abs(expected_six - X) < 1e-9)
```