



RsiingSI · Independent Research

New treatment directions for allergic rhinitis: mechanisms already tested in adjacent type-2 diseases but never registered in rhinitis, crossed with a replicating nasal transcriptomic signature

Agent Team¹ · scalingX¹

¹RsiingSI — the open-source org scaling recursive self-improvement for bioscience.

Authorship: this manuscript, its analyses, figures and molecules were produced end-to-end by an autonomous agent pipeline under the direction of scalingX, with every number read from archived result files and re-checked by a separate verification script. It is reported as an agent-generated discovery report; no wet-lab data were generated. All code and derived tables are archived with the manuscript.

Hypothesis-generating computational study · 2026-10-20 · Every number is machine-generated from `results/paper4_facts.json`; figures in `figures/fig18-fig24`; the registry pull, the phase-3 outcome audit, the end-organ outcome audit, the PubMed literature cross-check, the inverse literature sweep, the signature replication test and the drug screen are each reproducible from public APIs/archives.

Abstract. Paper 3's transcriptome-driven repurposing screen returned nothing because its disease signature did not replicate. We therefore attacked the same question — what should be tried next in allergic rhinitis — from two directions that do not depend on that failed signature, and intersected them. **(i) Clinical registry.** We pulled every interventional study registered on ClinicalTrials.gov for allergic rhinitis and for seven adjacent type-2 / inflammatory diseases (10,007 studies in total, yielding 9,486 drug-like intervention arms and 6,939 canonical drug mentions over 1,999 distinct substances). Rhinitis (1,276 registered studies) has been dominated by a very short list — intranasal glucocorticoids, second-generation antihistamines, leukotriene antagonists and allergen immunotherapy — while 1,535 substances with registered studies in adjacent diseases have zero registered rhinitis studies; of those, 407 have ≥ 2 adjacent-disease studies and 34 carry an annotatable mechanism. **(ii) Transcriptome.** GSE261239 supplies a paired nasal allergen-challenge contrast in birch-allergic subjects. It replicates with GSE19187 ($p = 0.297$; 141 concordant vs 14 discordant genes, odds 10.1), unlike the paper-3 cohort pair ($p = -0.090$). Building a v2 signature from those two contrasts and re-running the identical screen, the drug-level score becomes *transferable* between cohorts ($p = 0.302$ vs -0.090) and nominal hits rise to 129 against 40 expected — but **still no drug survives FDR control** (minimum FDR 0.154). **(iii) Intersection.** Crossing the two sources gives 34 ranked hypotheses that combine a real clinical precedent with a target direction that the nasal transcriptome does not contradict; the strongest are mepolizumab (IL5), benralizumab (IL5RA), lebrikizumab (IL13), tralokinumab (IL13), crisaborole (PDE4). A third check on the same registry grades the evidence behind those precedents: re-querying the 34 mechanism-annotated gaps by intervention name finds 220 phase-3 studies, 146 of which posted results; 78 of their primary-endpoint analyses are significant, 20 are not, and 48 report no numeric p . The verdicts split 11 positive / 7 mixed / 2 negative / 14 with no posted phase-3 p — and the negative set contains the DP2 antagonist fexiprant, whose 7 posted phase-3 studies include **no** significant primary endpoint: registration counts measure commitment, not efficacy. **(iv) The end organ.** The mirror audit of rhinitis itself shows why even that reading needs outcomes attached: of 1,276 registered rhinitis studies 269 posted results (93 significant, 38 not, 138 without a numeric p), and the mechanisms the nose has already rejected include phenylephrine and a histamine-H3 antagonist. One further sweep isolates a third kind of gap that is neither "untried" nor "failed" but **registry-silent**: 3 completed phase-3, placebo-controlled rhinitis trials of a sublingual bacterial lysate with symptom endpoints at $n = 76-80$, none of which posted results, while the thymic/spleen-peptide immunomodulators return **zero** registered rhinitis studies. **(v) Literature cross-check.** PubMed searches for the same candidates show that a registry gap is not a literature gap: 22 of 34 candidates have at least one allergic-rhinitis publication. Reading every retrieved record (80 records, 69 unique documents) shows that only 2 of them are human studies in which the candidate was given and rhinitis symptoms were an endpoint — a 2001 roflumilast trial that was never registered on ClinicalTrials.gov, and a 2024 upadacitinib comorbidity analysis — while 38 are reviews and 25 are studies in other diseases; 12 candidates have no rhinitis publication at all. The same check finds two randomised rhinitis trials of spleen aminopeptide, a compound with **zero** registered studies, and identifies the publication that resolves one of the three unposted lysate trials. **(vi) The reverse sweep.** Inverting the same question — which drugs have published human rhinitis trials but appear in none of the 849 registered rhinitis studies — and matching the titles of all 3,391 PubMed records indexed as allergic-rhinitis trials against a 6,565-name drug lexicon finds 79 such drugs. Hand-reading every match for the 35 of them that carry at least two population trials (every matched title is listed in `results/lit_inverse_title_audit.csv`) shows that 25 are genuine AR-population programmes (314 trial publications), that 21 are still marketed drugs and 3 were withdrawn for cardiotoxicity; the pattern is explained by the pre-registration era, not by concealment. The value of that sweep is the calibration it supplies: "zero registered in rhinitis" is a statement about the registry's clock as much as about the drug — roflumilast's unregistered 2001 trial is one member of a large class, not an anomaly. We emphasise that the entire output is **hypothesis generation**: it identifies mechanisms that were demonstrably skipped and are testable in rhinitis, not drugs shown to work.

Keywords: allergic rhinitis; drug repurposing; ClinicalTrials.gov; evidence gaps; nasal transcriptome; type-2 inflammation; hypothesis generation.

1 Background: why a registry view is necessary

Paper 3 screened 4,010 compounds against a three-cohort allergic-rhinitis signature and found 0 FDR-significant drugs, with 25 nominal hits — *fewer* than the 40 expected by chance — because the discovery and validation cohort signatures were effectively uncorrelated ($p = -0.090$). A negative result of that kind has a specific diagnostic meaning: the limiting factor is the signature, not the library. Two responses are available. One is to repair the signature (section 2). The other is to change evidence source entirely: what has actually been *tried* in rhinitis, and what has been tried in diseases close enough that the result should generalise, is recorded in a registry and does not depend on any expression measurement. Section 3 does that. The two are then crossed in section 4.

Framing. "Not registered in rhinitis" is a statement about registration records, not about biology. A mechanism may have been tested in rhinitis without ever being registered on ClinicalTrials.gov — for instance in an academic nasal-challenge study, or in a Chinese registry that this API does not cover. Every gap named below should be read as a *testable, documented omission from the registered trial record*, which is exactly the kind of omission a sponsor can act on, and never as "nobody has ever tried this".

2 Route A — repairing the signature with a paired allergen-challenge contrast

GSE261239 profiled nasal mucosa in birch-pollen-allergic and non-allergic subjects under both saline and birch-pollen challenge (11/11 allergic, 12/12 non-allergic), processed here with the same DESeq2 pipeline used throughout this series. It contributes two contrasts: a between-subject "AR-saline vs control-saline" contrast, which is the direct analogue of the paper-3 discovery contrast, and a within-subject "allergen vs saline" contrast, which is a clean pharmacological activation of the type-2 pathway. Figure 18 and the table below show that the second contrast is the one that replicates.

Contrast pair	Shared genes	ρ	Concordant genes	Discordant genes	Odds concordance
GSE19187 AR vs HC — GSE261239 allergen challenge	14,987	0.297	141	14	10.1
GSE43523 SAR vs CTRL — GSE261239 AR vs CTRL	14,119	0.150	318	136	2.3
GSE329964 AH+AR vs HC — GSE261239 allergen challenge (opposite sign)	15,996	-0.133	255	469	0.5
GSE19187 — GSE329964 (the paper-3 pair that failed)	15,782	-0.105	151	148	1.0
GSE19187 — GSE43523 (paper-3 pair)	16,841	0.044	57	12	4.8

The between-subject contrast does *not* rescue the signature: its correlation with GSE19187 is $\rho = -0.066$ and with GSE329964 $\rho = -0.088$, i.e. as weak as the paper-3 pairs. The allergen-challenge contrast, by contrast, reaches $\rho = 0.297$ against GSE19187 with an odds of concordance of 10.1 among genes significant in both. That is still a modest correlation, and it is *not* reproduced by every cohort: against GSE329964 the challenge contrast is anti-correlated ($\rho = -0.133$) and against GSE43523 it is flat ($\rho = -0.002$). One paper-3 pair, GSE19187-GSE43523, does show sign concordance among the genes significant in both (57 versus 12, odds 4.8) even though its overall ρ is only 0.044; with just 69 genes significant in both contrasts, that is a low-powered observation, so we record it as suggestive rather than as a third replicating pair. We therefore used the challenge contrast where it is supported — combined with GSE19187 into signature v2 (27,434 genes; weights $\sqrt{32}$ and $\sqrt{2}$) — and treat GSE43523 and GSE329964 as held-out external checks rather than as screening inputs.




Figure 18. Signature replication. (A) Pairwise Spearman ρ between every nasal contrast used in this series. The only strongly positive off-diagonal entry is GSE19187 $\times$ GSE261239 allergen challenge ($\rho = 0.297$); every paper-3 pair sits between -0.105 and +0.05. (B) Of the genes significant at $p < 0.05$ in both contrasts, the concordant and discordant counts give an odds ratio of 10.1 for the challenge contrast, versus 1.02 for the pair that failed in paper 3. (C) Gene-level scatter for the replicating pair. The cloud is far from tight: v2 is an improvement in transferability, not a repaired signature.

Re-running the identical target-inversion screen (same drug library, same 3,995-drug space, same size-stratified permutation null, same $+1/(n+1)$ convention, same BH-FDR) on signature v2 changes the *behaviour* of the screen without changing its headline:

Signature	Drugs scored	FDR<0.05	$p < 0.01$ (observed)	$p < 0.01$ (expected)	$\rho(S_{\text{discovery}}, S_{\text{validation}})$
paper 3 (GSE19187 / GSE329964)	4,010	0	25	40.1	-0.090
paper 4 v2 (GSE19187 / GSE261239 challenge)	3,995	0	129	40.0	0.302




Figure 19. The screen before and after. (A) With signature v2 the drug score computed from one cohort transfers to the other ($\rho = 0.302$); (A') under the paper-3 signature the same statistic is -0.090, i.e. the ranking carried no transferable information. (B) The permutation p -value distribution shows a clear excess at small p (129 drugs at $p < 0.01$ vs 40 expected), yet the minimum FDR is 0.154 — the signal is broad and weak, not concentrated in any individual drug.

Three things should be said plainly. First, **the negative headline stands**: minimum FDR 0.154, so no candidate is FDR-significant. Second, the improvement is real but limited to transferability: the ranking is no longer anti-correlated between cohorts, and the excess of nominal hits grew from 25 (below chance) to 129 (about 3.2 $\times$ chance). Third, external checks are mixed: v2 agrees weakly with the held-out GSE43523 cohort ($\rho = 0.036$) and is anti-correlated with GSE329964 ($\rho = -0.137$), so v2 is not a universal signature either. Calibration against known anti-rhinitis drugs does not separate the two screens (mean S -percentile 0.707 in paper 3 vs 0.730 here; 24 of 32 known drugs are in the target table), and because most of those drugs have a single mapped target they cluster onto a handful of identical scores. Restricting to drugs whose targets are sign-concordant between the two input cohorts leaves 442 drugs; the top of that list is shown below, and it is dominated by one approved intranasal antiallergic drug followed by cell-cycle kinase inhibitors that are not plausible rhinitis therapeutics.

Drug	S (sign-concordant targets only)	Targets tested	Targets agreeing	S (all targets)	Mechanism of action
olopatadine	3.70	6	3	1.87	histamine receptor antagonist
CYC116	2.36	3	3	2.36	Aurora kinase inhibitor
JNJ-7706621	2.30	4	3	2.15	CDK inhibitor

Drug	S (sign-concordant targets only)	Targets tested	Targets agreeing	S (all targets)	Mechanism of action
PHA-680632	2.12	3	3	2.12	Aurora kinase inhibitor
CCT129202	2.12	3	3	2.12	Aurora kinase inhibitor
SNS-314	2.12	3	3	2.12	Aurora kinase inhibitor
AMG900	2.12	3	3	2.12	Aurora kinase inhibitor
CCT137690	2.12	3	3	2.12	Aurora kinase inhibitor
MK-5108	2.12	3	3	2.12	Aurora kinase inhibitor
BI-847325	2.03	5	5	2.03	Aurora kinase inhibitor MEK inhibitor

What the transcriptome does and does not contribute here. Signature v2 sharpens a direction for genes whose mRNA is measurable in nasal mucosa — HRH1 ($z = 3.09$), CCR3 (2.82), IL4R (1.65), BTK (1.13) — and it is silent about the cytokines that several of the biologics below actually neutralise: IL5 (0.77) and IL13 (1.20) are measured only weakly, TSLP is effectively flat (0.06), and IgE is not on the array at all. A whole-mucosa bulk measurement also cannot report the abundance of the cells that matter for those drugs (eosinophils, mast cells). That is why section 4 uses the registry as primary evidence and the transcriptome only as a direction check.

3 Route B — the registered trial landscape and its gaps

We pulled all interventional studies registered for allergic rhinitis and for seven adjacent type-2 / inflammatory diseases via the ClinicalTrials.gov v2 API, parsing every intervention arm and mapping each free-text arm name to canonical drug identities (brand names resolved to generics, combination arms credited to each active, formulation and dose words discarded).

Condition (ClinicalTrials.gov query .cond)	Studies pulled
allergic rhinitis	1,276
asthma	5,299
atopic dermatitis	1,786
chronic rhinosinusitis	542
nasal polyps	332
eosinophilic esophagitis	250
chronic urticaria	280
allergic conjunctivitis	242
total	10,007

From 9,486 drug-like arms we obtain 6,939 canonical mentions of 1,999 distinct substances across 4,441 studies. The rhinitis landscape is narrow and repetitive:

Rank	Intervention (canonical)	Registered rhinitis studies	Phases
1	mometasone	74	e1 / 1 / 2 / 3 / 4 / NA
2	azelastine	50	1 / 2 / 3 / 4 / NA
3	montelukast	44	e1 / 1 / 2 / 3 / 4 / NA
4	fluticasone propionate	40	1 / 2 / 3 / 4 / NA
5	cetirizine	35	1 / 2 / 3 / 4 / NA
6	levocetirizine	34	1 / 2 / 3 / 4 / NA
7	desloratadine	33	2 / 3 / 4
8	ciclesonide	31	1 / 2 / 3 / 4
9	loratadine	28	2 / 3 / 4 / NA
10	fexofenadine	27	2 / 3 / 4
11	fluticasone furoate	21	2 / 3 / 4 / NA
12	budesonide	20	e1 / 2 / 3 / 4 / NA

Against that, 1,535 substances have registered studies in the adjacent diseases and none in rhinitis; 407 of them have at least two adjacent-disease studies, 102 have five or more, and 34 carry a mechanism we could annotate by hand. Of those 34, 22 have reached phase 3 somewhere else — and reaching phase 3 is a milestone, not an outcome, so section 3.1 audits what those phase-3 programmes actually reported. The single most striking asymmetry is in the anti-eosinophil and anti-IL-13 monoclonals, several of which have dozens of registered studies across six adjacent diseases and not one in rhinitis:

Intervention	Mechanism	Class	Adjacent studies (condition query)	Distinct diseases	Phase 3 reached	Example NCT
mepolizumab	IL5	anti-IL-5 mAb	109	6	yes	NCT00244686
benralizumab	IL5RA	anti-IL-5Ralpha mAb	70	6	yes	NCT00768079
lebrikizumab	IL13	anti-IL-13 mAb	39	2	yes	NCT00781443
tralokinumab	IL13	anti-IL-13 mAb	29	2	yes	NCT01402986
crisaborole	PDE4	PDE4 inhibitor	27	1	yes	NCT03233529
upadacitinib	JAK1	JAK1-selective inhibitor	23	2	yes	NCT02925117
ruxolitinib	JAK1/JAK2	JAK1/2 inhibitor	22	1	yes	NCT03011892
reslizumab	IL5	anti-IL-5 mAb	21	4	yes	NCT00587288
roflumilast	PDE4	PDE4 inhibitor	21	2	yes	NCT00073177
abrocitinib	JAK1	JAK1-selective inhibitor	21	1	yes	NCT03358693
depemokimab	IL5	anti-IL-5 mAb (long-acting)	17	3	yes	NCT04718389
nemolizumab	IL31RA	anti-IL-31RA mAb	15	1	yes	NCT01986933




Figure 20. The registered trial landscape. (A) What rhinitis has already tried: a short list dominated by intranasal glucocorticoids, antihistamines, a leukotriene antagonist and allergen immunotherapy. (B) Biologics and small molecules with registered studies in adjacent type-2 / inflammatory diseases and zero registered rhinitis studies; each drug's mechanism is in brackets and the number of distinct adjacent diseases follows the study count. (C) Pull depth per condition. Rhinitis is not under-studied in absolute terms (1,276 registered studies) — the gap is in which mechanisms get tested, not in how many trials are run.

Only three biologics have ever been registered in rhinitis — omalizumab, dupilumab and tezepelumab — while other antibodies against the same type-2 axis have not. That contrast is what makes the absences above informative rather than a blanket statement about biologics: the barrier is not "biologics in rhinitis", it is which mechanism is

tried.

Biologic	Adjacent-disease precedent	Registered rhinitis studies
omalizumab	yes	10
dupilumab	yes	4
tezepelumab	yes	2
benralizumab	yes	0
ligelizumab	yes	0
mepolizumab	yes	0
reslizumab	yes	0

3.1 Was the adjacent-disease precedent actually positive? A programme-outcome audit

A count of registered studies measures how hard a mechanism was pushed; it says nothing about whether it worked. To separate the two, the 34 mechanism-annotated gaps were re-queried on the same API *by intervention name* — a second counting convention, not a correction of the first. Section 3 counts studies per condition (query.cond, one study per condition it lists); this audit counts studies per drug (query.intr) and then restricts them to the eight conditions of this study. The two conventions disagree substantially and neither is authoritative: benralizumab has 70 adjacent-disease studies by condition query and 100 by intervention query; the respective figures are 3 and 15 for fevipiprant and 6 and 19 for remibrutinib. Both are reported in the table rather than merged into one number.

For every phase-3 study of these candidates that posted results on ClinicalTrials.gov (hasResults), the p-values of the *primary* outcome analyses were parsed from the results module and reduced to one mechanical statement per study: significant (any primary analysis $p < 0.05$), not significant (all reported primary $p \geq 0.05$), or no numeric p. Over the whole set this yields 220 phase-3 studies, 146 with posted results, and 78 significant vs 20 not significant primary analyses, with 48 studies posting no numeric p. Reducing each candidate to a single verdict gives 11 positive, 7 mixed, 2 negative and 14 with no posted phase-3 p at all.

Intervention	Adj. studies cond. / intr.	Phase-3 studies	Posted results	Primary p significant	Not significant	No numeric p	Verdict
benralizumab	70 / 100	28	21	9	3	9	mixed
lebrikizumab	39 / 44	26	16	6	1	9	mixed
mepolizumab	109 / 108	17	14	7	1	6	mixed
tralokinumab	29 / 39	15	13	6	3	4	mixed
reslizumab	21 / 24	13	10	4	3	3	mixed
roflumilast	21 / 29	13	5	4	0	1	positive
upadacitinib	23 / 42	12	6	5	0	1	positive
abrocitinib	21 / 35	11	7	6	0	1	positive
remibrutinib	6 / 19	10	4	2	0	2	positive
depemokimab	17 / 14	9	6	4	1	1	mixed
baricitinib	13 / 20	8	8	6	0	2	positive
crisaborole	27 / 34	8	6	4	1	1	mixed
rocatinlimab	13 / 13	8	6	4	0	2	positive
ruxolitinib	22 / 25	8	5	4	0	1	positive
fevipiprant	3 / 15	7	7	0	6	1	negative
amlitelimab	12 / 12	6	0	0	0	0	no phase-3 p reported
ligelizumab	9 / 15	5	5	2	0	3	positive
nemolizumab	15 / 16	5	2	2	0	0	positive
itepekimab	10 / 10	4	0	0	0	0	no phase-3 p reported
tapinarof	6 / 9	4	3	2	0	1	positive
etrasimod	3 / 3	1	1	0	1	0	negative
lirentelimab	4 / 3	1	1	1	0	0	positive
ustekinumab	3 / 5	1	0	0	0	0	no phase-3 p reported

The counterexample that motivated this section. The DP2/CRT2 antagonist **fevipiprant** is, by registration count, an unremarkable entry on the gap list (3 adjacent-disease studies by condition query, 15 by intervention query) — yet it has 7 phase-3 studies, 7 of them with posted results, and **none** with a significant primary endpoint: the informative posted p-values run from 0.059 to 0.665 (1 further study posted no numeric p), and 26.7% of its adjacent-disease studies were stopped early. A decade ago the same registry view would have placed oral DP2 antagonism near the top of this list. Its presence here is the reason the outcome layer exists: registration depth and phase-3 reach are prioritisation inputs, not evidence, and a candidate list built from them alone would rank a failed in-class programme beside a successful one. Which is exactly the difference the audit exposes — anti-IL-5 (mepolizumab 7 significant vs 1 not significant among informative posted endpoints; benralizumab 9 vs 3) versus DP2 (0 vs 6).

The verdict column is coarse by construction and must be read with section 5: a "positive" verdict is uncorrected for multiplicity across a study's several primary endpoints, "no numeric p" is dominated by studies whose primary endpoint is pharmacokinetic or safety rather than efficacy, and a study-level result is not a programme-level result — a drug can fail one indication and succeed in another, which is what the mixed rows are.

3.2 The nasal endpoint record: what rhinitis itself has already tested and rejected

Sections 3 and 3.1 read the trial record around rhinitis. The mirror question is what the nose has already answered. Of the 1,276 registered rhinitis studies, 269 posted results on ClinicalTrials.gov (21%); 179 of those are phase 3 or 4. Applying the identical primary-endpoint parser: 93 reached a significant primary endpoint, 38 did not, and 138 posted results without a numeric primary p. 147 distinct substances have at least one posted-result rhinitis study and 41 have two or more. The top of that table is exactly the narrow field of section 3, now with outcomes attached:

Intervention	Posted studies	Phase 3+	Significant	Not significant	No p
mometasone	33	29	21	5	7
fluticasone furoate	13	12	6	3	4
azelastine	20	20	6	1	13
beclomethasone	8	7	6	0	2
ciclesonide	15	14	5	2	8
olopatadine	15	12	5	1	9
fluticasone propionate	18	16	4	3	11
loratadine	16	15	4	1	11
aerosol	5	5	4	0	1
levocetirizine	10	9	3	3	4
montelukast	17	17	3	2	12
desloratadine	8	4	3	1	4




Figure 22. What rhinitis itself has tested. (A) Interventions with at least two posted-result rhinitis studies, split by the verdict of the primary endpoint (green = significant, red = not significant, grey = results posted without a numeric p); the top of the list is the narrow field of figure 20 with outcomes attached. (B) All 269 rhinitis studies that posted results: 93 reached a significant primary endpoint, 38 did not, and 138 posted none. A few entries are canonicalisation artefacts rather than drug entities ('aerosol', 'allergen', 'vitamin', 'comparator') and are shown for transparency, not as verdicts.

The negative side of the same table matters more for prioritisation, because a candidate list that re-proposes what the end organ has already rejected is worse than useless. 20 substances have posted non-significant rhinitis results and no significant one. Most of them are allergen-extract product labels with one or two studies each ("timothy grass", "pangramin SLIT", "2lalerger") and are not interpretable as drug verdicts; the clean, drug-like entries are what the layer is for. **Phenylephrine** has 4 posted rhinitis studies, 2 of them non-significant and none significant (largest reported $p = 0.561$; the remainder posted no numeric p) — this is the face-validity check for the whole outcome layer, since oral phenylephrine's decongestant efficacy is precisely the claim that has collapsed. The histamine-H3 antagonist **PF-03654746**, tested on top of fexofenadine, is the more interesting one for this paper: a *novel* mechanism that was actually tried in rhinitis, phase 2, and did not separate from its comparator ($p = 0.097$, NCT00562120). And **esomeprazole** (NCT00745849, $p = 0.450$) is the reflux-rhinitis hypothesis, also already tested and not confirmed.

Two further observations from the same table. First, the biologics that did get into rhinitis have thin posted records: omalizumab has 4 posted studies (1 significant, 3 without a numeric p) and dupilumab 2 — of which 1 is non-significant — so even for the one biologic the field has adopted, the registered evidence base in the nose is a handful of studies. Second, allergen immunotherapy is the only non-pharmacological class with volume here (22 of the 269 posted-result studies involve immunotherapy arms) — which is why the gap in section 3 is a gap in *mechanisms*, not in trial activity.

3.3 A non-type-2 direction with phase-3 data nobody can read: the immunomodulator gap

An earlier question in this series asked about a thymic/spleen-peptide immunomodulator for rhinitis. The registry answer is blunt, and it is worth stating with the counts: by intervention name there are **0** studies for "spleen aminopeptide", 4 for "spleen polypeptide" (none in rhinitis), 1 for pidotimod, 10 for thymopentin and 65 for thymalfasin — and **zero** of any of them registered in rhinitis. A zero here means the compound is absent from this registry, not that the class is inert; but it does mean that no registered conclusion about the nose can be drawn in either direction. The literature layer of section 3.4 shows how far even that statement has to be qualified: **2 randomised trials of spleen aminopeptide in allergic rhinitis** are indexed in PubMed and 2 for pidotimod (which has 1 registered study in total, in another indication), while the registered rhinitis count for both is 0.

The one immunomodulator family that *did* reach phase 3 in rhinitis is a sublingual bacterial lysate (Ismigen, "polyvalent mechanical bacterial lysate"), and its record adds a third kind of gap to the two above: not "never tried" and not "tried and failed", but **silent in the registry**. 3 of 6 registered lysate studies are in rhinitis; all 3 are phase 3, all are completed, all three carry symptom-based primary endpoints (total nasal symptom score, peak nasal inspiratory flow, VAS) — and **3 of 3 never posted results**:

NCT	Phase	Status	Enrolment	Symptom primary endpoints / all primaries	Results posted
NCT04270552	3	COMPLETED	76	5 / 5	no
NCT04637425	3	COMPLETED	70	1 / 1	no
NCT04802616	3	COMPLETED	80	3 / 8	no

This is a mechanistically distinct direction — innate-immune stimulation rather than type-2 cytokine blockade — and it already has completed phase-3, placebo-controlled, symptom-endpoint trials in seasonal rhinitis at $n = 76$ to 80. What the registry cannot supply is the direction of the result. Three caveats travel with the finding: the trials come from a single academic sponsor, the enrolment is small by phase-3 standards, and a fourth lysate trial in the same programme used a microbiological primary endpoint (nasal *S. aureus* carriage) rather than a symptom score, so the programme is a mixture of mechanistic and efficacy designs. As a research target it is unusually concrete: the data exist, they were never published to the registry, and the class sits outside the type-2 axis that occupies the rest of this paper. For completeness, the same sweep shows the lysate class under other names (OM-85, Broncho-Vaxom) has 0 rhinitis studies by that name, so the entire registered immunomodulator presence in rhinitis is these three trials.

The registry silence, however, is not the whole record. Querying PubMed for the three NCT numbers returns a linked publication for one of them: NCT04802616 is the trial reported in *Front Immunol* 2023 (PMID 36733480), a randomised, double-blind, placebo-controlled study in 50 children with grass-pollen seasonal rhinitis. In that report the lysate **did** separate from placebo on the symptom scores — TNSS significantly lower vs placebo ($p=0.01$ TC1; $p=0.009$ V2); VAS significantly lower ($p=0.04$ TC1; $p=0.04$ V2; $p=0.03$ TC2) — while peak nasal inspiratory flow showed no significant difference and the measured immune subsets showed no significant change ($\gamma\delta T$, iNKT, CTL). The second linked record is an immune-subset analysis of the same trial (PMID 42512095, 2026: Th1/Th10/Treg up, Th2/Th17 down). The other two trials (NCT04270552, NCT04637425) return no publication at all. Two qualifications keep this in proportion: the analysed set (50 children) is smaller than the registered enrolment ($n = 80$), and the trial remains a single-centre, single-sponsor study in one pollen season. The corrected statement of this gap is therefore narrower and more interesting than either version alone: **one of the three registry-silent phase-3 trials is readable, its symptom endpoints went in the sponsor's favour, and it was still never posted to the registry** — which is the cleanest illustration in this paper that the registry is a record of what sponsors filed, not of what has been tested.

3.4 Does "zero registered" mean "never studied"? A literature cross-check

Every statement so far has been about the registry. The natural objection is that a mechanism can be studied without being registered, so we ran the complementary search: PubMed for each of the 34 candidates, first under a deliberately broad nasal query (rhinitis OR nasal OR nose) and then under strict allergic-rhinitis phrases ("allergic rhinitis" OR seasonal OR perennial OR rhinoconjunctivitis OR "hay fever"). The broad query returns hits for 30 of 34 candidates — including 4 with none — but inspection shows it is not usable as an evidence count: its top hits are asthma and chronic-rhinosinusitis reviews and case reports that merely mention the nose (the largest,

for mepolizumab, returns several hundred records). The strict query is the honest one: **22 of 34 candidates have at least one strict-phrase publication**, and 12 have none at all.

Because that still does not distinguish "studied in rhinitis" from "mentioned in a review", every retrieved record was read at title level (80 records, 69 unique documents across the 22 candidates; figure 23A). The partition is the result of this section:

Candidate	Human AR endpoint	AR as comorbidity stratum	Preclinical AR model	Human AR, no drug
roflumilast	1	—	1	—
upadacitinib	1	—	—	—
tralokinumab	—	2	1	—
fevipiprant	—	—	1	1
apremilast	—	—	1	—
lebrikizumab	—	—	1	—
infliximab	—	—	1	—
ruxolitinib	—	—	1	—
tofacitinib	—	—	1	—
benralizumab	—	1	—	—
mepolizumab	—	1	—	—
crisaborole	—	1	—	—
abrocitinib	—	1	—	—

Only 2 records are human studies in which the candidate was given and rhinitis symptoms were an endpoint — roflumilast (an AR-first randomised trial) and upadacitinib (a comorbidity analysis in atopic dermatitis, not an AR-first trial). 6 more records are human studies of the drug in another disease that report rhinitis as a comorbidity or effect modifier, 8 are animal, in-vitro or in-silico work on nasal tissue, and 1 is a human rhinitis study in which the candidate was not the intervention. The remaining 63 records (38 reviews or unrelated topics, 25 human studies in other diseases) carry no rhinitis evidence at all. In other words: **the registry gap is not a literature gap, but it is very nearly an evidence gap** — the mechanism-level candidates are mentioned in the rhinitis literature mainly through reviews and comorbidity analyses, not through trials.

The one exception deserves to be stated with its numbers, because it is the strongest single item in this paper's evidence base. **Roflumilast**, an oral PDE4 inhibitor, was tested in allergic rhinitis in a randomised, double-blind, placebo-controlled crossover study of 25 subjects (oral 500 µg/day for 9 days, with daily intranasal allergen challenge in AR subjects); it was published in *J Allergy Clin Immunol* in 2001 (PMID 11590377, Clinical Trial; Randomized Controlled Trial) and the reported outcomes were favourable: nasal airflow (rhinomanometry) higher than placebo on day 9; nasal itching and rhinorrhoea significantly reduced; subjective congestion significantly reduced within 4 days. That trial is **not registered on ClinicalTrials.gov** (roflumilast has 0 registered rhinitis studies), its adjacent-disease programme reads positive, and its target is up in the v2 nasal signature (PDE4B z = 1.387). The same limitations that attach to any single study apply, and the section 4 entry carries them: n = 25 crossover challenge model, not a field trial; 9-day acute dosing, not chronic therapy; oral PDE4 inhibition carries GI and weight-loss adverse effects.

The same check revises the immunomodulator story in the opposite direction. Spleen aminopeptide has **0 registered rhinitis studies** but 2 randomised rhinitis trials indexed in PubMed (PMID 39524381, 39142892), pidotimod has 1 registered study in total (in another indication) against 2 published rhinitis studies, and the OM-85 name returns 2 human rhinitis studies in the literature while registering no rhinitis study under that name (figure 23B). The two blind spots are therefore complementary rather than merely absent data: **trials that were registered and never posted (the lysates, section 3.3) and studies that were published and never registered (roflumilast; spleen aminopeptide)**. Neither kind is visible from a registry count, and both are visible from the two-source check used here.




Figure 23. The literature cross-check. (A) Every PubMed record retrieved for the 34 gap candidates under strict allergic-rhinitis phrases, classified by reading its title (80 records, 69 unique documents). Green = human study in which the candidate was given and rhinitis was an endpoint (2 records, 2 candidates); orange = preclinical or in-silico only; blue = human study reporting rhinitis as a comorbidity; light grey and grey = human studies in other diseases and reviews; the 12 rows with no bar at all have no strict-phrase rhinitis publication. Most candidates are represented by reviews and comorbidity analyses, not by rhinitis trials. (B) The immunomodulator side of the same check: registered rhinitis trials (red) vs published human rhinitis studies (green) for spleen aminopeptide, pidotimod, the oral lysate OM-85 and the sublingual lysate Ismigen — registry silence and published evidence are dissociated in both directions. Counts are title-level, hand-curated, and are lower bounds on the true literature.

3.5 The same question asked backwards: published in rhinitis, never registered

Section 3.4 checked the candidate list one drug at a time. The same check can be run at scale, in the opposite direction, and it answers the question a reader of any such list actually has: **which drugs have published human rhinitis trials and yet appear in none of the 849 registered rhinitis studies?** We retrieved every PubMed record indexed as a trial in allergic rhinitis (3,391 records), took its title, and matched titles against the repurposing-hub drug lexicon plus a supplement for compounds the hub does not carry (6,565 names). 1,883 of the 3,391 titles name at least one lexicon drug; after merging the synonyms and salt forms that the registry files under one canonical name,

144 distinct labels matched, of which 63 also appear in a registered rhinitis study and **81 do not** — 79 drugs plus 2 non-drug labels (acupuncture, and substance P, which appears in rhinitis studies as a measured mediator rather than as an intervention).

A title match alone over-counts, because a drug's name also appears in programme papers — dose-finding in healthy volunteers, pharmacokinetics, disposition — that are not rhinitis trials in any population sense (66 such titles across 32 labels; setipirant has three AR-trial titles and no AR-population study among them). Excluding them leaves 376 trial titles for the 79 literature-only drugs; 35 of those drugs carry at least two. All 35 were then read by hand (every matched title is listed in `results/lit_inverse_title_audit.csv`) and split into those whose hits are AR-population trials (25 drugs, 314 trial publications), those whose hits are comorbid or mechanistic rather than AR-endpoint work (5: cimetidine, deflazacort, furosemide, lodoxamide, terbutaline), and those whose hits belong to a different disease altogether (5: amoxicillin, azithromycin, clarithromycin, clavulanate, transfer factor — for the antibiotics the rhinitis term sits in the abstract of a rhinosinusitis trial). Figure 24A shows the largest members of the first group; figure 24B the composition of the audited set.

Drug (lit-only, hand-audited as AR trials)	Published AR trials	Years	Status today	Mechanism of action	v2 screen S
terfenadine	94	1980-2004	withdrawn	histamine receptor antagonist	1.593
sodium cromoglycate	41	1973-2015	active	—	n/a
astemizole	38	1983-1998	withdrawn	histamine receptor antagonist	1.593
ebastine	23	1991-2009	regional	histamine receptor antagonist	1.706
mizolastine	15	1996-2007	regional	histamine receptor antagonist	3.091
nedocromil	14	1986-2006	active	histamine receptor antagonist	0.581
clemastine	13	1975-2003	active	histamine receptor antagonist	3.091
acrivastine	10	1988-1996	active	histamine receptor antagonist	3.091
oxatamide	8	1980-2002	regional	histamine receptor antagonist	3.091
ipratropium	8	1986-1998	active	acetylcholine receptor antagonist	0.106
triprolidine	7	1975-2021	active	histamine receptor antagonist	3.091
pranlukast	5	2008-2017	regional	leukotriene receptor antagonist	0.606

The point of the sweep is **not** that these are hidden drugs, and reading it that way would repeat the error this paper argues against. Nearly every name at the top is old, and the reason it has no registered rhinitis trial is that there was no registry to register in: the rhinitis trial literature in this set begins in the 1960s and 1970s, while routine trial registration — and the posting of results — only became general after the mid-2000s. The audit column makes the same point from the drug's side: the 35 hand-audited drugs, 21 are still marketed (several over the counter), 7 are marketed only in some regions (ebastine and mizolastine without US approval; pranlukast and suplastat in Japan), 3 were withdrawn for cardiotoxicity (terfenadine, astemizole, fenoterol), 1 is restricted for hepatotoxicity (nimesulide), 1 has left the market (azatadine) and 1 stopped in development (the TRPV1 antagonist SB-705498, whose two 2013–2014 trials are the most recent entries in the list). Roflumilast's 2001 trial, the strongest single item in section 3.4, is therefore a normal member of a large class rather than an anomaly — "published but never registered" is what the pre-registration era looks like from inside a registry.

Read that way, the inverse sweep calibrates the whole paper rather than adding candidates to it. The 1,535 substances of section 3 have zero *registered* rhinitis studies, and the 79 drugs here have zero registered rhinitis studies as well: the first group is "not registered yet", the second is mostly "registered too late to be recorded", and only the publication record separates them. The one check the sweep does supply runs against the screen rather than for it: among the literature-only drugs that are AR-population trials and carry a v2 score, the highest score in this group is 3.09, attained by 7 drugs (mizolastine, clemastine, acrivastine, oxatamide, triprolidine, azatadine, dimetindene), i.e. by old H1 antihistamines that are already standard rhinitis therapy (their hub annotations are cholinergic receptor antagonist / histamine receptor antagonist). The top of this group is therefore occupied by established pharmacology rather than by novelty, which is the same lesson as the known-drug calibration of section 2. A high score reproduces established pharmacology; it does not signal newness, and it does not make the 79 drugs of this section candidates.




Figure 24. The inverse sweep: the rhinitis literature contains drugs the rhinitis registry does not. (A) The largest literature-only drugs — those with published AR trial titles and zero registered rhinitis studies — by hand-audited AR-population trial count, coloured by where the compound stands as a product today (3 withdrawn for cardiotoxicity, 7 marketed regionally only, 21 still marketed, plus discontinued, restricted and development-stopped members). The largest single entry, terfenadine, has 94 published AR trial titles and no registered rhinitis study; nothing here is hidden, and the pre-registration era explains the pattern. (B) What the 35 hand-audited literature-only drugs with at least two population trials turn out to be: top bar, product status today; bottom bar, whether their hits are AR-population trials (25), comorbid or mechanistic work (5) or another disease altogether (5). Counts are title-level and are lower bounds; the sweep is limited to records MEDLINE indexes with a trial publication type.

4 Cross-hypotheses: precedent × nasal transcriptome direction

For each of the 34 mechanism-annotated gaps we ask whether the nasal transcriptome contradicts it. Every mechanism in this set is an inhibitor, antagonist or blocking antibody of its target, so the concordant configuration is a target that is *up* in rhinitis. Figure 21 positions all 34 candidates by that criterion; 15 are concordant, 2 mixed, 13 discordant and 4 involve a target that is not detectable on the nasal array (a "not measurable" verdict, which we deliberately do not treat as evidence against). 21 of 34 are

therefore not contradicted by the transcriptome. The programme-outcome audit of section 3.1 is carried alongside: the 2 candidates with a negative posted phase-3 record (etrasimod, fevipirant) are shown for completeness and are not recommendations, and 14 candidates have no posted phase-3 primary endpoint at all, so their precedent rests on phase-2 or surrogate data.

Hypothesis	Target	z (v2 nasal)	Adjacent studies (condition)	Diseases	Phase 3	Transcriptome
mepolizumab	IL5	0.77	109	6	yes	concordant
benralizumab	IL5RA	-2.24	70	6	yes	discordant
lebrikizumab	IL13	1.20	39	2	yes	concordant
tralokinumab	IL13	1.20	29	2	yes	concordant
crisaborole	PDE4	n.d.	27	1	yes	not measurable
upadacitinib	JAK1	-0.90	23	2	yes	discordant
ruxolitinib	JAK1/JAK2	-0.90	22	1	yes	mixed
reslizumab	IL5	0.77	21	4	yes	concordant
roflumilast	PDE4	n.d.	21	2	yes	not measurable
abrocitinib	JAK1	-0.90	21	1	yes	discordant
depemokimab	IL5	0.77	17	3	yes	concordant
nemolizumab	IL31RA	2.10	15	1	yes	concordant




Figure 21. Cross-hypotheses. Each row is a drug whose mechanism is registered in at least one adjacent type-2 / inflammatory disease and which has no registered rhinitis study. Dot size is the number of adjacent-disease studies (condition query); colour is whether the nasal transcriptome supports the direction (concordant = the target is up in rhinitis, so blocking it is the therapeutically expected direction). Stars mark mechanisms that reached phase 3 elsewhere — reach, not outcome. The two candidates whose posted phase-3 primary endpoints are all non-significant (section 3.1) are ringed in red and retained for completeness only. 'n.d.' means the target is not present on the nasal array — mostly cytokines and IgE, which are exactly the molecules the strongest biologics neutralise. This displays evidence, not expected effect size.

4.1 The hypotheses we would act on first

Ranked by clinical precedent and filtered on not being contradicted by the nasal transcriptome, the leading directions are — with the only candidate that already has direct human rhinitis evidence placed first:

- PDE4 inhibition (roflumilast).** The one entry in this set for which the question "does blocking this target do anything in the human nose?" has already been answered once: an oral PDE4 inhibitor was tested in allergic rhinitis in a randomised, double-blind, placebo-controlled crossover trial (25 subjects, 500 µg/day for 9 days against daily nasal allergen challenge) and the reported symptom and airflow outcomes were favourable — nasal airflow (rhinomanometry) higher than placebo on day 9; nasal itching and rhinorrhoea significantly reduced; subjective congestion significantly reduced within 4 days. That trial (PMID 11590377) is **not in the registry**, so a registration-count view would have called this mechanism untested; PDE4B is up in the v2 signature ($z = 1.387$) and the adjacent-disease programme verdict is positive. It ranks first on evidence, not on novelty, and the caveats are explicit: $n = 25$ crossover challenge model, not a field trial; 9-day acute dosing, not chronic therapy; oral PDE4 inhibition carries GI and weight-loss adverse effects. The open questions it leaves are therefore formulation and chronic-exposure questions rather than questions about whether the mechanism reaches the nose.

2. **Anti-IL-5 / anti-IL-5R α (mepolizumab, benralizumab, depemokimab, reslizumab).** Neither mepolizumab nor benralizumab has a single registered rhinitis study, against 109 and 70 studies in up to six adjacent diseases including chronic rhinosinusitis and nasal polyps, where they reached phase 3. This is the one mechanism whose adjacent-disease record is strong on outcomes as well as volume: among informative posted phase-3 endpoints, 7 of 8 are significant for mepolizumab and 9 of 12 for benralizumab. Tissue eosinophilia is a defining feature of the severe rhinitis endotype, and the transcriptome is consistent (CCR3 $z = 2.82$). The discordant flag on benralizumab comes solely from IL5RA mRNA ($z = -2.24$), which in bulk mucosa reflects epithelial rather than eosinophil expression and should not be read as biological opposition.
3. **Anti-IL-13 (lebrizumab, tralokinumab; and the IL-4R α axis already used in rhinitis via dupilumab).** IL-4R α ($z = 1.65$) is up in the v2 signature and dupilumab already has 4 registered rhinitis studies, so the pathway is not in doubt — and the posted phase-3 record for selective IL-13 blockade is more positive than not (lebrizumab 6 significant vs 1 not; tralokinumab 6 vs 3), though both are mixed rows because they failed in some indications. What is untested is whether selective IL-13 blockade adds anything in rhinitis or in rhinitis with conjunctival involvement.
4. **BTK inhibition (remibrutinib, rilzabrutinib).** The only mechanism in this set that targets the IgE/mast-cell axis with a small molecule; its phase-3 chronic-urticaria programme posted significant primary endpoints (remibrutinib 2 significant vs 0 not significant among 4 posted studies), BTK is up in the nasal signature ($z = 1.13$) and the drugs are oral. For patients who cannot take biologics this is the most actionable gap.
5. **Anti-OX40L / anti-OX40 (amlitelimab, rocatinlimab).** TNFRSF4 (OX40) is up ($z = 1.07$), the target is a T-cell costimulatory axis that sits upstream of cytokine production, and the mechanism is in late-phase development in atopic dermatitis with early asthma data — a mechanistically distinct bet from another type-2 cytokine. The outcome layer splits the two drugs and this should be stated plainly: rocatinlimab has a positive posted record (4 significant vs 0 not significant), whereas amlitelimab has 6 phase-3 studies registered and **none** with posted results, so its precedent currently rests on phase-2 data.
6. **JAK1 inhibition (upadacitinib, abrocitinib).** The largest precedent in atopic dermatitis and eosinophilic oesophagitis after the type-2 antibodies, oral, and upstream of the same IL-4/IL-13 signal that dupilumab blocks; its posted phase-3 endpoints are uniformly significant (upadacitinib 5 vs 0, abrocitinib 6 vs 0). The transcriptome is unhelpful here: JAK1 mRNA is down ($z = -0.90$), which for a kinase is uninformative about signalling activity.
7. **Next-generation anti-IgE (ligelizumab).** Omalizumab is the only anti-IgE registered in rhinitis (10 studies); a higher-affinity successor has been developed in chronic urticaria but never registered in rhinitis. Its phase-3 urticaria record is positive but thin: 2 significant and 0 not significant among informative posted endpoints, with 3 of 5 posted studies reporting no numeric p and 33.3% attrition. IgE itself is not on the array, so no transcriptomic claim is made.

What this list is not. None of these entries is evidence of efficacy in rhinitis. Each is a mechanism with (i) a documented clinical precedent in a disease that shares type-2 biology, (ii) no registered rhinitis study, and (iii) no transcriptomic contradiction. The section 3.1 verdicts do not change this reading in either direction: a "positive" verdict means that the posted primary analyses reached $p < 0.05$, not that the mechanism would work in the nose, and a "negative" one (as for fexipirant) means the opposite. The literature layer of section 3.4 sharpens the same warning from the other side: of 34 candidates, exactly 2 (roflumilast, upadacitinib) has a human rhinitis study at all, and the roflumilast trial of item 1 is a 25-year-old challenge-model study of 25 subjects. Two further caveats travel with the list. The mechanism labels are hand-curated, not derived from a database, because biologics are absent from the small molecule target table used in paper 3. And several gap candidates are dermatological or inhaled formulations that are not administrable to the nose — the route column of `results/trials_gap_candidates.csv` records this, and those entries should not be read as nasal-relevant regardless of their precedent.

5 Limitations

- The registry covers **registered** interventional studies only, and is weighted toward industry and US/EU academic sponsors. Chinese-language registries are largely absent from this API. "Zero rhinitis studies" therefore means "not in this record", and at least one known historical example — a CRTh2 antagonist developed with a nasal-challenge design — does not appear as a rhinitis study here.
- Arm names are free text. Canonicalisation is rule-based (dose and formulation words dropped, brand names resolved, combination arms credited to each active) and was validated by inspection of the top of the frontier and gap tables, not exhaustively. Substances that are not single drug entities (allergen extracts, "house dust") appear in the frontier and are not gap candidates.
- Mechanism annotations for biologics are curated by hand; the target-direction check can only be evaluated for targets present on the array, and absence from a bulk nasal array is not evidence of irrelevance.
- Signature v2 is an improvement in transferability, not a repaired signature. It is anti-correlated with one held-out cohort (GSE329964), and the screen it feeds still produces no FDR-significant drug.
- Clinical precedent in an adjacent disease does not transfer automatically: asthma, atopic dermatitis and chronic urticaria differ from rhinitis in target tissue, exposure and endotype mix, and a mechanism that works elsewhere can fail in the nose.
- The programme-outcome layer (section 3.1) is bounded by what sponsors post. Primary p-values exist only where a sponsor filed results; a single study can have several primary endpoints that are not corrected for multiplicity here; and the "no numeric p" bucket is dominated by studies whose primary endpoint is pharmacokinetic or safety rather than efficacy. A study-level p-value is also not a programme-level verdict — a drug can fail one indication and succeed in another, which is why 7 candidates are classed mixed. The "significant" rule is the mechanical one (any primary analysis $p < 0.05$) and was not tuned.
- Two counting conventions are used and they disagree. Section 3 counts studies per condition (`query.cond`), section 3.1 counts studies per intervention name (`query.intr`) restricted afterwards to the eight conditions; benralizumab is 70 vs 100 and fexipirant 3 vs 15 under the two. Neither is a complete count of what was run, and both are reported rather than reconciled. The adjacent-condition assignment itself is keyword-based, so broader terms can pull in related but non-identical entries (any "dermatitis", any "sinusitis"), which inflates rather than deflates the gap list.
- The end-organ outcome layer (section 3.2) inherits every limit of section 3.1 and adds its own: only 21% of registered rhinitis studies posted results, so the record is a biased sample of what was run (positive trials are likelier to be posted); a per-intervention verdict requires the drug to be a single entity, and the tables retain non-drug labels ("aerosol", "allergen", "vitamin", allergen-extract product names) whose counts must not be read as drug verdicts; and combination arms are credited to every active ingredient, so a study of an antihistamine-plus-decongestant combination appears under both names. Section 3.3 reports the existence and shape of unposted phase-3 trials: for two of the three lysate trials the direction is genuinely unknown to us and is not inferred anywhere in this paper, and the one direction that is reported (NCT04802616) comes from its peer-reviewed publication, cited as such, not from a registry result.
- The literature layer (section 3.4) is a title-level search, not a systematic review. It retrieved records for 80 candidate-by-drug hits and 11 of those were the same document retrieved for more than one candidate; the strict phrase set will miss studies that describe rhinitis without using one of its standard names; the classification was done by reading titles (and abstracts for the records that carry claims) rather than full texts; and the counts are therefore lower bounds on the true literature. Absence of a publication for a candidate is not evidence that none exists. The two headline contrasts it produces — roflumilast published but never registered, spleen aminopeptide published twice with zero registered trials — were verified against the abstracts and the registry probe respectively.
- The inverse sweep (section 3.5) inherits the title-level limits of section 3.4 and adds two. It can only see records MEDLINE indexes with a trial publication type, so it is a lower bound in both directions: the forward search of section 3.4 found 2 randomised spleen-aminopeptide rhinitis trials, while this sweep sees 1 of them. And its lexicon is the repurposing hub plus a hand-written supplement, so a drug the hub does not carry and the supplement omits cannot be found by construction. The audit verdicts are typed by hand from titles; the status column ("withdrawn", "regional", "still marketed") is a manual reading of the public record and was *not* checked against a regulatory database in this work, so those labels should be read as annotations and not as regulatory facts.
- No new experimental data were generated; this is a secondary analysis of public transcriptomes and registry records.

6 Conclusion

Asking "what should be tried next in rhinitis" is more productive through the trial registry than through a non-replicating expression signature. The registry shows a field concentrated on a dozen topical and oral agents while 407 substances with at least two registered studies in adjacent type-2 diseases — including anti-IL-5, anti-IL-13, BTK and OX40-pathway agents, several at phase 3 — have never been registered in rhinitis.

A second pass through the same registry turns those counts into outcomes, and tempers the read. Re-queried by intervention name, the 34 mechanism-annotated gaps have 220 phase-3 studies; 146 posted results, in which 78 primary analyses are significant and 20 are not (a further 48 post no numeric p). The verdicts are 11 positive, 7 mixed, 2 negative — the latter being a single in-class failure, fexipirant, whose 7 posted phase-3 studies contain no significant primary endpoint — and 14 with nothing posted. Reporting that layer alongside the counts is what keeps a failed programme such as oral DP2 antagonism from being ranked beside a successful one such as anti-IL-5: registration depth measures industrial commitment, not efficacy.

The same discipline applied to the nose gives the second result of this paper. Of 1,276 registered rhinitis studies, 269 posted results; 93 reached a significant primary endpoint and 38 did not, and the rejected list includes phenylephrine and a histamine-H3 antagonist — mechanisms that a registration-count view would still call "untested". Two non-type-2 openings survive that filter. The first is a completed phase-3 programme of a sublingual bacterial lysate in seasonal allergic rhinitis (3 trials, $n = 76-80$, symptom endpoints, **no results posted**): the cheapest open question in this paper, and the literature layer shows it is already partly answered — one of the three trials was published in 2023 with symptom scores favourable to the lysate (TNSS significantly lower vs placebo ($p=0.01$ TC1; $p=0.009$ V2)) while never being posted to the registry, so what remains open is replication, not existence. The second is the thymic/spleen-peptide immunomodulator class, which returns zero registered rhinitis studies and 2 published randomised rhinitis trials — the mirror image of the same blind spot. Crossing the gap list with a signature that does finally replicate between two cohorts ($p = 0.297$, odds 10.1) leaves 21 hypotheses that are both testable and mechanistically coherent, of which the one with existing human nasal evidence is roflumilast. The transcriptomic

screen itself remains negative under FDR control even after the signature is repaired, and we report it that way: the value of this paper is a ranked, falsifiable list of untested mechanisms, not a drug.

Read together, the evidence layers agree on one methodological point that generalises beyond rhinitis: a single-source literature view is misleading in both directions. A registry count classified roflumilast and spleen aminopeptide as never tested and the lysate programme as unreadable; the publication record reverses part of that and adds nothing to the rest. The inverse sweep of section 3.5 supplies the scale of that error: 79 drugs have published rhinitis trial titles and no registered rhinitis study at all, of which 25 are confirmed AR-population programmes — a population that is mostly old and mostly explained by the registry's late start, but one in which a drug with 94 published AR trials and zero rhinitis registrations (terfenadine) is unremarkable. The practical recommendation for a field deciding what to try next is therefore procedural, not pharmacological — cross the registry against the literature before treating a gap as white space, read the direction of every trial that exists in either source, and treat "not registered" as a statement about filing rather than about evidence.

References

1. Fokkens WJ, et al. Allergic rhinitis and its impact on asthma / EUFOREA position documents — current standard-of-care drug classes in rhinitis.
2. ClinicalTrials.gov v2 API (retrieved 2026-10-20); all study counts, phases and NCT numbers in this work derive from it.
3. GSE19187, GSE329964, GSE43523 — allergic rhinitis cohorts used for the disease signature (processing as in paper 1).
4. GSE261239 — paired nasal allergen-challenge cohort in birch-allergic and non-allergic subjects; source of the replicating contrast.
5. Corsello SM, et al. The Drug Repurposing Hub. Nat Med 2017 — target table underlying the screen.
6. Love MI, Huber W, Anders S. Moderated estimation of fold change and dispersion for RNA-seq data with DESeq2. Genome Biol 2014 (pydeseq2 implementation used here).
7. Smyth GK. Linear models and empirical Bayes methods for assessing differential expression in microarray experiments. Stat Appl Genet Mol Biol 2004.
8. Benjamini Y, Hochberg Y. Controlling the false discovery rate. J R Stat Soc B 1995.
9. PubMed E-utilities (esearch/efetch, NCBI; retrieved 2026-10-20) — source of every literature count in section 3.4.
10. Schmidt BM, et al. The phosphodiesterase 4 inhibitor roflumilast is effective in the treatment of allergic rhinitis. J Allergy Clin Immunol 2001;108:530–6 (PMID 11590377) — the unregistered human rhinitis trial discussed in sections 3.4 and 4.1.
11. Stelmach I, et al. Effect of immunostimulation with bacterial lysate on the clinical course of allergic rhinitis ... : a randomized controlled trial. Front Immunol 2023;14:1073788 (PMID 36733480) — the published outcome of the registry-silent lysate trial NCT04802616.
12. US Food and Drug Administration Amendments Act of 2007 (Public Law 110-85), Title VIII — the registration and results-posting requirements whose mid-2000s start dates the inverse sweep of section 3.5 measures against.
13. PubMed E-utilities trial-type query (`randomized controlled trial[pt] OR clinical trial[pt]` over allergic-rhinitis title/abstract terms, 3,391 records, retrieved 2026-10-20) — source of the inverse sweep in section 3.5 and of `results/literature_only_ar_drugs.csv` and `results/lit_inverse_title_audit.csv`.

paper 3 的转录组老药新用之所以一无所获，根源在于疾病签名不复现（发现与验证队列的基因级相关 $\rho = -0.090$ ）。本文换两条不依赖该失败签名的证据线，并把它们交叉。一、注册试验全景。我们用 ClinicalTrials.gov v2 接口抓取了过敏性鼻炎及 7 个邻近 2 型 / 炎症疾病（哮喘、特应性皮炎、慢性鼻窦炎、鼻息肉、嗜酸性食管炎、慢性荨麻疹、过敏性鼻炎）的全部干预性研究，共 10,007 项研究，解析出 9,486 条药物 / 生物制剂干预臂、6,939 条规范化药物记录、1,999 个不同物质。鼻炎本身并不缺试验（1,276 项），但机制极度集中：鼻用糖皮质激素、第二代抗组胺药、白三烯受体拮抗剂与过敏原免疫治疗占据榜首。关键发现：有 1,535 个物质在邻近疾病已注册研究、在鼻炎为 0；其中 407 个邻近研究 ≥ 2 项，34 个机制可人工注释，22 个已在别处做到 III 期。

二、可复现的鼻黏膜签名。GSE261239 提供配对的鼻过敏原激发对比（桦树花粉过敏者 11 对 11，非过敏者 12 对 12）。该对比与 GSE19187 显著复现： $\rho = 0.297$ ，两队列同向基因 141 个、反向仅 14 个（优势比 10.1）；相比之下，paper 3 所使用的队列对只有 $\rho = -0.090$ 。用这两个对比构建 v2 签名（27,434 个基因）并沿用完全相同的筛选机器（同一药物库、同一分层置换零模型、同一 $+1/(n+1)$ 与 BH-FDR）重跑后：药物打分的跨队列可迁移性从 $\rho = -0.090$ 变为 $\rho = 0.302$ ，名义命中从低于随机的 25 个升至 129 个（期望 40 个、约 3.2 倍），但仍然没有任何药物通过 FDR 控制（最小 FDR 0.154）。这一阴性结论如实保留：v2 是可迁移性的改善，不是签名的修复（对留出队列 GSE329964 仍为 $\rho = -0.137$ ）。

三、交叉假说。把「邻近疾病已注册、鼻炎 0 注册」的机制与 v2 签名中该靶点的方向交叉，得到 34 条排序假说：15 条方向一致、2 条混合、13 条方向相反、4 条靶点不在鼻黏膜芯片上（不视为反对证据），共 21 条未被转录组否定。经第六节文献层校验后，第一位是 PDE4 抑制剂 roflumilast——本文唯一有直接人体鼻部试验证据的候选（2001 年随机交叉试验、结果阳性、从未注册；PDE4B $z = 1.387$ ）；其余按邻近疾病先例排序：抗 IL-5 / IL-5Ra（mepolizumab、benralizumab，各自在含慢性鼻窦炎 + 鼻息肉的 6 个邻近病种中做到 III 期，鼻炎注册数为 0；已贴出的 III 期主要终点里，mepolizumab 7 显著对 1 不显著，benralizumab 9 对 3）；抗 IL-13（lebrikizumab、tralokinumab，IL-4Ra $z = 1.65$ 支持该通路）；BTK 抑制剂（remibrutinib，在慢性荨麻疹做到 III 期、靶点 BTK $z = 1.13$ ，是唯一小分子 IgE / 肥大细胞轴选择）；抗 OX40L / OX40（amlitelimab、rocatinlimab，TNFRSF4 $z = 1.07$ ；注意 amitelimab 的 6 项 III 期尚无结果贴出）；JAK1 抑制剂（upadacitinib、abrocitinib，已贴出的 III 期主要终点全部显著，且 upadacitinib 在 AD 共病分析中报告 AR 症状改善）；以及新一代抗 IgE（ligelizumab；omalizumab 是唯一在鼻炎注册过的抗 IgE，共 10 项）。

四、项目结局审计。注册数量衡量的是投入，不是疗效，所以我们把 34 个候选按干预名重查了一遍（与前述按条件查询的计数口径不同，两者并列报告、不合并）：III 期研究 220 项，已贴出结果 146 项，主要终点分析显著 78 项、不显著 20 项、未报 p 值 48 项；按药物判读为 11 个阳性、7 个混合、2 个阴性、14 个无 III 期 p 值可读。最典型的阴性反正是 DP2 / CRTh2 拮抗剂 fevipiprant：7 项 III 期全部贴出结果，无一显著主要终点（可读 p 值 0.059–0.665，另有 1 项未报 p），26.7% 的邻近疾病研究提前中止——同一条注册记录十年前会把它排在榜首。这层审计正是为了不让已失败的项目与真正成功的项目并列，也提醒读者：候选清单的排序是优先级输入，不是疗效证据。

五、鼻腔自身的结局记录。反过来审计鼻炎本身：1,276 项注册研究中有 269 项贴出结果（其中 179 项为 III / IV 期），判读为 93 项主要终点显著、38 项不显著、138 项未报 p 值。鼻腔已否定的机制里，除一大堆过敏原制剂标签外，干净的药物例子是去氧肾上腺素（phenylephrine）（4 项贴出结果、2 项不显著、0 项显著，最大 $p = 0.561$ ；这正是本层结局审计的表面效应度检验）与 H3 拮抗剂 PF-03654746（II 期， $p = 0.097$ ，在鼻炎中已试且未做出显著性）。另有第三类缺口——既非「从未试过」也非「试过失败」，而是在注册库中没有痕迹：舌下细菌溶解物（Ismigen/PMBL）在季节性鼻炎有 3 项已完成、安慰剂对照、以症状评分（TNSS / PNIF / VAS）为主要终点的 III 期试验（ $n = 76-80$ ）、3 项全部未贴出结果——其中 NCT04802616 已在 2023 年发表（PMID 36733480，50 例儿童季节性鼻炎，TNSS 与 VAS 显著优于安慰剂、PNIF 无差异），见第六节；而「脾氨肽 / 胸腺肽」类免疫调节剂按干预名在注册库中为 0 项（脾氨肽、氨肽酶类、胸腺多肽均查无研究；spleen polypeptide 4 项、匹多莫德 1 项、胸腺五肽 10 项、胸腺法新 65 项，均无一在鼻炎注册），但脾氨肽在文献中已有 2 项随机鼻炎试验。这是一个机制上独立于 2 型细胞因子轴的方向：注册库里没有结论，而文献层（第六节）给出了部分方向的证据。

六、文献层交叉校验（「0 注册」是否等于「从未被研究」）。对 34 个候选逐一检索 PubMed：宽口径（rhinitis / nasal / nose）有 30 个命中，但读题后确认多为哮喘与慢性鼻窦炎的综述与个案（mepolizumab 一项即返回数百条），不能当作证据计数；严口径（“allergic rhinitis” / seasonal / perennial / rhinoconjunctivitis / hay fever）命中 22 / 34，12 个零命中。对全部 80 条命中记录（69 篇唯一文献）逐条读题分类后，只有 2 条属于「药物给了、鼻炎症状为终点」的人体研究：① roflumilast——2001 年发表于 *J Allergy Clin Immunol* 的随机、双盲、安慰剂对照交叉试验（25 例，口服 500 $\mu\text{g/day}$ ，9 天，配每日鼻内花粉激发；鼻气流、鼻痒与流涕、主观鼻塞均优于安慰剂；PMID 11590377）。该试验从未在 ClinicalTrials.gov 注册（roflumilast 鼻炎注册数为 0），而它在邻近疾病中的项目判读为 positive，靶点 PDE4B 在 v2 签名中上调（ $z = 1.387$ ），因此在本清单中被提为首位；限制是样本小、交叉激发模型、急性给药、口服 PDE4 抑制的胃肠与体重不良反应；② upadacitinib 2024 年在重度特应性皮炎人群中的共病分析（PMID 38329235，*J Eur Acad Dermatol Venerol*）报告合并过敏性鼻炎 / 哮喘的症状改善——该药是 AD 的适应症用药，AR 症状属共病终点。其余命中为 6 条共病分层研究、8 条动物 / 体外 / 分子模拟研究、1 条「未给该药」的人体 AR 研究、25 条其他疾病的人体研究与 38 条综述 / 无关条目。结论：注册缺口不等于文献缺口，但已接近于证据缺口；反向命题同样不成立——脾氨肽注册 0 项、文献有 2 项随机鼻炎试验（PMID 39524381, 39142892），匹多莫德注册 1 项（另一适应症）而文献有 2 项鼻炎研究，OM-85 命名下注册 0 项鼻炎研究而文献有 2 项人体鼻炎研究。两类盲区互为镜像：注册了却没有结果（细菌溶解物）与发表了却从未注册（roflumilast、脾氨肽）；单看任一来源都会判错，这也是本文方法论层面的主要结论。

七、反向文献扫描（文献里有、注册库里没有）。把第六节的问题倒过来、放到全量尺度上：我们抓取 PubMed 中所有被标引为「过敏性鼻炎试验」的文献（3,391 条），取其标题，与 6,565 词的药物词典（老药新用库 + 人工补充）做 n-gram 匹配，再与注册库鼻炎药物集合对照。1,883 条标题提到至少一个词典药物；合并同义名 / 盐型后匹配到 144 个标签，其中 63 个也出现在已注册鼻炎研究中，81 个完全没有——即 79 个药物 + 2 个非药物标签（针刺；substance P 是被测量的介质而非干预）。剔除仅出现在药代 / 健康志愿者等「项目类论文」标题中的命中（66 条、涉及 32 个标签）后，文献独有药物共 376 条试验标题，其中 35 个药物有 ≥ 2 项。对这 35 个药物逐条读题人工判读（全部命中标题见 results/lit_inverse_title_audit.csv）：25 个确为 AR 受试人群试验（合计 314 篇），5 个实为共病 / 机制激发研究（cimetidine, deflazacort, furosemide, lodoxamide, terbutaline），5 个命中的其实是别的疾病（抗生素类的「鼻炎」出现在鼻窦炎试验摘要里）。关键判读：这份名单不是「被埋没的新药」——榜首多为 1970–1990 年代的老药，没有注册鼻炎试验的原因是当年尚无注册制度（本文文献起点为 1960–70 年代，而试验注册与结果公开在 2000 年代中期才普及）；21 个至今仍在售、7 个仅区域上市、3 个因心脏毒性撤市（特非那定、阿司咪唑、非诺特罗）、1 个研发终止（TRPV1 拮抗剂 SB-705498）。因此这一层的价值是校准而非扩充候选：roflumilast 那类「发表了却从未注册」的试验是一个大类的常态，而不是异常；「鼻炎 0 注册」只说明注册库的时间起点，不说明药物没被测过。

必须声明的限制：「鼻炎 0 注册」仅指本注册记录中没有，不等于从未被测试（本接口基本不含中文注册库）；生物制剂的机制标签为人工注释；「靶点不在芯片上」不等于靶点无关；邻近疾病的阳性结果不能自动外推到鼻腔；p 值只在申办方贴出结果时可得，同一研究的多个主要终点未做多重校正，“未报 p 值”多为药代 / 安全性终点，研究级结果不等于项目级结论；两种计数口径（按条件 / 按干预名）不一致，本文并列报告；逐药结局表里混有非药物实体（“aerosol”、“allergen”、“vitamin”、过敏原制剂商品名），其计数不可当作药物判读。文献层（第六节）为标题级检索与人工分类，计数是下界而非穷举，某候选未见文献不等于不存在，且宽口径计数明显高估；roflumilast 与脾氨肽两处关键对照已分别核对原始摘要与注册库探针。反向扫描（第七节）受 MEDLINE 试验类型标引限制（正向检索到脾氨肽 2 项随机鼻炎试验，本扫描只看到 1 项），且词典以老药新用库为主，库外且补充表未收的药物按构造无法被发现；其中「撤市 / 区域上市 / 仍在售」为人工对公开记录的判读，未逐一核对监管数据库。本文全部输出均为假说生成，不是任何药物在鼻炎中的疗效证据。

关键词：过敏性鼻炎；药物重定位；ClinicalTrials.gov；证据缺口；鼻黏膜转录组；2 型炎症；假说生成。