All Successful Oncology Targets Are Alike: A Rule of Five for Target Selection

By Pawel Zawadzki, PhD, Founder & CEO, Gordion Bioscience

Why We Need a Rule for Targets

In 1997, Christopher Lipinski and colleagues published a short paper in *Advanced Drug Delivery Reviews* that would quietly become one of the most consequential contributions to modern drug discovery. The paper described simple physicochemical parameters: molecular weight under 500, logP under 5, hydrogen bond donors under 5, and hydrogen bond acceptors under 10, which together predicted whether a small molecule stood a reasonable chance of being orally bioavailable. The paper has since been cited more than 24,000 times.

Lipinski's contribution was not that the individual parameters were novel. Every one of them was discussed in the medicinal chemistry literature of the early 1990s. His contribution was that he looked at which molecules successfully made it from Phase 1 to Phase 2 and which ones quietly died in development, and then he distilled the pattern into numbers that fit on a napkin.

Oncology needs the same thing, but for a different problem.

Medicinal chemistry has rules for asking whether a molecule will behave well in the body. Oncology has no equivalent rule for asking whether a target is worth pursuing in the first place. We have decades of failed programs, billion-dollar bets on targets that looked compelling in a cell line and disappointed in a patient. We have a handful of spectacular successes where the targeted agent walked into the clinic and delivered 80% response rates. And we have almost no principled way of distinguishing the two categories before clinical readout delivers the verdict.

This essay proposes a Rule of Five for oncology target selection, five criteria that separate the targets likely to produce durable clinical responses from those that will likely disappoint. Like Lipinski's rule, it is not a hard cutoff but a framework for disciplined thinking. Like Lipinski's rule, it emerges from looking at what worked and what didn't across hundreds of oncology programs.

The analysis presented below relies on exploring treatment-naive patient molecular data. Cell line and mouse models are extremely important, but I'll argue that at the target selection stage, patient data is so much more translatable into the clinical reality than any model system one can think of.

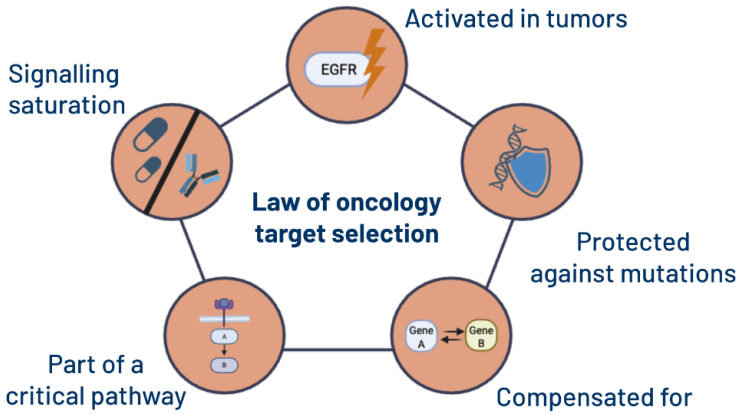

Fig. 1. ***Rule-of-five****. Inspired by Lipinski's rule of 5, and the experience of assessing hundreds of oncology drug programs, a handful of simple rules characterising successful drug targets emerged.*

The Observation That Prompted the Framework

Over the past two decades, a small number of oncology targets have produced therapies that genuinely transform patient outcomes. Osimertinib in EGFR-mutant lung cancer delivers response rates near 80%. Alectinib in ALK-rearranged lung cancer delivers response rates over 83% and median overall survival exceeding 81 months. Trastuzumab (or Enhertu) in HER2-positive breast cancer has fundamentally changed the natural history of the disease.

Then there are dozens, if not hundreds, of targets with extremely strong preclinical data that simply did not translate into clinical reality.

When you place the successes and failures side by side and ask what distinguishes them, a pattern emerges. The successes share five biological features that the failures lack. The features are not individually obscure; what is striking is that no one has systematically codified them into a decision framework. This essay is an attempt to do so.

The five rules are:

1. **Activation** - the target is activated in the intended patient population, and there are multiple ways tumors use to activate a gene.
2. **Evolutionary protection** - the gene encoding the target shows a footprint of negative selection in the intended indication: deleterious mutations are underrepresented.

3. **No compensating paralogs** - the target has no paralog capable of substituting for its function.
4. **Pathway centrality** - the target sits within a canonical cancer-sustaining pathway or hallmark.
5. **Pathway saturation** - inhibiting the target produces enough pathway downregulation to trigger a therapeutic effect.

A target that satisfies all five is a candidate for monotherapy with high response rates. A target that satisfies fewer rules is a candidate for rationally designed combinations or a modality like a multispecific antibody.

Rule 1: Activation.

The first question to ask about a candidate target is not whether the protein is present in the tumor, but whether and how it is hyperactivated there.

EGFR is the canonical illustration. EGFR is expressed in most epithelial tissues throughout the body. But in a subset of tumors, EGFR is hyperactivated by mutations in the kinase domain that cause the receptor to signal constitutively in the absence of ligand. Osimertinib with >80% ORR and > 47months OS proves the rule no. 1.

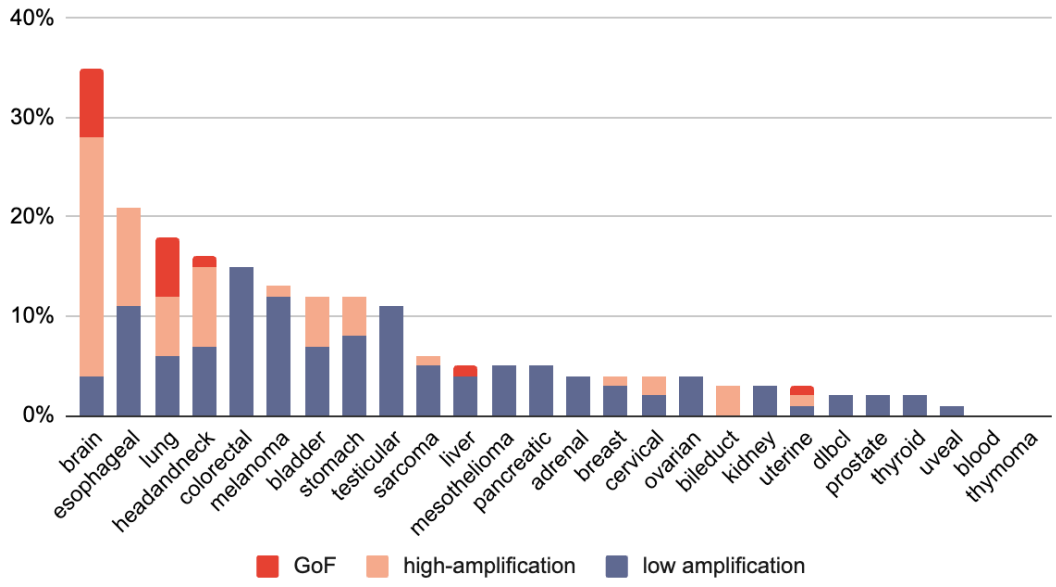

Figure 2. Different ways tumors can activate EGFR across tissues. GoF (Gain-of-Function) mutation; high amplification (≥7 additional EGFR gene copies); low amplification (2-6 additional gene copies).

BRAF tells the same story with even greater precision. The Phase 1 trial of vemurafenib in unselected melanoma was stopped early because the responses in V600E patients were so

dramatic that continuing to enroll wild-type patients became ethically difficult. When BRAF in melanoma is hyperactivated by this particular mutation, blocking it is transformative.

HER2 is even more remarkable. Today, we have drugs successfully targeting HER2 activated by mutations but also trastuzumab or enhertu showing remarkable clinical benefit in patients overexpressing HER2, not necessarily activating it by mutation.

Unfortunately, HER2 success led to a conclusion that every gene overexpressed in tumor is a good drug target. The vast majority of drug targets identified as overexpressed in tumor vs. normal tissue failed to deliver clinical efficacy comparable to that observed with trastuzumab. Furthermore, recent lessons from the ADC field show that merely considering the membrane receptor expression does not translate into clinical efficacy.

The first rule: demonstrate that the target is hyperactivated above baseline, whether by mutation, amplification, ligand-driven signaling, fusion, or loss of a negative regulator.

Rule 2: Evolutionary Protection

The observation underlying rule no. 2 is that tumor cells are subject to Darwinian selection. A tumor accumulates mutations at a rate exceeding that of healthy cells, and those mutations are randomly (not quite, let's simplify) distributed across the genome. In the absence of selective pressure, any given gene will acquire a predictable rate of mutations: premature stop codons, frameshifts, etc.

When you look at a target gene across a large cohort of tumors and compare the observed rate of mutations to the rate expected from background mutagenesis, one of two things will be true. If the gene is expendable, mutations will accumulate at or above the expected rate, sometimes dramatically above, because loss of the gene provides a fitness advantage (tumor suppressor). If the tumor depends on a gene to sustain the malignant phenotype, then mutations will be systematically underrepresented. The gene is under negative selection. The tumor cells that happen to acquire damaging mutations do not proliferate; only cells with an intact, functional copy are visible in the sequenced population.

This signature, the underrepresentation of deleterious mutations relative to expectation, is the footprint of a good target visible at the level of population genomics. And critically, it is indication-specific. The same gene may show strong negative selection in one tumor type and no selection at all in another. The indications where negative selection is strongest are precisely the indications where targeting that gene will produce the best clinical responses.

But the challenge is which mutations one is looking at. It was shown that the typical population genomics approach of comparing Synonymous vs. Nonsynonymous mutations can be applied to target assessment (https://pubmed.ncbi.nlm.nih.gov/29056346/), yet to reliably assess targets in smaller indications, hundreds of thousands of samples would need to be investigated. We are

happy to share that a lot is to be learned when different mutation types are investigated separately or when grouped together. Furthermore, we discovered that detailed analysis of a gene's copy numbers can provide another level of the evaluation of evolutionary pressure, allowing to separate good drug targets from poor ones.

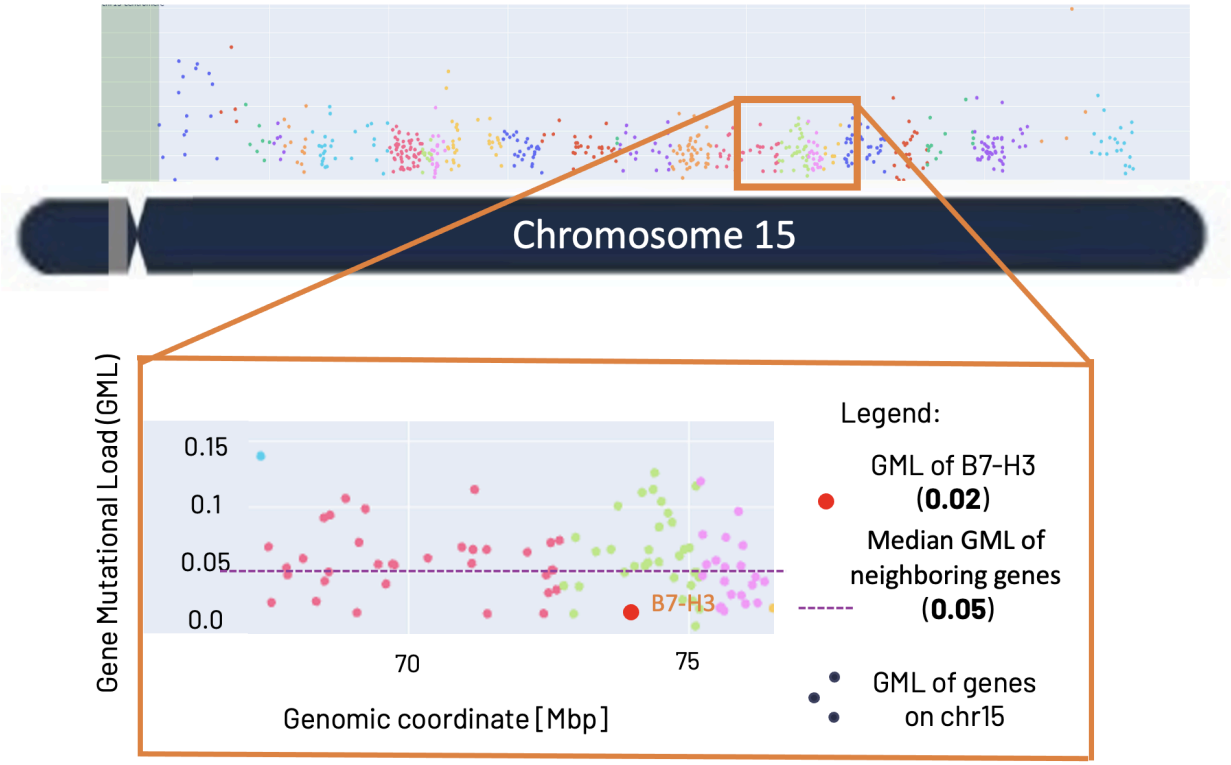

Fig. 3. ***Analysis of Gene Mutational Load (GML)****, showing the burden of point mutations in genes on the q arm of chromosome 15. Each dot represents a distinct protein-coding gene, and the horizontal line is the median for all genes. Here, evidence shows that B7-H3 is a promising oncology target, perhaps currently tested in the wrong patient population.*

The second rule: look at whether the gene encoding the target shows a footprint of negative selection in the intended indication.

Rule 3: No Compensating Paralogs

The third rule concerns what happens to the pathway when the target is blocked. Biological networks contain redundancy, and the most common form of redundancy is a paralog, a closely related protein that can perform the same molecular function when the primary protein is unavailable.

The RAS family is the clearest illustration. KRAS, NRAS, and HRAS are three paralog small GTPases, all capable of activating the MAPK cascade through identical downstream biochemistry. In most RAS-addicted tumors, one paralog dominates, typically KRAS but the other paralogs remain present and expressed at baseline. When KRAS is inhibited, NRAS in particular can be upregulated to compensate. The pathway continues to signal, at reduced but not eliminated levels, through the paralogs.

VEGFA has three paralogs (VEGFB, C, D). Typically, cell lines are used to evaluate the relationship between paralogs, but patient data can also provide this information. Ask a simple question: when my target (VEGFA) is downregulated (via genomic mutation or expression downregulation) will tumor cells upregulate any of the paralogs? The answer can have massive implications in terms of what you will see in the clinic and for VEGFA, I can confirm that no compensation was detected in patient data.

The third rule: a target is a monotherapy candidate only if it has no close compensating paralog. Targets with paralog are not wrong targets; they require a molecule that will block the compensating paralog as well.

Rule 4: Pathway Centrality

The fourth rule is about where the target lives in the broader biology of malignancy. Cancer, across its bewildering heterogeneity, reliably uses a limited set of cellular machineries to sustain itself: MAPK signaling for proliferation, PI3K/AKT/mTOR signaling for growth and survival, the cell cycle machinery, the apoptotic machinery for death evasion, and a smaller set of pathways for angiogenesis, immune evasion, metabolic reprogramming, and metastasis. This framework is captured in the Hallmarks of Cancer literature (https://pubmed.ncbi.nlm.nih.gov/41616779/).

The rule is simple: targets that sit within these hallmark pathways produce clinical responses when engaged. Targets outside them produce, at best, modest incremental benefit.

The success stories obey this rule with striking consistency. EGFR, HER2, ALK, ROS1, RET, BRAF, KRAS - all sit within the MAPK or closely coupled receptor tyrosine kinase signaling axis. PI3K, AKT, mTOR — hallmark growth signaling. CDK4/6 — hallmark cell cycle regulation. BCL-2 — hallmark apoptosis regulation. PD(L1) - hallmark of immune evasion. Every major oncology target success of the last two decades has been a direct hit on a major hallmark pathway.

The principle behind rule 4 is that cancer cells do not tolerate disruption of major hallmark pathways. These pathways are the ones that actually sustain the malignant phenotype. Here, the preclinical models are particularly dangerous as the models are optimised to show strong dependency regardless of whether the dependency is tier 1 or tier 3 importance for cancer survival. Also, novel targets showing high dependency in a limited number of artificial cell line models should be taken with a grain of salt.

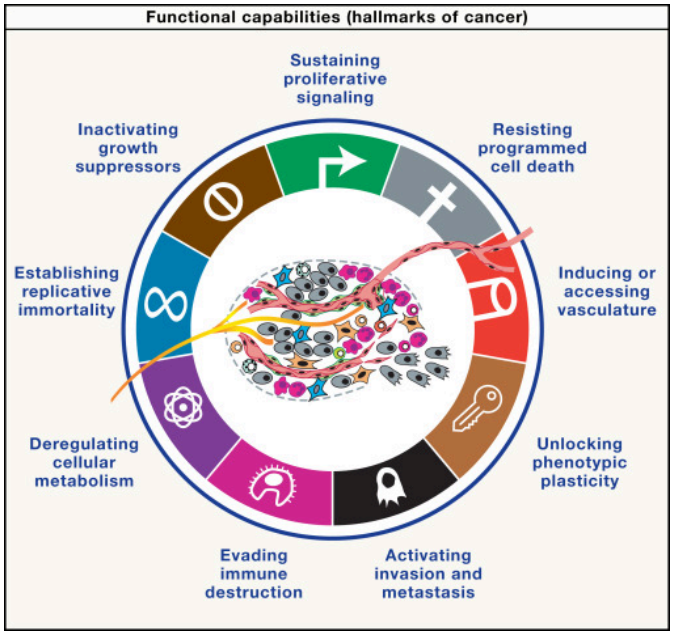

Fig. 4. ***The hallmarks of cancer*** *(Hanahan 2026). Not all hallmarks are created equal; some hallmarks delivered multiple successful therapies while others still remain elusive.*

The fourth rule: if a target does not sit within one of the most important hallmark cancer pathways, a very high bar of evidence should be required before a lead oncology program is initiated.

Rule 5: Pathway saturation - Does inhibiting the Target Collapse the Pathway Enough?

The fifth rule is, in my view, the most underappreciated in oncology drug development. It has been misread as a potency question or a dose-response question. It is a question about the biology of the driver itself: does the target produce enough pathway activity, *on its own*, to sustain the tumor, or does it require additional amplifying inputs to reach the level of pathway activity that tumorigenesis requires?

The logic is this. Every cancer-sustaining pathway has a threshold. Below some level of pathway output, the cell cannot sustain the proliferation rate, survival, or metabolic reprogramming that defines the tumor phenotype. Above that threshold, the pathway is

functional for tumorigenesis. The question for any candidate target is: when the target is inhibited, does the pathway drop below the tumorigenesis threshold, or does it merely drop to a lower but still-sufficient level?

The answer depends on whether the target is driving the pathway sufficiently on its own or whether it is operating in a "supercharged" configuration that requires additional amplifying inputs.

EGFR activating mutations in NSCLC are the cleanest example of driving sufficiency. The mutant EGFR produces MAPK output that is, on its own, sufficient to sustain tumorigenesis. The tumor does not require additional upstream or parallel inputs. When the mutant EGFR is inhibited with osimertinib, MAPK output collapses not to some intermediate level but below the tumorigenesis threshold. The result is durable, deep response.

ALK fusions show the same pattern. The ALK fusion protein produces constitutively active kinase signaling that, on its own, drives MAPK and downstream pathways sufficiently for tumorigenesis. Inhibiting the fusion with alectinib or lorlatinib drops the pathway output below the threshold. The responses are deep and durable.

BCR-ABL in CML is the archetype. The fusion protein is the sole driver of the disease. Inhibiting it with imatinib drops the pathway output below the threshold needed to sustain the malignant clone, and the clinical responses are so profound that CML has been transformed from a rapidly fatal leukemia into a chronic, manageable condition.

Now consider **KRAS activating mutations in PDAC**. The mutant KRAS is constitutively active, and it is unquestionably essential, but it is not, on its own, sufficient. To reach the pathway output required for aggressive tumorigenesis, the cell requires additional amplifying inputs: MET amplification, EGFR ligand-driven signaling, TGFbeta, or other amplifiers that "supercharge" the pathway above what mutant KRAS alone would produce. When KRAS is inhibited with daraxonrasib or sotorasib, MAPK output drops, but because the supercharging inputs remain active, the pathway does not drop below the tumorigenesis threshold in most patients. The tumor slows but does not collapse. The clinical result is the now-familiar pattern: 35–41% response rates, median PFS of 6–9 months, and progression as the supercharged pathway reasserts itself.

PI3K mutations in solid tumors tell the same story. PI3K activating mutations are real drivers, but in most contexts they require supercharging. PI3K monotherapy outside of specific contexts (PIK3CA-mutant, hormone receptor-positive breast cancer with alpelisib) has disappointed, because inhibiting PI3K alone does not drop downstream pathway output below threshold.

The big question for rule 5, therefore, is not about the drug target itself. It is about the tumor biology: **does** the target drive the pathway sufficiently on its own or requires supercharging by other players? The rule is testable in patient data. The recipe is simple (not sure whether easy):

in a treatment-naive patient biopsies with activated target of interest, measure all the ways the tumor activate the same pathway the target is driving.

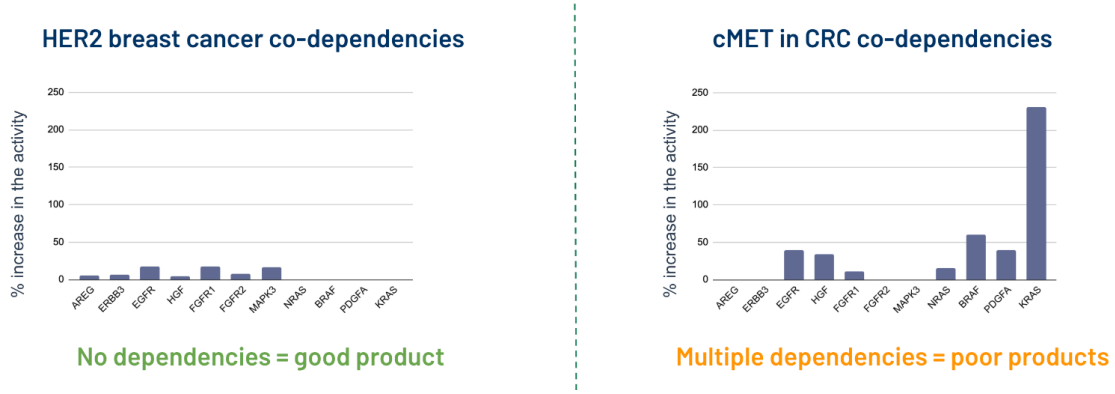

Fig. 5. ***Compensatory network.*** *Here, the genetic events (activations, gene amplifications) co-occurring with HER2 and cMET (high according to IHC) were measured for selected genes. Even without transcriptomic analysis or sophisticated pathway deconvolution, the message from these data is clear: HER2 inhibition will reduce signalling significantly more compared with cMET inhibition.*

The fifth rule: before pursuing a target as a monotherapy, ask whether inhibiting it will drop pathway output below the tumorigenesis threshold. Supercharged drivers like KRAS, cMET and PI3K struggle, not because they are wrong targets, but because the therapeutic strategy is not accounting for the compensatory network.

How to Use the Framework

A target that passes all five rules is a monotherapy candidate with the potential for 70–80% response rates and durable benefit. Pursue it as such. A target that fails some rules can be pursued, but only with the specific strategy designed against the failing rule.

In the quarter-century since Lipinski's rule was published, it has helped prevent countless molecules from progressing into expensive development only to fail on bioavailability. The analogous goal for a Rule of Five in oncology target selection is to prevent the expensive clinical disappointment that has characterized too much of the past two decades.

All successful oncology targets are alike. Every failed target fails in its own way. Recognizing the pattern is the first step toward producing fewer of the latter.

Pawel Zawadzki, PhD, is the founder and CEO of Gordion Bioscience, a Cambridge, Massachusetts-based biotechnology company translating real-patient cancer dependency data into first-in-class multispecific antibody therapeutics. The Rule of Five is the analytical framework Gordion uses to select and develop its pipeline.