Guide selection transfers depletion effects across noncanonical ORF screens

Robin Gounder[1], Samuel Schmidt[2,3]

[1]University of Cambridge

[2]Vaionex Corporation

[3]Cornell University

Correspondence: robin@gounder.com

Samuel Schmidt: samuel@vaionex.com

29 October 2026

Abstract

CRISPR fitness screens help prioritize noncanonical open reading frames (nORFs) for biological follow-up, but cross-study agreement can reflect both product dependencies and shared guide responses. We reconstructed complete reference ORFs and matched exact guide sequences across published human screens to quantify how earlier guide selection changes later depletion estimates. Among 105 matched products with 471 shared guides in three cell lines, relative guide responses remained correlated after subtracting each product-by-cell mean (equal-product weighted $r = 0.643$; nominal 95% dependence-cluster interval 0.576-0.696). In a retrospective 87-product scenario, selecting the two most-depleting earlier guides shifted the later score by -0.202 $\log_2$ units relative to the expected score under uniform random two-guide selection (nominal 95% interval -0.246 to -0.160). A post hoc sensitivity excluding the evaluated cell from the earlier selection score retained a shift. An analytical and synthetic positive control shows that shared on-target guide efficacy acting on a true dependency can generate both centred correspondence and selection shifts. The count-balanced difference in product-level agreement between reused and newly designed guides remained uncertain (0.086; 97.5% interval -0.094 to 0.268). A third publication provided product-level rank agreement across 196 exact products (Spearman $\rho = 0.462$), with guide reuse unresolved. These results quantify a consequence of guide-selection history in matched nORF screens; attributing depletion to a peptide, RNA or local regulatory mechanism requires complementary evidence.

Keywords: functional genomics; microproteins; CRISPR; replication; guide RNA

1 Introduction

Noncanonical open reading frames (nORFs) are reading frames not represented as conventional protein-coding products in reference annotations, including products within untranslated regions, other reading frames and transcripts annotated as noncoding (Mudge et al., 2022). Noncanonical describes their relationship to annotation and does not imply that a product was newly discovered in the present study. Ribosome profiling first enabled genome-wide measurement of translation in yeast (Ingolia et al., 2009), and mammalian applications revealed unannotated products and alternative initiation at both AUG and non-AUG codons (Ingolia et al., 2011; Lee et al., 2012). Human peptidomics detected products of short ORFs outside conventional annotations (Slavoff et al., 2013). Proteogenomic profiling also detected regulated peptides from RNAs classified as noncoding in mouse neurons (Prabakaran et al., 2014), while ribosome-profiling analyses extended translated regions to lncRNAs, untranslated regions and pseudogenes (Ji et al., 2015). Tissue studies subsequently identified previously unrecognized microproteins in the human heart and other organs (van Heesch et al. 2019). These studies establish a substantial discovery space, but the measurements supporting translation do not all address the same question as a functional perturbation. Here we examine how guide identity and earlier guide selection contribute to agreement between nORF fitness screens.

Defining the product being studied is itself an analytical task. ORF-RATER uses a regression-based framework to identify translated coding sequences (Fields et al., 2015); RiboTaper uses ribosome-footprint periodicity to identify translated regions (Calviello et al., 2016); and PRICE models experimental noise to resolve overlapping ORFs and noncanonical initiation (Erhard et al., 2018). Community annotation guidance and expanded reference catalogues provide a shared basis for representing these products (Mudge et al., 2022; Chothani et al., 2026). Current proteome annotation also distinguishes detected translation products from evidence establishing their physiological function (Deutsch et al., 2026). When functional studies are combined, shared gene names, starts or genomic regions therefore need to be distinguished from complete reference-product identity.

Mechanistic studies demonstrate that some previously unannotated products have specific biological activities. In mouse muscle, myoregulin inhibits SERCA-mediated calcium uptake, whereas DWORF enhances SERCA activity by opposing inhibitory peptides (Anderson et al., 2015; Nelson et al., 2016). These findings combine a defined product with molecular and physiological measurements. They illustrate the value of nORF research while also showing why a pooled fitness phenotype and a product-specific mechanism represent different steps in an evidence chain.

Pooled CRISPR fitness screens target DNA regions and track changes in guide abundance as the cells carrying those guides expand or become depleted. The endpoint is a relative cell-fitness proxy, rather than a direct measurement of an nORF peptide or its host canonical protein. Functional screening has expanded the number of products that can be investigated. Chen et al. (2020) combined systematic CRISPR perturbation with translation, localization and interaction measurements, while Prensner et al. (2021) coupled cancer-cell fitness screening with selected molecular follow-up. Disease-focused work identified ASNSD1-uORF/ASDURF as a medulloblastoma-associated dependency (Hofman et al., 2024). Larger subsequent studies examine additional candidate microproteins and integrate perturbation outcomes with other annotation evidence (Schlesinger et al., 2025; Deutsch et al., 2026). Cross-study comparisons can help prioritize follow-up, provided they specify both the product and the intervention whose

phenotype is being reproduced.

A phenotype after targeting an nORF region can involve local DNA or RNA changes, regulation associated with translation, or activity of the encoded peptide. For an upstream ORF, translation can regulate expression of a downstream canonical coding sequence without establishing an autonomous function of the upstream peptide (Calvo et al., 2009). A host-gene label alone does not resolve these alternatives or identify the transcript involved. The present reanalysis measures screen correspondence; it does not directly measure changes in the host canonical protein. Selected source-study experiments address mechanisms beyond the screen outcomes reanalysed here.

Differences among guides are established features of CRISPR screening. Library evaluation measures reagent performance (Hart et al., 2017), and JACKS jointly models guide efficacy across screens using the same library (Allen et al., 2019). Recent library-design work explicitly balances on-target activity with off-target avoidance (Drepanos et al., 2026). Reproducibility statistics must also isolate the signal relevant to the biological question (Billmann et al., 2023); common-guide correspondence has been examined directly in noncoding regulatory screens (Yao et al., 2024). Removing a product mean does not remove differences in guide efficacy acting on a true product dependency. The empirical question addressed here concerns the magnitude and consequence of persistent relative guide responses within precisely matched nORF products. In particular, how much correspondence remains after removing a product's mean response, and how does selecting guides by earlier activity affect its later depletion estimate?

We address this question across the Prensner et al. (2021) and Deutsch et al. (2026) screens by matching complete reference products and exact guide sequences. Conditional guide-response analysis separates correspondence within a product from agreement in its average phenotype. An earlier-screen-only two-guide scenario then quantifies the resulting change in a later score, with a post hoc sensitivity excluding the evaluated cell from selection. A count-balanced reused/fresh comparison and a third publication provide complementary product-level evidence (Schlesinger et al., 2025). Together, these analyses establish a quantitative basis for interpreting and reporting repeated nORF perturbations.

2 Materials and methods

Data sources, analytical scope and timing

We used original released guide-level data associated with Prensner et al. (2021) and the current-screen data in Supplementary Table 13 of Deutsch et al. (2026). The primary comparison uses the 553-candidate older inventory and 2,196 newer Phase-1 catalogue nORFs. The additional follow-up comparison uses Prensner et al. Supplementary Tables 24 and 27 and the separately released secondary-screen count matrix. The latter contains 6,996 sequenced constructs and includes regions and controls outside our matched-product cohort. These original data remain distinct from the newer publication's reprocessed meta-analysis. Third-study aggregate results and identity metadata come from the Schlesinger et al. author-repository table `gene_summaries_final_v4.csv` (Schlesinger et al., 2025), pinned to commit 6ca8f3961e8a94fb116984b9ac6e1f4f08678006 and checked against its Git blob hash. The table has 11,957 rows before identity filtering. The separately listed individual-guide results and library design supplements could not be retrieved, and no guide sequence independence is inferred.

The analyses were developed retrospectively in stages. Underlying source outcomes and product averages were known before the conditional shared-guide question and the count-balanced sensitivity were specified. The additional follow-up and transferred guide-selection rules were recorded before their new outcome comparisons; all main-screen outcomes were already known. These records document analysis choices rather than prospective preregistration. Intervals are nominal and conditional on the selected question, cohort and fixed cell contexts. The extensions are exploratory; their intervals are not adjusted across all analyses in this retrospective study. The third-study protocol was recorded after metadata-only matching and before opening its numeric fitness columns. Results from all earlier analyses were already known.

Complete reference products and relationship classes

Older source ORF definitions and newer Phase-1 BED coding blocks were reconstructed against the pinned hg38 reference. Coding blocks are zero-based, half-open intervals in RNA orientation. Source stop-codon conventions, non-AUG initiation and literal standard-code translation were retained rather than silently replacing source products. Complete coding-chain/DNA identity is the first relationship class, followed by complete amino-acid identity, same-start nonidentical products, different-start same-phase coding-body overlap, other-frame overlap and same-gene relationships without body overlap. Gene labels do not resolve co-transcription or isoform identity. The original DNA-to-printed-protein discrepancy flags remain available alongside literal reconstructed products.

Guide qualification and outcome construction

Exact 20-nucleotide strings identify shared reagents. Geometry is inherited from the complete frozen source-guide scans: one primary-genome NGG site, a cut inside the assigned nORF, one source target assignment and no intersection between the whole protospacer and any current annotated protein-coding CDS. Qualification does not use the direction or magnitude of the fitness response. Finite outcomes in all three common cells are required. The newer screen additionally requires plasmid abundance of at least 10 in each of its three source samples. Shared guide coordinates must agree between inventories. Current CDS and measured CDS-bearing RNA splice geometry use the pinned GENCODE v50 annotation (GENCODE, 2026).

A unique primary-genome NGG site denotes uniqueness of the exact guide match under this geometry rule. It does not exclude mismatch-tolerant off-target sites, copy-number-dependent cutting effects or perturbation of unannotated regulatory elements. Avoiding annotated canonical CDS limits direct overlap with those coding regions; it does not establish peptide-specific disruption. No new editing-efficiency, off-target or copy-number correction is fitted. These filters define the comparison cohort and should not be interpreted as a molecular specificity assay.

The older endpoint is the source-provided day-21 LFC relative to plasmid DNA. The newer endpoint is the mean of $\log_2(A+1)$ across three endpoint samples minus the mean of $\log_2(A+1)$ across three plasmid samples, where A denotes released processed abundance. These are released processed abundance measurements, not raw integer read counts. They are not a joint-normalized or copy-number-corrected Chronos endpoint. Direct source joins and reconstruction from all required newer count rows reproduce the 1,413 conditional-analysis guide-cell outcomes.

Conditional shared-guide estimand and uncertainty

Products enter the main conditional cohort if at least three guide sequences qualify in both screens. For screen s, product p, guide g and cell c, let G_p be the product's eligible shared-guide set and $m_p = |G_p|$. The residual and guide-cell weight are defined in Equation (1). The summation index denotes the guides being averaged; it is distinct from the guide whose residual is evaluated. The reported Pearson correlation uses these weights, giving each product and each cell the same total weight. Product-average response and any additive product-by-cell term have been removed separately in both screens.

$$\begin{aligned} e_{s,p,g,c} &= \mathrm{LFC}_{s,p,g,c} - \frac{1}{m_p} \sum_{g' \in G_p} \mathrm{LFC}_{s,p,g',c} \\ w_{p,g,c} &= \frac{1}{3m_p} \end{aligned} \tag{1}$$

Dependence groups join source-gene and exonic-locus relationships from the complete frozen inventory. All products, guides and cells in a group are resampled together. The 105 main products occupy 105 such groups. We generated 2,000 cluster samples with replacement using NumPy `default_rng` seed 20261015 and report the 2.5th and 97.5th percentiles. These are nominal conditional intervals over the observed product pool and fixed contexts, not estimates of variability across independently repeated laboratories or patient populations.

A separate permutation check uses `default_rng` seed 20261016 and 5,000 random guide-identity permutations within each product. The same permutation is used across all three cells, preserving each guide's cross-cell profile. Equation (2) defines the two-sided add-one permutation P value: B is the number of permutations, the observed correlation is compared with each permuted correlation in absolute value, and I(condition) equals one when its condition holds and zero otherwise. Validity of this pairing check depends on the stated within-product exchangeability model; the biological mechanism is not identified by rejection.

$$P_{\mathrm{perm}} = \frac{1 + \sum_{b=1}^{B} \mathrm{I}\left(\left|r^{(b)}\right| \geq |r_{\mathrm{obs}}|\right)}{B+1}, \qquad B = 5000 \tag{2}$$

Sensitivity analyses and descriptive covariance decomposition

Declared point sensitivities include each cell separately, at least four shared guides, at least three guides jointly more than 50 nucleotides from the measured CDS-bearing RNA splice boundaries in both inventories, centred within-product/cell guide ranks, and descriptive GC-by-cell adjustment. In the GC analysis, each screen's centred response is separately regressed on product-centred guide GC fraction within each cell using the same weights, then the adjusted residuals are correlated. All 105 single-product omission estimates are supplied. The sensitivity points have no additional multiplicity-adjusted discovery interpretation.

For the covariance decomposition, each screen's cell mean is removed with equal product weights. Its cell-centred guide response is the sum of a product-by-cell component and the already defined within-product residual. Weighted cross terms are zero because each within-product/cell residual sums to zero. Total old-new covariance therefore equals the sum of between-product and within-product covariance. We report their numerical values and the within/total ratio descriptively; this is not a causal variance partition or an estimate of technical error.

Shared-efficacy positive control

To clarify what product-by-cell centring can establish, we derived its behaviour under a shared-guide-efficacy model motivated by JACKS (Allen et al., 2019). We then generated a fixed synthetic positive control, alongside equal-efficacy and within-product later-label-permutation controls. The protocol was recorded before this control was computed, after all empirical manuscript outcomes were known. No parameters were fitted to the empirical data. Complete parameters, outputs and interpretation limits are supplied in Supplementary Note 5, Figure S3 and Tables S33-S34. A separate endpoint implementation checks every synthetic condition and random-pair expectation.

Additional follow-up assay

The follow-up cohort inherits the 127 exact complete-product matches and newer guide eligibility. Shared 20-nucleotide sequences must occur in both the source secondary library and count matrix. Secondary Table 27 must label the guide `ORF_exon`, and its Gene Symbol must equal the older source name after removing only a terminal `_exon` followed by digits; aliases and version-like gene-name differences are not reconciled. We require one secondary-library assignment, plasmid abundance of at least 10, at least three eligible shared guides per product and finite outcomes in A375 and HEPG2. These metadata rules leave 77 guides before the product minimum and 72 guides in 16 eligible products. Eligibility and exclusions are supplied for all 720 newer guide-product rows and all 127 products.

The primary follow-up endpoint is the original Table 27 day-21 LFC relative to plasmid DNA. It is reproduced as the mean of three final-sample $\log_2(\mathrm{CPM} + 1)$ values minus plasmid $\log_2(\mathrm{CPM} + 1)$, using the full released library to calculate sample depths. Newer outcomes are unchanged. Product-by-cell centring and weights $1/(2m_p)$, where m_p here counts eligible shared guides for that product in the follow-up cohort, define the conditional correlation. We resample the 16 inherited dependence groups 2,000 times with `default_rng` seed 20261016 and use 2.5th and 97.5th percentiles. The 5,000 identity permutations use seed 20261017 and move guide profiles jointly across both cells. Declared point sensitivities include at least four guides, adequate counts in all six early samples, early-to-final source LFC, direct count reconstruction, individual cells and product omissions.

Transferred guide-selection scenario

The selection scenario uses the 87 main-cohort products with at least four eligible shared guides. We average older LFC equally over the three common cells, select the two lowest-scoring guides per product and resolve exact ties by lexicographic guide sequence. We then average their newer LFC across cells and subtract the newer all-eligible-guide mean. Averaging this contrast equally across products is the primary selection endpoint. For each product-cell, enumerating every possible two-guide subset confirms that the uniform random-pair expectation equals the all-guide mean. This is a comparison of equal-count strategies in expectation, not a comparison of two-guide precision with all-guide precision.

We resample the 87 inherited dependence groups 2,000 times with seed 20261018, retaining cells jointly, and report a nominal 95% percentile interval. Per-cell means and every product omission are descriptive. For rank sensitivity, the fixed priority budget is the most-depleted 20% of products rounded upward (18 products), ordered by the equal-cell newer mean; exact ties

use the pair identifier. The older-selected and all-guide rankings are compared descriptively. A further 1,000 uniform random-pair plans use seed 20261019 and retain the same guide choice across cells. Their spread describes reagent-selection sensitivity rather than biological sampling uncertainty.

The context-exclusion sensitivity was specified on 26 October 2026 after the original selection result and underlying outcomes were known. For each product and evaluated cell, the guide-selection score is the older LFC averaged over the other two cells. The two lowest-scoring guides, with the same lexicographic tie rule, are evaluated only in the excluded cell of the later screen. The 87 products, 417 eligible guides and random-pair expectation are unchanged. Guide choices may differ across evaluated cells. We report all three cell means and their equal-cell mean, resampling the 87 product dependence groups 2,000 times with `default_rng` seed 20261026 and retaining cells jointly. Nominal 95% percentile intervals describe this post hoc sensitivity. The evaluated cell is absent only from the older selection score: the contexts were already observed, so this is not validation in a biologically unseen cell line. Median product shifts and counts by shift direction are descriptive (Tables S31-S32).

Paired reused/fresh comparisons

The 88-product paired cohort requires an eligible older score and at least one eligible newer reused and fresh guide in every common cell. Within each cell, we calculate Spearman correlations across eligible products between the older medians and, separately, the newer reused-guide means and fresh-guide means. The rank endpoint is the reused-minus-fresh correlation difference averaged equally across cells. The LFC endpoint is the equal-product, equal-cell mean of fresh minus reused newer LFC. The 2,000 target-cluster samples produce 97.5% per-endpoint percentile intervals, using 1.25th and 98.75th percentiles for the originally specified two-endpoint family. These nominal bootstrap intervals do not claim exact finite-sample family coverage.

The explicitly post hoc count-balanced analysis selects the smaller group size from each newer group without replacement, retaining complete guide profiles across cells. All 1,000 fixed guide selections are averaged inside every one of the 2,000 target-cluster samples. This combines guide-selection averaging with target-sampling uncertainty. The guide-selection spread alone is not a biological confidence interval. Complete older-mean, geometry, expression, source-identity, pseudocount and cell-omission sensitivity estimates remain supplied.

Third-study product-level comparison

The third-study identity pass used only complete amino-acid sequences, genomic blocks, guide identifiers/counts and CCDS-overlap metadata. Whitespace and a terminal stop marker were normalized; internal sequence differences were retained. Exact genomic blocks were compared with the newer reference coding body, allowing only whether the terminal stop is included. The primary cohort requires one-to-one identities, source CCDS-overlap FALSE, at least four third-study designed guides and at least four newer guides meeting the existing geometry, source-assignment and plasmid-coverage rules. Products are joined into dependence groups if they share a stable source gene, overlap genomically or share third-study guide identifiers.

For untreated A375, we average the reported `A375_r1_1_0_lfc` and `A375_r2_1_0_lfc` product summaries without skipping missing values. We compare this score with the mean newer LFC across qualified guides. No eligible outcome was missing. The primary endpoint is equal-

product Spearman correlation; 2,000 cluster samples with replacement use `default_rng` seed 20261021 and retain complete clusters. Ranks are recomputed on explicitly duplicated rows in every sample; the interval is a nominal 95% percentile interval. Declared descriptive sensitivities use each third-study replicate, a median newer score, at least four splice-safe newer guides, one lexicographically first product per cluster, the original 127 exact-product inventory and all cluster omissions. Ranking overlap sets $k = \lceil 0.2n \rceil$, where n is the number of eligible products, breaks score ties by source ORF identifier and reports k^2/n as the uniform independent-selection expectation. It is not a significance test or clinical target qualification.

Computational verification

Separate implementations reproduce source joins, complete reference products, guide geometry and each reported endpoint. They replay all 2,000 cluster samples for each new extension, all 5,000 follow-up permutations and all 1,000 random-pair plans. The follow-up endpoint is also reconstructed from original counts. For the main conditional analysis, all 2,000 bootstrap values and a fixed 100-draw permutation replay were checked. These checks establish implementation consistency, not external expert or biological validation. For the third-study comparison, a separate implementation reconstructs all 1,147 newer guide outcomes from processed counts, checks all 196 source product summaries directly from the retrieved CSV, and reproduces the primary rank statistic and all 2,000 cluster bootstrap values using an independent average-rank calculation.

The context-exclusion analysis uses only frozen Table S25 and relative file paths. A separate implementation reconstructs every selected pair using Python loops, checks all 522 original/sensitivity product-cell contrasts and all 16,000 bootstrap values, and verifies summary means, medians and direction counts. Enumerating every possible two-guide subset confirms the random-pair expectation for each product and cell. Analyses use Python 3.11.5, NumPy 1.26.4, pandas 2.0.3 and SciPy 1.11.1; plotting uses Matplotlib 3.7.2. Exact environment versions, seeds, input hashes and complete new outcomes accompany the code.

3 Results

Complete-product matching defines distinct comparison cohorts

We first established the reference products and guide sequences represented in each source. The Prensner et al. inventory contains 553 candidates, of which 551 have resolved reference coding models; the two unresolved candidates remain in the inventory. The Deutsch et al. analytical universe contains 2,196 Phase-1 catalogue nORFs represented in its screen, rather than every target in the larger library. Literal source-gene identifiers or same-strand coding-body overlap defined 274 candidate relationships involving 236 older and 268 newer targets. These are relationships constructed for this analysis, not 274 full-product replication claims made by the source authors.

Only 127 relationships match both the complete reference coding chain and DNA. Two further relationships encode identical complete amino-acid sequences despite different coding chains or DNA. Other relationships have different starts, frames or complete products, or share a source gene without coding-body overlap (Tables S01-S03, S19-S20). Figure 1 distinguishes the reference correspondence from the outcome cohorts it supports.

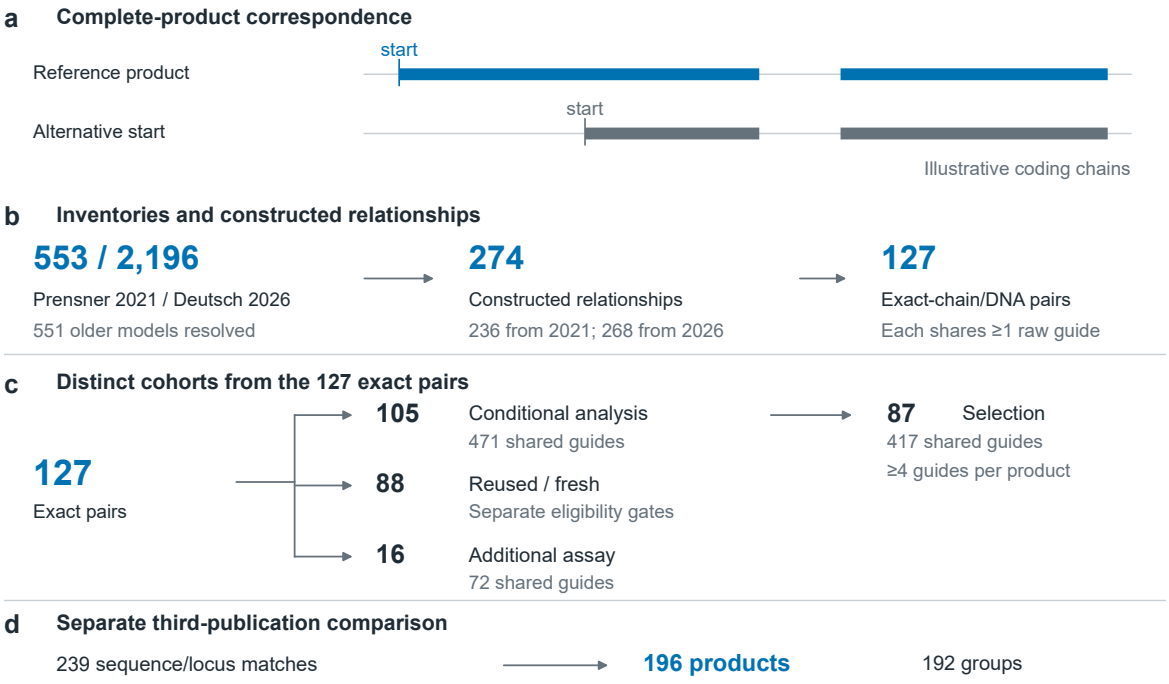

Figure 1: Product correspondence and analytical cohorts. (a) Illustrative coding chains distinguish products with different starts. (b) The Prensner 2021 and Deutsch 2026 analytical inventories yield 274 constructed relationships and 127 exact coding-chain/DNA pairs; every exact pair shares at least one raw guide. (c) Separate eligibility gates define the conditional, reused/fresh and additional-assay cohorts. The 87-product selection cohort is a subset of the 105-product conditional cohort. (d) The third-publication comparison is formed separately. Counts refer to the labelled products or relationships; these branches are not a single attrition sequence.

Across the two analytical inventories, 744 exact 20-nucleotide guide strings are shared, and every one of the 127 exact-chain pairs shares at least one raw guide before eligibility filtering (Table S04). The newer source explicitly incorporated targets from an earlier guide library (Deutsch et al., 2026). Shared reagents are therefore a property to quantify when interpreting agreement. After guide qualification, separate gates define 105 products for conditional shared-guide analysis, 88 for reused/fresh comparison, 16 for the additional-assay comparison and 87 for the selection scenario. The third-publication comparison is formed separately (Figure 1; Table 1).

Table 1: Analytical questions and cohort definitions.

Comparison	Source studies	Fixed cell contexts	Products / sampling units	Guide requirement
Conditional guide responses	Prensner primary / Deutsch	A375, A549, HEPG2	105 / 105	At least 3 shared eligible guides; 471 total
Additional assay	Prensner additional / Deutsch	A375, HEPG2	16 / 16	At least 3 shared eligible guides; 72 total
Transferred selection	Prensner primary / Deutsch	A375, A549, HEPG2	87 / 87	At least 4 shared eligible guides; select 2
Reused/fresh products	Prensner primary / Deutsch	A375, A549, HEPG2	88 / 88 targets	At least 1 eligible newer guide in each group
Third-publication products	Schlesinger / Deutsch	A375	196 / 192	At least 4 designed / qualified guides

The cohorts are defined by different metadata and measurement gates. Product and cell observations move together under their stated resampling scheme. The additional assay belongs to the earlier source study; the third-publication comparison does not establish guide-sequence independence. Full definitions and all exclusions are provided in the Methods and Supplementary Data.

Relative guide responses persist within complete-product matches

We asked whether guide correspondence remains after removing each matched product's average phenotype in each cell. Guides had to qualify in both sources, with one primary-genome NGG site, a cut inside the assigned nORF, one source target assignment and a full protospacer avoiding current annotated canonical CDS. Newer plasmid coverage had to be at least 10 in each of three source samples. These gates leave 471 shared guides in 105 products with at least three eligible shared sequences, measured in A375, A549 and HEPG2 (Tables S11-S12).

For each product and cell, we subtracted the mean eligible shared-guide LFC separately in each screen. Equal total product and cell weights prevent products with more guides from dominating. The 1,413 dependent guide-cell pairs retain weighted Pearson $r = 0.64324$, with nominal 95% dependence-cluster interval 0.57584-0.69649 (Figure 2b). Thus a common product-average response cannot by itself account for the observed correspondence. None of 5,000 within-product guide-identity permutations was as extreme as observed (add-one two-sided $P = 0.000200$). This correspondence test does not identify its biological or technical cause. Shared on-target efficacy acting on a true dependency remains one compatible explanation, as demonstrated by the positive control in Supplementary Note 5; centring does not establish an effect beyond that explanation.

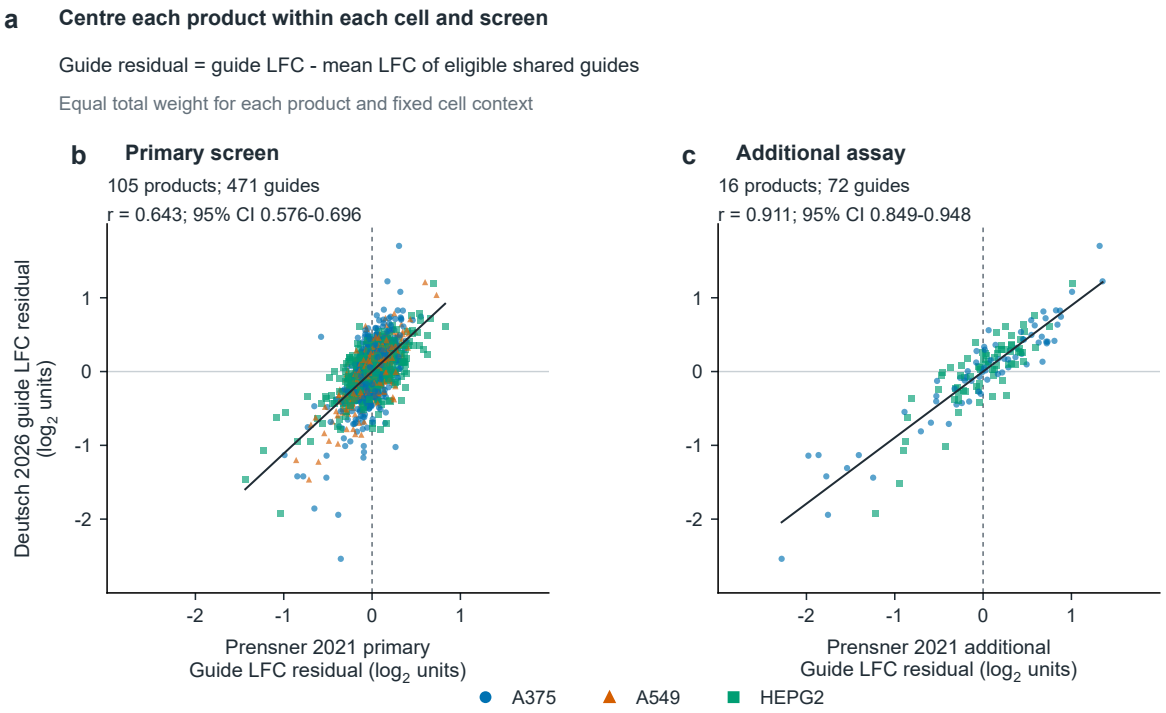

Figure 2: Persistent within-product guide responses. (a) Each guide response is centred on the mean of eligible shared guides for the same product and cell, separately in each screen. (b) The primary comparison contains 105 products, 471 guides and 1,413 dependent guide-cell pairs. (c) The additional assay from the earlier study contains 16 products, 72 guides and 144 pairs. Both scatterplots use identical axis limits. Colours and marker shapes identify cells; each point is a guide-cell pair. Lines are weighted least-squares fits through the origin. Reported Pearson correlations use equal total product/cell weight; nominal 95% intervals resample dependence groups. The additional assay is not a third independent publication. Neither correlation identifies a molecular mechanism.

The estimate varies from 0.62686 to 0.65246 across all single-product omissions. Requiring at least four shared guides gives $r = 0.65115$, and restricting distance from measured CDS-bearing RNA splice boundaries gives $r = 0.64616$. GC adjustment and within-product guide ranks retain positive correspondence (Figure S1; Tables S13-S14). These are sensitivity analyses, rather than independent validations. A descriptive covariance decomposition is supplied in Supplementary Note 3 and Table S15.

The separately performed additional assay from Prensner et al. supports the same pattern in a smaller selected cohort. Sixteen products with 72 eligible shared guides provide 144 guide-cell pairs in A375 and HEPG2. After the same centring and weighting, $r = 0.91124$, with nominal 95% dependence-cluster interval 0.84865-0.94802 (Figure 2c). None of 5,000 guide-identity permutations was as extreme as observed. This assay belongs to the earlier publication and enriches for previously investigated targets; it is not a third independent study (Tables S21-S24).

Selecting guides in the earlier screen shifts the later depletion score

Persistent relative guide responses suggest a practical consequence for follow-up design. We examined a retrospective scenario in the 87 main-cohort products with at least four eligible shared guides, comprising 417 guides. For each product, we ranked guides using only their equal-cell mean Prensner primary-screen LFC and selected the two most-depleting guides. The

same choice was retained across the three Deutsch cell contexts (Figure 3a; Table S25). This rule was not established to have been used by either source study, and later outcomes were already known during the wider research program.

The reference is the exact expected later-screen score under uniform selection of two eligible guides without replacement, which equals that product's all-eligible-guide mean. The comparison therefore holds guide count at two in expectation; it does not equate the precision of a realized two-guide measurement with the all-guide mean. The equal-product selected-minus-reference shift in the later screen is −0.20180 $\log_2$ units, with nominal 95% dependence-cluster interval −0.24565 to −0.16004 (Figure 3b; Tables S26-S27).

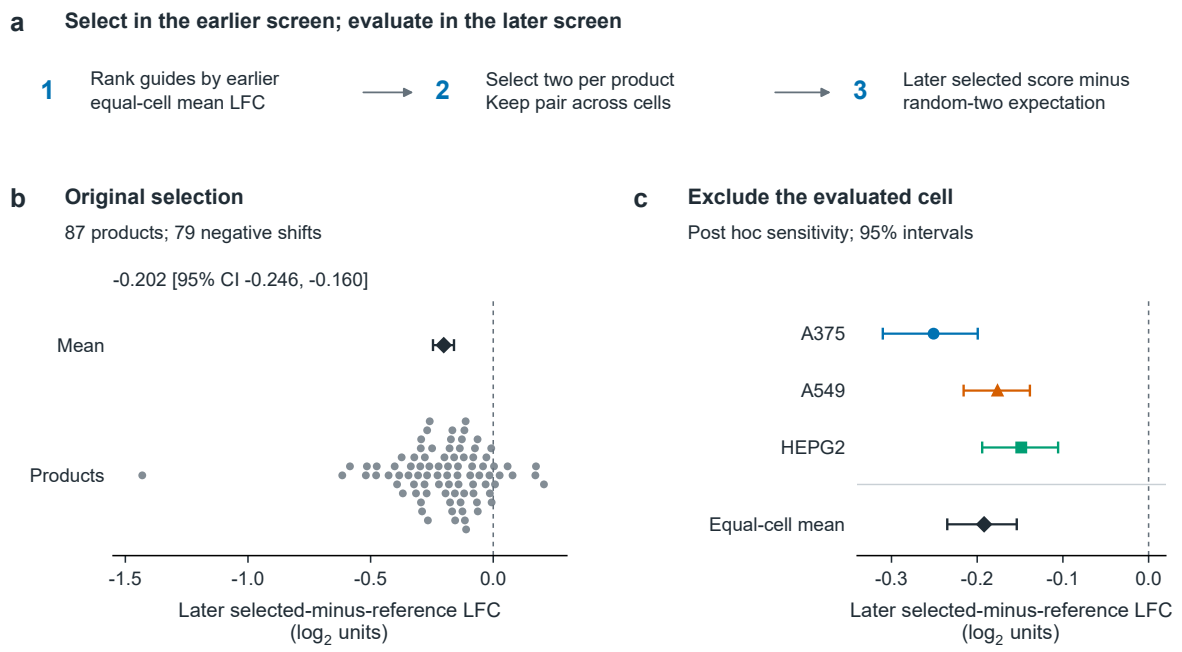

Figure 3: Earlier guide selection shifts later depletion estimates. (a) The original rule selects two guides by earlier-screen equal-cell mean LFC and keeps the pair across three later contexts. The reference is the exact expected later score under uniform random two-guide selection. (b) All 87 product-average shifts; vertical placement separates observations without changing their values. The diamond is the mean with its original nominal 95% dependence-cluster interval. Seventy-nine shifts are negative. (c) Post hoc sensitivity excluding each evaluated cell from the older selection score: three cell estimates and their equal-cell mean, with nominal 95% cluster intervals. Guide pairs can differ between evaluations. Resampling retains cells jointly. Negative values indicate greater depletion. Neither rule is established author practice; the sensitivity uses existing contexts and establishes no molecular mechanism or biological generalization.

The shift is negative in each fixed cell context: −0.25380 in A375, −0.19649 in A549 and −0.15512 in HEPG2. Across products, the median shift is −0.17262 $\log_2$ units and 79 of 87 shifts are negative (Table S27). Single-product omissions range from −0.20655 to −0.18751. The effect therefore extends beyond a small number of extreme products. Greater on-target disruption remains one possible explanation; the score shift does not establish a false-positive mechanism or quantify cell killing. Candidate-rank sensitivity, including ordinary random two-guide variability, is reported in Supplementary Note 2.

Excluding the evaluated cell from the older guide-selection score retains a shift in every context: −0.25090 in A375, −0.17629 in A549 and −0.14859 in HEPG2 (Figure 3c). The equal-product, equal-cell mean is −0.19193 $\log_2$ units, with nominal 95% dependence-cluster interval −0.23492

to −0.15383. Its product-average median is −0.15640, and 81 of 87 shifts are negative (Tables S31-S32). Thus the observed transfer does not require the evaluated cell's older measurements in the selection score. This post hoc sensitivity addresses dependence on the selection information within the existing contexts; it does not supply new biological replication.

A product-level reused-guide advantage remains uncertain after balancing

A separate question is whether an older product score agrees more with newer reused guides than with newer fresh guides. This paired cohort contains 88 exact-chain products with eligible older guides and at least one eligible newer guide in each group, including 345 reused and 138 fresh newer guides. Older scores are eligible-guide medians; newer scores are the corresponding group means. The cohort and endpoints differ from the conditional guide and transferred selection analyses (Tables S05-S06).

The equal-cell mean reused-minus-fresh Spearman agreement gap is 0.20238, with a 97.5% target-cluster interval of 0.00161-0.40243. Unequal guide counts affect the noise of the group averages. In the explicitly post hoc count-balanced analysis, the gap decreases to 0.08632, with interval −0.09388 to 0.26794 (Figure 4a). The balanced interval includes zero, establishing neither a general reused-guide agreement advantage nor equivalence.

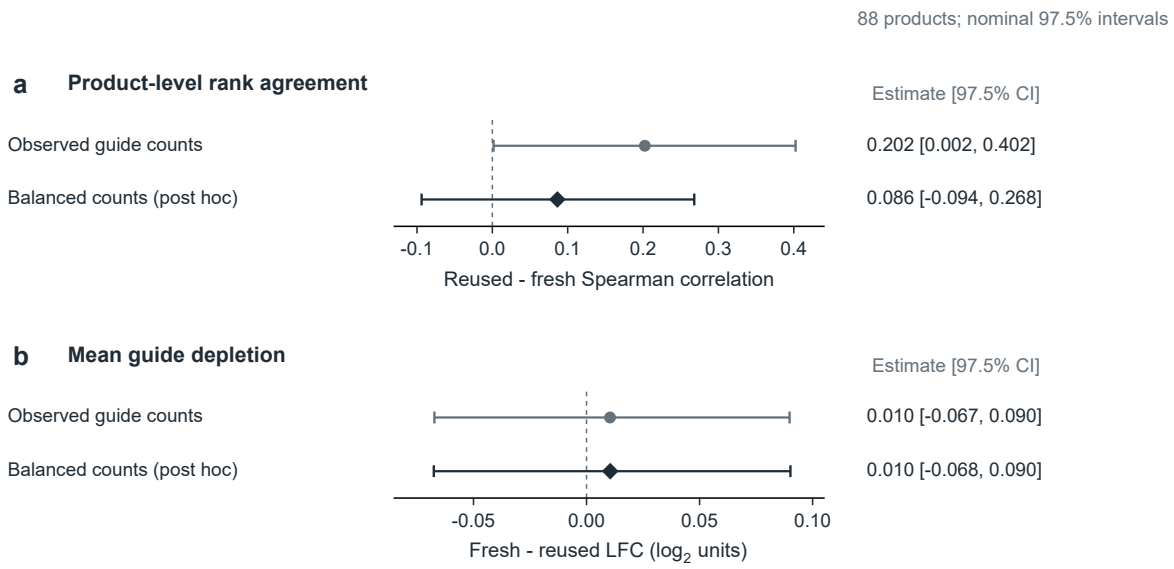

Figure 4: Product-level reused/fresh uncertainty. The separate paired cohort comprises 88 exact products. (a) Reused-minus-fresh Spearman agreement with the older product score. (b) Fresh-minus-reused newer guide LFC. Observed guide counts and the post hoc count-balanced sensitivity retain their original nominal 97.5% target-cluster intervals, with exact rounded values alongside. Circles show observed-count estimates; diamonds show balanced estimates. Guide-selection averaging is included within each target-cluster sample. The count-balanced rank interval crosses zero, supporting neither a general reuse advantage nor equivalence. These endpoints differ from the conditional guide correlation and transferred selection scenario.

The separate fresh-minus-reused mean LFC is 0.01039, with 97.5% interval −0.06729 to 0.08981. Its count-balanced estimate is 0.01047, with interval −0.06756 to 0.09019 (Figure 4b). These paired-endpoint intervals retain the original 97.5% level rather than the 95% level used elsewhere. Complete comparison variants and bootstrap outcomes are supplied in Tables

S07-S10. Persistent relative guide responses and the selection consequence do not resolve this uncertain product-level comparison.

Exact products retain aggregate agreement in a third publication

We next assessed complementary product-level evidence from the Schlesinger et al. screen in untreated A375 cells. Among 239 exact amino-acid and compatible-locus matches to the newer catalogue, metadata rules retained 196 one-to-one products in 192 dependence groups. Eligibility required no source-reported CCDS overlap, at least four designed third-study guides and at least four qualified newer guides. The cohort was fixed using metadata before opening the third-study fitness columns (Table S28).

The mean of the two third-study product LFC summaries correlates with the newer qualified-guide mean: Spearman $\rho = 0.46242$, with nominal 95% dependence-cluster interval 0.34941-0.56949 (Figure 5a). Each source replicate retains positive agreement, and the splice-restricted sensitivity remains positive in 140 products (Tables S29-S30). The source baselines and product summaries differ, so this estimate is not directly comparable with the conditional within-product Pearson correlations.

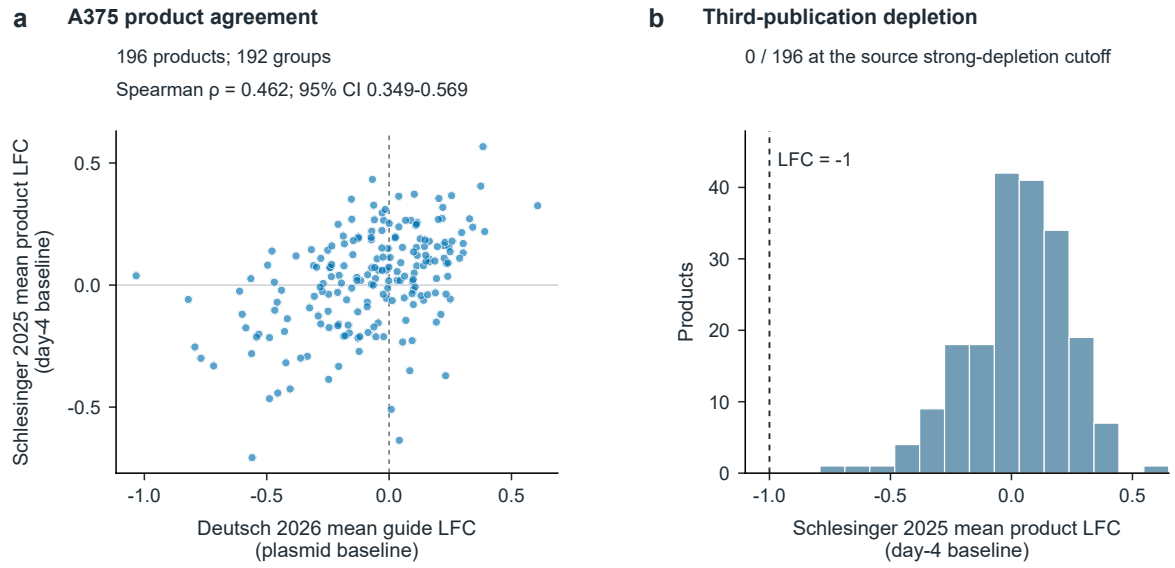

Figure 5: Third-publication aggregate product agreement. Exact sequence-and-locus matches provide a third-publication aggregate A375 comparison. The primary cohort contains 196 products in 192 dependence groups. Rank correspondence is shown with its nominal 95% dependence-cluster interval; source baselines and product summaries differ. The source mean-LFC distribution is displayed alongside its fixed strong-depletion cutoff of -1, which no eligible product reaches. Guide reuse cannot be established from the retrieved data. The result therefore describes aggregate ranking agreement rather than independent-reagent validation or replicated strong individual dependencies.

None of the 196 eligible products reaches the source strong-depletion cutoff of mean LFC at or below -1 (Figure 5b). Agreement therefore concerns relative product rankings in this cohort, rather than replication of strong individual dependencies. The top-40 overlap is descriptive (Figure S2). The retrieved source table contains guide identifiers but no guide sequences or individual-guide outcomes. Guide reuse remains unresolved; these aggregate measurements cannot establish independent-reagent validation or peptide-specific causality.

4 Discussion

Earlier-screen guide selection carries a measurable consequence for later nORF depletion estimates. After complete-product matching and removal of each product-by-cell mean, relative guide responses remain correlated across the two primary screens. Selecting two guides by their earlier activity shifts the later score by -0.202 $\log_2$ units relative to an equal-count random-selection expectation, with negative product-average shifts in 79 of 87 products. Excluding the evaluated cell from the earlier selection score retains a mean shift of -0.192. These results quantify how the history of the intervention can contribute to an apparently repeated phenotype within the observed products and contexts.

The underlying principle has clear precedents. JACKS models guide efficacy shared across screens (Allen et al., 2019), library evaluation compares reagent performance (Hart et al., 2017), and current design strategies balance on-target activity and off-target avoidance (Drepanos et al., 2026). Selecting active guides can improve an experiment. Our contribution is the product-resolved nORF measurement and the size of its consequence for a later depletion estimate. The worked selection rule is retrospective and is not established to have been used by either source study. Its score shift therefore characterizes a strategy within these data, rather than estimating bias in a particular published validation.

The molecular explanation remains open. Greater on-target disruption, sequence-dependent repair, RNA or local regulatory effects and activity at other sites can all contribute to persistent guide responses. Exact-match uniqueness and canonical-CDS avoidance do not exclude mismatch-tolerant off-targets. Copy-number-dependent, gene-independent responses to Cas9 targeting are also documented (Aguirre et al., 2016); our processed endpoints have no newly fitted copy-number correction. Product-by-cell centring removes additive mean responses, but retains multiplicative guide-efficacy differences acting on a true product dependency. Our analytical and synthetic positive control produces both persistent centred responses and a selection shift using only that mechanism (Supplementary Note 5; Figure S3). This compatibility result does not fit or validate an efficacy explanation of the observed data, quantify a technical fraction, or identify which mechanism generates the empirical guide variation. The measurements should therefore be interpreted as intervention responses until a product-specific mechanism is established.

Existing nORF biology shows how that attribution can be addressed. Chen et al. (2020) reported initiation-codon-dependent rescue for selected lncRNA- and uORF-encoded products; Prensner et al. (2021) included selected experiments addressing translation-dependent activity. Hofman et al. (2024) linked the medulloblastoma-associated ASNSD1-uORF/ASDURF dependency to the prefoldin-like chaperone complex. Matsumoto et al. (2017) separated SPAR polypeptide loss from expression of its host lncRNA, while D'Lima et al. (2017) connected No-Body to the decapping complex and P-body abundance. The calcium-regulatory activities of myoregulin and DWORF provide further defined mechanisms in mouse muscle (Anderson et al., 2015; Nelson et al., 2016). These studies supply evidence beyond pooled-score correspondence. Our reanalysis concerns the repeated perturbation component and does not reassess all source experiments supporting individual examples.

Complete-product correspondence is an additional requirement. Initiation mapping and ORF-calling methods distinguish products that a host-gene label alone cannot resolve (Lee et al., 2012; Fields et al., 2015). Annotation guidance provides a shared vocabulary (Mudge et al., 2022), while evolutionary reconstruction and interaction mapping contribute different informa-

tion (Vakirlis et al., 2022; Sandmann et al., 2023). Our 274 relationships were constructed from literal source-gene identity or coding-body overlap; 127 match the complete coding chain and DNA. Those counts describe the comparison made here and should not be read as claims originally made by the source authors. The released product and guide tables allow the units of a cross-study comparison to be inspected directly.

Translation, stable protein detection and phenotype attribution also require distinct measurements. BAG6-associated triage can direct products of noncoding translation towards degradation or membrane targeting (Kesner et al., 2023), and microprotein mass-spectrometry benchmarking highlights variation in identification support (Wacholder et al., 2026). Cancer immunopeptidomics identifies another relevant output: presentation of noncanonical peptides (Chong et al., 2020; Ouspenskaia et al., 2022), with differences between proteome and immunopeptidome detection (Ruiz Cuevas et al. 2021). Antitumour vaccination responses to selected tumour-specific antigens have been demonstrated in mice (Laumont et al., 2018). Presentation and immune recognition need not require a stable protein or a fitness dependency. Consequently, screen agreement informs functional prioritization while therapeutic targeting also needs evidence appropriate to the proposed modality.

The product-level comparisons limit how broadly the selection result can be generalized. In the separate 88-product reused/fresh cohort, balancing guide counts reduces the agreement gap to 0.086, with an interval spanning -0.094 to 0.268. This uncertainty supports neither a general reuse advantage nor equivalence. The conditional guide analysis and the product-level comparison ask different questions. The third publication adds positive aggregate ranking agreement across 196 exact products (Schlesinger et al., 2025), but guide reuse remains unresolved and none reaches its source strong-depletion cutoff. Its product-level Spearman correlation and the conditional within-product Pearson correlations summarize different signals, consistent with the need to match a reproducibility statistic to its intended question (Billmann et al., 2023).

A useful cross-study report should specify the complete reference product, exact guide sequences and reuse, how those guides were selected, guide counts, and the outcome being compared. The comparison can then be described as repeated response to shared interventions or agreement obtained with different reagents. Shared-guide evidence remains informative, including in regulatory screens (Yao et al., 2024). To determine molecular attribution, follow-up should connect a defined translated product to the phenotype through complementary interventions, such as rescue and translation-disrupting controls that preserve the relevant RNA. Disease-specific prioritization additionally requires appropriate normal-cell comparisons. Our correspondence resource supports these choices by exposing which products and reagents have actually been compared.

The inference is conditional on selected products, published measurements and fixed cancer-cell contexts. The analysis developed retrospectively; intervals are nominal and do not account for selection across the entire research program. The context-exclusion extension uses already observed cells, and its guide pairs can differ between evaluations. The additional assay is drawn from the earlier study, while the third publication lacks retrieved guide-resolved outcomes. Source baselines, normalization and product summaries differ. Computational re-expression checks numerical consistency; substantive scientific review and complementary experiments remain necessary. Within these limits, the study supplies a reproducible quantitative account of persistent guide responses and the transferable effect of guide selection in matched nORF screens.

Generative AI disclosure

OpenAI Codex agents contributed substantially to source discovery, literature synthesis, and mathematical exposition

Data availability

The accompanying Supplementary Data and Code archive contains complete product and guide correspondence inventories, eligibility and exclusion records, all analytical outcomes, fixed protocols, figure-source data and reproduction scripts. Source URLs, versions, byte counts and hashes are recorded in the evidence manifest. Original source data remain available from the cited publications and author repositories under their respective terms. The hg38 genome and GENCODE annotation are external pinned references rather than bundled large files. The public data and code release is available at `https://github.com/Kohze/norf-guide-selection` (commit f519e1fb31d248e2e4ae2dc9330c4572a5467c2d). The tested pipeline reproduces the statistical analyses and figures from frozen analysis-ready inputs.

Funding

This work was supported by BBSRC funding.

Conflicts of interest

Robin Gounder declares no competing interests.

Author contributions

Robin Gounder: study design, analysis and writing. Samuel Schmidt: co-writing.

Research scope

This study reanalyses publicly released data and reports no new patient recruitment or experimental measurements.

References

Aguirre AJ, Meyers RM, Weir BA, et al. 2016. Genomic copy number dictates a gene-independent cell response to CRISPR/Cas9 targeting. Cancer Discovery 6(8):914-929. `https://doi.org/10.1158/2159-8290.CD-16-0154`

Allen F, Behan F, Khodak A, et al. 2019. JACKS: joint analysis of CRISPR/Cas9 knockout screens. Genome Research 29:464-471. `https://doi.org/10.1101/gr.238923.118`

Anderson DM, Anderson KM, Chang CL, et al. 2015. A micropeptide encoded by a putative long noncoding RNA regulates muscle performance. Cell 160(4):595-606. https://doi.org/10.1016/j.cell.2015.01.009

Benjamini Y, Hochberg Y. 1995. Controlling the false discovery rate: a practical and powerful approach to multiple testing. Journal of the Royal Statistical Society, Series B 57:289-300. https://doi.org/10.1111/j.2517-6161.1995.tb02031.x

Billmann M, Ward HN, Aregger M, et al. 2023. Reproducibility metrics for context-specific CRISPR screens. Cell Systems 14(5):418-422.e2. https://doi.org/10.1016/j.cels.2023.04.003

Calvo SE, Pagliarini DJ, Mootha VK. 2009. Upstream open reading frames cause widespread reduction of protein expression and are polymorphic among humans. Proceedings of the National Academy of Sciences of the United States of America 106(18):7507-7512. https://doi.org/10.1073/pnas.0810916106

Calviello L, Mukherjee N, Wyler E, et al. 2016. Detecting actively translated open reading frames in ribosome profiling data. Nature methods 13(2):165-170. https://doi.org/10.1038/nmeth.3688

Chen J, Brunner AD, Cogan JZ, et al. 2020. Pervasive functional translation of noncanonical human open reading frames. Science (New York, N.Y.) 367(6482):1140-1146. https://doi.org/10.1126/science.aay0262

Chong C, Müller M, Pak H, et al. 2020. Integrated proteogenomic deep sequencing and analytics accurately identify non-canonical peptides in tumor immunopeptidomes. Nature communications 11(1):1293. https://doi.org/10.1038/s41467-020-14968-9

Chothani S, Ruiz-Orera J, Tierney JAS, et al. 2026. An expanded reference catalog of translated open reading frames for biomedical research. Nucleic Acids Research 54(6):gkag234. https://doi.org/10.1093/nar/gkag234

D'Lima NG, Ma J, Winkler L, et al. 2017. A human microprotein that interacts with the mRNA decapping complex. Nature chemical biology 13(2):174-180. https://doi.org/10.1038/nchembio.2249

Deutsch EW, Kok LW, Mudge JM, et al. 2026. Expanding the human proteome with microproteins and peptideins. Nature 654:813-825. https://doi.org/10.1038/s41586-026-10459-x

Drepanos LM, Srikanth S, Kaplan EG, et al. 2026. Balancing off-target and on-target considerations for optimized CRISPR-Cas9 knockout library design. Cell Genomics 6(5):101190. https://doi.org/10.1016/j.xgen.2026.101190

Erhard F, Halenius A, Zimmermann C, et al. 2018. Improved Ribo-seq enables identification of cryptic translation events. Nature methods 15(5):363-366. https://doi.org/10.1038/nmeth.4631

Fields AP, Rodriguez EH, Jovanovic M, et al. 2015. A Regression-Based Analysis of Ribosome-Profiling Data Reveals a Conserved Complexity to Mammalian Translation. Molecular cell 60(5):816-827. https://doi.org/10.1016/j.molcel.2015.11.013

GENCODE. 2026. Phase-1 Ribo-seq ORF catalogue and GENCODE v50 human transcript annotations. Accessed 15 October 2026. https://www.gencodegenes.org/pages/riboseq_orfs/ ; https://www.gencodegenes.org/human/

Hart T, Tong AHY, Chan K, et al. 2017. Evaluation and Design of Genome-Wide CRISPR/SpCas9 Knockout Screens. G3 7:2719-2727. https://doi.org/10.1534/g3.117.041277

Hofman DA, Ruiz-Orera J, Yannuzzi I, et al. 2024. Translation of non-canonical open reading frames as a cancer cell survival mechanism in childhood medulloblastoma. Molecular cell 84(2):261-276.e18. https://doi.org/10.1016/j.molcel.2023.12.003

Ingolia NT, Ghaemmaghami S, Newman JR, et al. 2009. Genome-wide analysis in vivo of translation with nucleotide resolution using ribosome profiling. Science (New York, N.Y.) 324(5924):218-223. https://doi.org/10.1126/science.1168978

Ingolia NT, Lareau LF, Weissman JS. 2011. Ribosome profiling of mouse embryonic stem cells reveals the complexity and dynamics of mammalian proteomes. Cell 147(4):789-802. https://doi.org/10.1016/j.cell.2011.10.002

Ji Z, Song R, Regev A, et al. 2015. Many lncRNAs, 5'UTRs, and pseudogenes are translated and some are likely to express functional proteins. eLife 4:e08890. https://doi.org/10.7554/elife.08890

Kesner JS, Chen Z, Shi P, et al. 2023. Noncoding translation mitigation. Nature 617(7960):395-402. https://doi.org/10.1038/s41586-023-05946-4

Laumont CM, Vincent K, Hesnard L, et al. 2018. Noncoding regions are the main source of targetable tumor-specific antigens. Science translational medicine 10(470):eaau5516. https://doi.org/10.1126/scitranslmed.aau5516

Lee S, Liu B, Lee S, et al. 2012. Global mapping of translation initiation sites in mammalian cells at single-nucleotide resolution. Proceedings of the National Academy of Sciences of the United States of America 109(37):E2424-32. https://doi.org/10.1073/pnas.1207846109

Matsumoto A, Pasut A, Matsumoto M, et al. 2017. mTORC1 and muscle regeneration are regulated by the LINC00961-encoded SPAR polypeptide. Nature 541(7636):228-232. https://doi.org/10.1038/nature21034

Mudge JM, Ruiz-Orera J, Prensner JR, et al. 2022. Standardized annotation of translated open reading frames. Nature biotechnology 40(7):994-999. https://doi.org/10.1038/s41587-022-01369-0

Nelson BR, Makarewich CA, Anderson DM, et al. 2016. A peptide encoded by a transcript annotated as long noncoding RNA enhances SERCA activity in muscle. Science (New York, N.Y.) 351(6270):271-275. https://doi.org/10.1126/science.aad4076

Ouspenskaia T, Law T, Clauser KR, et al. 2022. Unannotated proteins expand the MHC-I-restricted immunopeptidome in cancer. Nature biotechnology 40(2):209-217. https://doi.org/10.1038/s41587-021-01021-3

Prabakaran S, Hemberg M, Chauhan R, et al. 2014. Quantitative profiling of peptides from RNAs classified as noncoding. Nature Communications 5:5429. https://doi.org/10.1038/ncomms6429

Prensner JR, Enache OM, Luria V, et al. 2021. Noncanonical open reading frames encode functional proteins essential for cancer cell survival. Nature Biotechnology 39:697-704. https://doi.org/10.1038/s41587-020-00806-2

Ruiz Cuevas MV, Hardy MP, Hollý J, et al. 2021. Most non-canonical proteins uniquely populate the proteome or immunopeptidome. Cell reports 34(10):108815. https://doi.org/10.1016/j.celrep.2021.108815

Sandmann CL, Schulz JF, Ruiz-Orera J, et al. 2023. Evolutionary origins and interactomes of human, young microproteins and small peptides translated from short open reading frames. Molecular cell 83(6):994-1011.e18. https://doi.org/10.1016/j.molcel.2023.01.023

Schlesinger D, Dirks C, Navarro C, et al. 2025. A large-scale sORF screen identifies putative microproteins involved in cancer cell fitness. iScience 28(3):111884. https://doi.org/10.1016/j.isci.2025.111884

Slavoff SA, Mitchell AJ, Schwaid AG, et al. 2013. Peptidomic discovery of short open reading frame-encoded peptides in human cells. Nature chemical biology 9(1):59-64. https://doi.org/10.1038/nchembio.1120

Vakirlis N, Vance Z, Duggan KM, et al. 2022. De novo birth of functional microproteins in the human lineage. Cell reports 41(12):111808. https://doi.org/10.1016/j.celrep.2022.111808

van Heesch S, Witte F, Schneider-Lunitz V, et al. 2019. The Translational Landscape of the Human Heart. Cell 178(1):242-260.e29. https://doi.org/10.1016/j.cell.2019.05.010

Wacholder A, Deutsch EW, Kok LW, et al. 2026. Community benchmarking and evaluation of human unannotated microprotein detection by mass spectrometry based proteomics. Nature communications 17(1):1241. https://doi.org/10.1038/s41467-025-68002-x

Yao D, Tycko J, Oh JW, et al. 2024. Multi-center integrated analysis of non-coding CRISPRi screens. Nature Methods 21:723-734. https://doi.org/10.1038/s41592-024-02216-7

Supplementary information

Supplementary reader index

Product and guide correspondence: Tables S01-S05, S11 and S19-S22. Conditional guide outcomes and sensitivities: S12-S15. Reused/fresh comparisons: S06-S10. Additional assay: S21-S24. Transferred selection: S25-S27. Third-publication correspondence and outcomes: S28-S30. The separate negative within-source-gene analysis is retained in S16-S18 and Supplementary Note 1. Table identifiers are preserved across manuscript versions. Context-excluded guide selection: S31-S32 and Supplementary Note 4. Synthetic shared-efficacy positive control: S33-S34 and Supplementary Note 5.

Tables S01-S34 are supplied as CSV files. S01-S32 contain empirical data and analyses; S33-S34 contain the synthetic positive control. Literal source identifiers and unavailable values remain explicit; missing values are not functional negatives. Coordinates are zero-based, half-open hg38 intervals. LFC outcomes and source fitness scores remain distinct. The 127 cross-screen exact-chain pairs in Tables S05-S06 and the 127 eligible within-source-gene pairs in Tables S16-S18 are different cohorts.

Table S01. Older source coverage

All 553 candidates and reference-product comparison denominators, including two unmapped candidates.

Table S02. Newer source coverage

All 2,196 newer analytical targets and relationship denominators.

Table S03. Complete product relationships

All 274 constructed relationships and the fixed identity hierarchy.

Table S04. Exact guide identities

All 744 exact 20-nucleotide strings shared across the two analytical candidate inventories.

Table S05. Reused/fresh outcome eligibility

All 127 exact-chain pairs, counts, dependence groups and primary/sensitivity inclusion flags.

Table S06. Product and cell guide-group scores

All 381 pair-cell records, including unavailable group scores for ineligible pairs.

Table S07. Reused/fresh analysis variants

All declared unbalanced comparison points and cohort sizes.

Table S08. Cell-specific product agreement

All declared source-cell Spearman correlations.

Table S09. Common-cell omissions

All three omissions in the paired reused/fresh analysis.

Table S10. Balanced target-bootstrap estimates

All 2,000 guide-plan-averaged target-bootstrap estimates and endpoints.

Table S11. Conditional shared-guide eligibility

All 127 exact-chain products and metadata eligibility for the conditional cohorts.

Table S12. Paired shared-guide residuals

All 1,413 main-cohort paired measurements, guide metadata, centring and weights.

Table S13. Conditional sensitivity points

Every declared conditional guide-correlation sensitivity point.

Table S14. Conditional product omissions

All 105 single-product omission estimates.

Table S15. Covariance decomposition

Every main-cohort guide-cell row and its cell-centred, between-product and within-product components.

Table S16. Within-source-gene pair eligibility

All 145 pairs among 248 ORFs in 117 source genes, with partner-exclusive guide counts and flags.

Table S17. Within-source-gene outcomes

All 1,160 pair-cell records, including strict-cohort and sensitivity flags.

Table S18. Exploratory Ewing contrasts

All 127 strict within-source-gene contrasts, joint BH results and every lead-qualification flag.

Table S19. Older complete models

All 553 older candidates with mapping, literal coding sequence, translation and source discrepancy fields.

Table S20. Newer complete models

All 2,196 newer complete reference-product models.

Table S21. Follow-up product eligibility

All 127 complete-product pairs and metadata guide counts for the additional assay.

Table S22. Follow-up guide eligibility

All 720 newer guide-product rows with secondary-library membership, literal-name correspondence, count coverage and exclusions.

Table S23. Follow-up paired residuals

All 144 main follow-up guide-cell pairs, source endpoints, centred responses and equal-product/equal-cell weights.

Table S24. Follow-up sensitivity estimates

All seven declared primary and sensitivity points, with their cohort sizes.

Table S25. Older-only guide selections

All 1,251 guide-cell records in the 87-product selection cohort, including older-only guide ranks and selection flags.

Table S26. Product-cell selection contrasts

All 261 product-cell contrasts between the older-selected two-guide score and the exact random-pair expectation.

Table S27. Product shifts and rank sensitivity

All 87 product-average selection shifts, selected sequences and descriptive top-20% ranking sensitivity, including random-pair selection fractions.

Table S28. Third-study complete-product correspondence and eligibility

All 239 exact-AA, compatible-locus pairs, identity multiplicities, designed and qualified guide counts, exact coding-block matches, eligibility and dependence groups before outcome comparisons.

Table S29. Third-study product outcomes and rank overlap

All 196 eligible product pairs, both source 2025 A375 replicate summaries, source 2026 mean/median scores, dependence groups, sensitivity flags and descriptive top-20% memberships.

Table S30. Third-study sensitivity estimates

All declared descriptive rank-correlation sensitivities, product counts and dependence-group counts.

Table S31. Context-excluded selection sensitivity

All three evaluated-cell estimates and their equal-cell mean, with nominal post hoc 95% dependence-group intervals, product medians and direction counts. The original primary interval is retained in the main text; new sensitivity intervals use seed 20261026.

Table S32. Context-excluded product-cell contrasts

All 261 product-cell selections and contrasts, including selection contexts, exact chosen guide sequences, selected and random-expectation scores, guide counts and dependence groups.

Table S33. Fixed synthetic efficacy-control summary

All 12 condition-by-noise summaries. These are synthetic positive-control outputs and are not additional observed nORF effects. An undefined correlation is left missing when all residuals are zero.

Table S34. All synthetic product-selection contrasts

All 24,000 synthetic product-by-condition-by-noise contrasts and selected guide indices. Complete synthetic guide-by-context outcomes are included in the accompanying compressed NumPy archive.

Supplementary Note 1. Negative within-source-gene discovery analysis

A separate retrospective protocol, saved before computing its new within-gene contrasts, compared distinct nORFs assigned to the same source gene across all eight newer cell contexts. The complete universe comprised 145 pairs among 248 ORFs in 117 genes. Strict eligibility required at least four qualified guides per product, with the whole protospacer avoiding the partner coding body, mapped references and finite expression/outcomes in every cell. This yielded 127 pairs across 103 genes and 102 dependence groups. This cohort is distinct from the 127 cross-screen exact-chain matches.

The specified weighted two-way fixed-effect association related within-pair depletion differences to within-pair differences in $\ln(1+x)$, where x is the source ribosome-abundance measurement, with equal source-gene weight and 2,000 dependence-cluster samples. The coefficient per residual predictor SD was 0.00171 depletion units (nominal 95% interval -0.00681 to 0.00983);

incremental total R^2 was 0.00005598. Cell omissions changed its sign, and minimum-three-guide, splice-safe and source-Chronos point sensitivities did not provide a useful consistent association. A separate weighted dummy-variable regression and explicit cluster-draw replays reproduced the main calculation.

The declared exploratory Ewing contrast compared A673, CADOES1 and RDES with five other cell contexts. Welch context-comparison P values were adjusted jointly across all 127 eligible pairs using BH (Benjamini and Hochberg, 1995). No pair had $q < 0.05$, and none passed the complete effect, expression and guide-stability lead criteria. Cell lines were treated as observational contexts, not patient or experimental replicates. No threshold was relaxed and no medical candidate was accepted. The complete protocol, denominators, outcomes, sensitivity points and null lead table are retained for transparency.

Supplementary Note 2. Candidate-rank sensitivity to guide choice

Using the fixed 18-product priority budget, the older-selected two-guide and all-guide scores retain 13 of the same products; five products are replaced. Across the 1,000 uniform random-pair plans, median overlap with the all-guide priority set is 12 products (range 7-16). The activity-selected overlap therefore lies within ordinary two-guide selection variability in this scenario. These descriptive comparisons expose dependence of a shortlist on reagent choice without identifying a gold-standard target set, showing unusually poor rank stability or declaring individual targets incorrect. All rankings, selection fractions and sequence identities are supplied in Table S27 and the complete random-plan outputs.

Supplementary Note 3. Descriptive covariance decomposition

After removing each screen's cell mean using equal product weights, total cross-screen covariance is 0.075105 in squared LFC units. The between-product component is 0.028436 and the within-product component is 0.046669; the components sum to the total by construction. Within-product responses contribute 62.1% of this covariance in the selected shared-guide cohort. This is neither a causal variance partition nor a technical-bias fraction. On-target disruption efficacy can contribute to the within-product term. Table S15 provides all rows used in this descriptive decomposition.

Supplementary Note 4. Selection without the evaluated cell

The unchanged 87-product cohort permits three evaluations, each excluding its evaluated cell from the older guide-selection score. All outcomes, selections and direction counts are retained in S31-S32. Different evaluations can use different guide pairs, so the equal-cell summary describes this context-specific strategy. The original strategy gives a median product-average shift of -0.17262 with 79 negative and eight positive shifts; the context-excluded strategy gives a median of -0.15640 with 81 negative and six positive shifts. No exact-zero product-average shift occurs. These counts describe the selected products and are not a binomial test of biological independence. Complete resampling draws and bootstrap estimates accompany the code. Original endpoints and their original bootstrap interval are preserved; additional resampling of the original strategy is retained solely as an implementation cross-check.

Supplementary Note 5. Shared efficacy as a positive control

A nonzero centred correlation need not reflect an effect separate from a genuine product dependency. Consider a model with guide efficacy a, screen- and cell-dependent product effect theta, and independent measurement noise epsilon (Equation S1). Efficacy is shared across screens and contexts. Bars denote means over the same eligible guide set used for centring. This is an illustrative multiplicative model motivated by established guide-efficacy modelling, rather than a fit of JACKS or an inferred editing-efficiency measurement (Allen et al., 2019).

$$\begin{aligned} \mathrm{LFC}_{s,p,g,c} &= a_{p,g}\theta_{s,p,c} + \epsilon_{s,p,g,c} \\ e_{s,p,g,c} &= (a_{p,g} - \bar{a}_p)\theta_{s,p,c} + (\epsilon_{s,p,g,c} - \bar{\epsilon}_{s,p,c}) \end{aligned} \tag{S1}$$

The centred response still contains efficacy multiplied by the product effect. In a noiseless example, each product has four efficacies, 0.25, 0.50, 0.75 and 1.00. The three older context effects are -0.4, -0.8 and -1.2; later effects are -0.6, -1.2 and -1.8. Because these effect vectors are proportional, the centred old-later correlation is exactly one. Selecting the two lowest older mean LFCs selects efficacies 0.75 and 1.00. Their mean efficacy is 0.875, compared with 0.625 over all four guides; the later equal-context mean effect is -1.2. Their exact selection contrast is therefore Equation S2.

$$\Delta = (0.875 - 0.625) \times (-1.2) = -0.300 \tag{S2}$$

We generated 2,000 synthetic products at each independent Gaussian noise standard deviation of 0, 0.25, 0.50 and 1.00, using `default_rng` seed 20261028. All products have the fixed four efficacies and three context effects above. The same realised noise arrays are used in the three comparison conditions at each noise level. The shared-efficacy condition retains the same guide identities. The equal-efficacy control sets every efficacy to 0.625. The label-permutation control retains heterogeneous efficacies but independently permutes each later product's four whole-guide profiles jointly across its three contexts. It breaks the correspondence being tested without removing the product-level dependency.

For each condition we apply the same centring and equal-product/equal-context weighting as in the empirical analysis, and select two guides by their older equal-context mean, resolving ties by guide index. The later selected-minus-all-guide contrast equals the selected-minus-uniform-random-pair expectation. Enumerating all six guide pairs verifies this equality for every synthetic product and context. A separate implementation reproduces all 12 aggregate summaries, all 24,000 product-selection contrasts and all 72,000 product-context random-pair expectations. With equal efficacy and zero noise, the correlation is undefined because residual variance is zero; its selection contrast is zero.

Figure S3 reports the complete fixed noise grid. It demonstrates that genuine dependencies combined with shared on-target efficacy can generate the two empirical estimands. The finite permuted and noisy equal-efficacy controls fluctuate around zero. The parameters were not tuned to reproduce the observed correlation or selection shift. This positive control does not estimate the cause, magnitude or fraction of the observed correspondence attributable to efficacy; failure of a particular synthetic model to match observed values would not establish off-target activity or peptide causality.

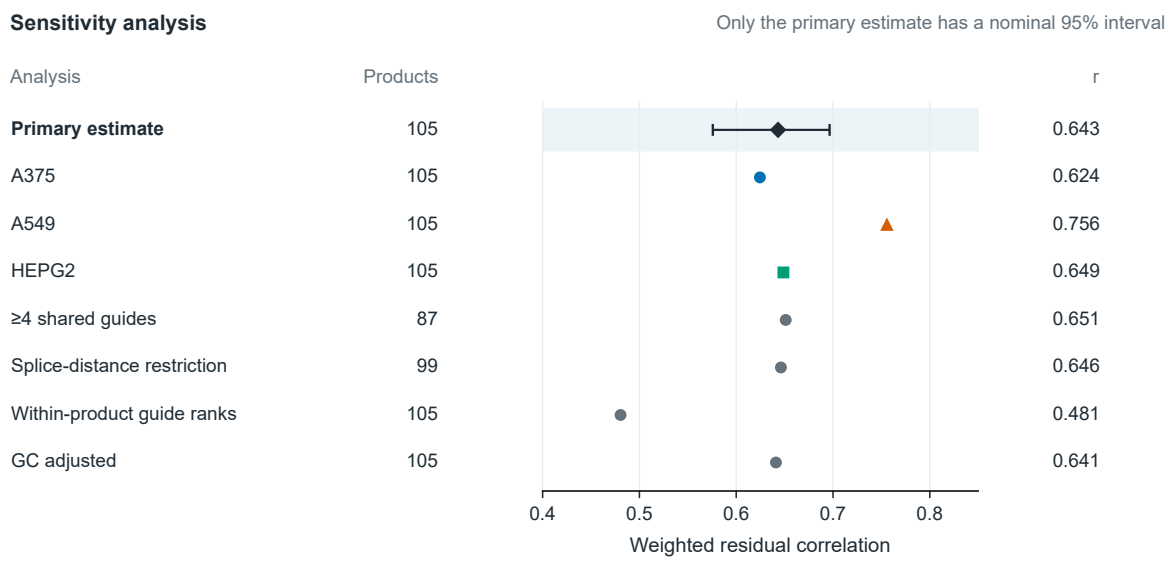

Figure S1: Declared conditional-guide sensitivities. Points show the main estimate, fixed-cell estimates and the previously specified guide-count, splice-distance, guide-rank and GC sensitivities. Only the main estimate displays its nominal 95% dependence-cluster interval; other points are descriptive. All products and fixed cells remain dependent under the stated design. Source: Table S13 and the frozen conditional analysis.

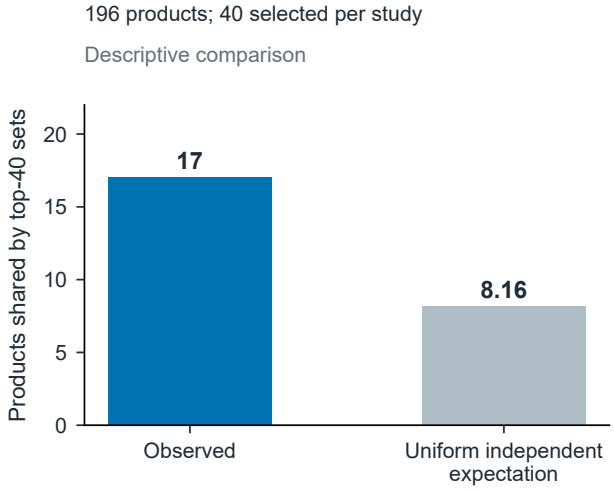

Figure S2: Descriptive third-publication priority overlap. Among 196 matched products, the two top-40 relative-depletion sets share 17 products, compared with 8.16 under uniform independent selection. This reference is an arithmetic expectation, not a calibrated biological null or an independent-reagent validation test. None of the 196 eligible products meets the source strong-depletion cutoff. Source: Table S29.

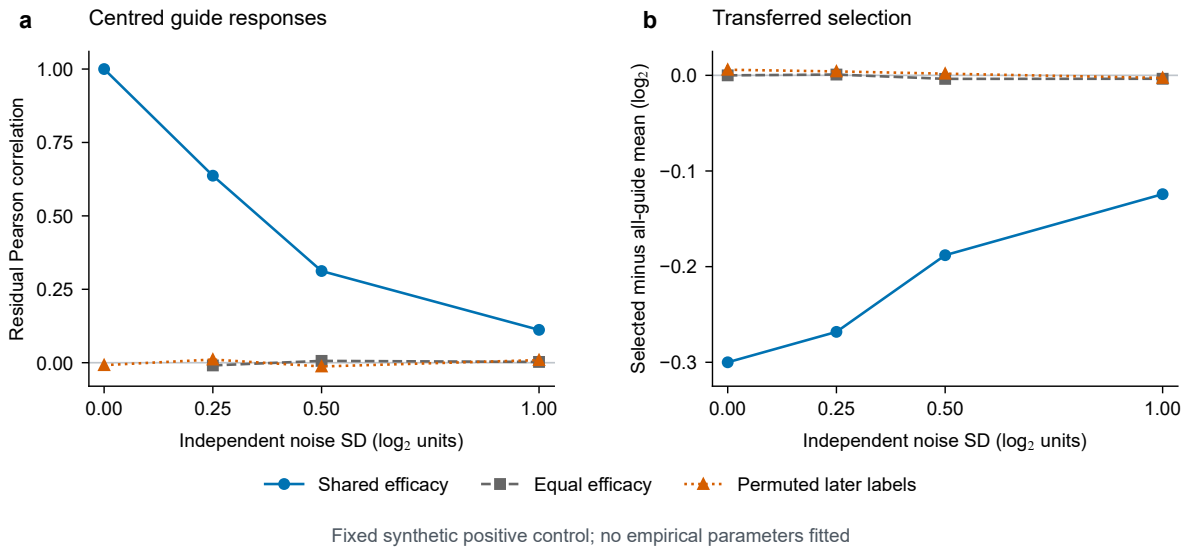

Figure S3: Fixed synthetic shared-efficacy positive control. (a) Product-by-cell-centred Pearson correlation. (b) Later selected-minus-all-guide mean after selecting the two most-depleting older guides. Each point summarizes 2,000 synthetic products with four guides and three fixed contexts. Shared efficacy is compared with equal efficacy and later-guide label permutation over the complete fixed noise grid. The noiseless shared-efficacy values are exactly 1 and −0.300; equal-efficacy correlation is undefined at zero noise and is omitted. Connecting lines guide the eye. Values are illustrative simulation summaries, not confidence intervals, a fitted explanation or additional empirical discoveries. No model parameter was fitted to the manuscript data. Source: Tables S33-S34; Supplementary Note 5.