

TOP-*k* READING: AGAINST POLISHING THE MEDIAN ESSAY

tags: [typography](#), [writing psychology](#), [RL exploration](#), [order statistics](#)

How much to edit? A top-k attention-window toy model of the publish-polish trade off: You should polish in inverse proportion to how much of your writing a reader will see; the more diverse your writing, the better off you are writing more rather than better.

2026-05-10–2026-05-11 · *finished* · certainty: *likely* · importance: 4

The Trade-Off

Self-Cannibalization

Order Statistics

Polish the Window

Heterogeneity

Eclectic Generalists


Similar Links

How should a writer trade off new writing (exploration) against polishing old writing (exploitation)?

I suggest a simple order statistics toy model of readers with a finite ‘window’ of reading time for an author’s best writing, where each new writing has a chance of being good enough vs polishing an old favorite. A new writing only matters if it enters the window; an edit only matters if it improves something in the window.

The implication is simple. Polish when a marginal edit improves work inside many readers’ windows: consensus hits, durable pieces, introductory essays, or high-stakes publications. Write more when readers have large windows, heterogeneous tastes, or when a new topic can create a different local hit for a different audience.

For eclectic generalists with diverse readerships like me, the strategy is *not* to perfect the median piece. Instead: polish the winners, ignore mid-tier work, and keep trying to create new local hits.



YEARS AGO, I DISCUSSED THE APPARENT “PERFECTION PREMIUM” IN POLISHING WRITING, and how a writer should quickly hook readers upfront in a new writing. But how should we think in the first place about polish in writing vs new writing?

THE TRADE-OFF

Most writers online, I think, give up relatively soon. This is not because we ran out of things to do—there’s *always* another typo we could fix... But when we choose to spend our limited time and energy hunting for typos or fixing formatting, there’s a definite sense in which we decide it’s no longer ‘worthwhile’ trying to improve old writing after some point, and we’re better off doing new writing.

But this is a little odd: if we’re not ‘done’ with old writing, why should we start something new? And if we have a large backcatalogue, we can hardly plead poverty of material; how can I argue with a straight face that *my* readers lack for things to read and so it is *imperative* I go write something new rather than polish an old popular essay?

The economics or optimal foraging answer would be the heuristic ‘you should switch between whichever has the highest marginal return at any given moment’—but marginal return of *what*? Even assuming we have some ‘utility’ definition we’re satisfied with, how, exactly, does writing turn into that? If I write an essay or I polish an old one, what does that mean for a reader?

A toy model that seems plausible to me is that what is going on is that each *reader* has a certain finite size or appetite or window of time, which a writer tries to fill up as best he can. (I’ll use *time* as my unit here instead of number of words or essays, because I think that increasingly, it is time and attention which is the bottleneck.) One reader has a window of 5 minutes of ‘Gwern reading time’ a year, while another has 50 minutes, and a big fan has 500... The absolute size of the backcatalogue doesn’t matter for each reader—only the best k minutes of your writings.

So, when you write, if a new piece doesn’t make it into the top- k and is only the $k+1^{\text{th}}$ best reading-minute, then a reader doesn’t care. They’re never going to see it and it doesn’t exist for them.

This helps account for why writing is such a “hit-driven business”: after a while, no matter how prolific you are, it makes zero difference to the reader, because you filled up the top- k in their window and nothing budes after that.

SELF-CANNIBALIZATION

To move the needle, you need to beat your best and push something out of their window. And frustratingly, once you do that, the reader is only slightly happier with you!

This is because you are cannibalizing yourself. If you wrote something which was a +10 in some Quality Unit, and then you write something which was +11, you feel proud of yourself for adding a +11 and thus writing a total +21 of text—that was hard work—but as far as the reader is concerned, it’s just a net +1 because your +11 pushed out the +10.

Perversely, the more you write, the harder it gets to improve your window. If a reader has a window of 10 essays and you’ve written 100 essays to date, then your next essay has to beat >90 of your prior essays to make it into the window at all!

Simply put, I think you should polish in inverse proportion to how much of your writing a reader will see; the more diverse your writing, the better off you are writing *more* rather than *better*.

ORDER STATISTICS

Rank entry · Equal odds · Lumpy returns

We can treat this statistically using order statistics; for example, if we imagine some normally-distributed quality statistic, our expected gain is only:

R simulation of the marginal gain from adding 1 normally-distributed piece to a 100-piece back catalog with a 10-piece reader window.

```
set.seed(2026-05-11)

windowGain <- function(backlog_n = 100, window_k = 10, simulations = 1e7) {
  mean(replicate(simulations, {
    old_quality <- rnorm(backlog_n)
    new_quality <- rnorm(1)
    cutoff <- sort(old_quality, decreasing = TRUE)[window_k]

    max(0, new_quality - cutoff)
  })))
}

windowGain()
# [1] 0.0474129394
```

$\sim 0.05\sigma$ better—which is almost invisible percentile-wise.

The simplification of treating pieces as independent samples is crude, but it resembles the equal-odds model summarized by Jung et al 2015: hit-rate is approximately stable, so quantity matters because it buys more rolls of the dice. For writing, this assumption is partly defended by offsetting forces: positive factors like experience are offset by negative ones like using up ideas.

More realistic distributions do not change the basic rank fact. With n old pieces and a k -piece window, a new independent piece enters the window with probability $k/(n+1)$, distribution-free. A power law or other heavy-tailed quality distribution changes the payoff conditional on entry, not the entry probability: the process becomes more hit-driven, with long stretches of zero marginal value punctuated by rare large replacements.

POLISH THE WINDOW

Guaranteed gain · Tiny windows · Decay

This also adds a little insight to the value of polish: while it's true fixing typos in your best writings doesn't make them much better, it does make them a little better... And importantly, it's making the writings *in the window* better. When you write a new one, you usually fail to replace something in the window, rendering it moot, and if you do, it's probably a modest improvement; but when you edit the writings in the window, you are getting a small but guaranteed improvement.

This sheds some light on why you might escalate your quality standards going from a blog

post to, say, a *New York Times* article. (Jasmine Sun mentioned in a recent Q&A at *Inkhaven* 2 that her *NYT* piece on peptides might have 10 million readers total, vastly exceeding her lifetime total of Substack readers—so of course she spent months sweating every detail!)

Meanwhile, a breaking-events periodical newsletter like *Zvi Mowshowitz's* will often not bother to fix typos or broken URLs, because the temporal decay guarantees even a good issue will soon drop out of the window, and it needs to rush to write the next one.

For the former, you are talking to a small steady audience with a large window; there's little point in polishing it. But for the latter, you may get only one, and so the window is small—you'd better make it as good as you can!

HETEROGENEITY

So, to the extent that your readers are homogenous and have a small window and a universally shared ranking, the better off you are, after a while, trying to polish your writing winners, because you aren't going to easily beat them but you can still make them a bit better.

But to the extent that they have large windows or heterogeneous different rankings (eg. readers prefer completely different articles), you're better off just writing more in the hopes of getting a new better one into their windows. The marginal return on every new essay may be quite low (because its topic will appeal to only a small fraction of readers) but it is also true that the return is even lower on polish (because those too are niche but the effort is constant). In the extreme case, every reader has a window size of 1 and prefers a different essay... So you are best off spending 100% of your time writing on a new topic every time, because each one will find a different audience.

In full generality, we could advise a writer to polish in proportion to:

$$(1/\text{Window}) \times \text{Homogeneity} \times \text{Shelf-life} \times \text{Stranger-fraction}$$

And to write more in inverse proportion to all 4.

ECLECTIC GENERALISTS

For my writing, I think it falls into the latter group: I do not write for prestige or popular periodicals, but on many topics for different readers (eg. darknets vs poetry vs AI vs genetics...). This also describes writers I am sometimes grouped with, like *Scott Alexander* or *Tyler Cowen's Marginal Revolution*. We all share, if nothing else, a generalist eclectic approach. Under this model, Scott Alexander, who has a recognizable 'beat', should probably invest more in polish, while I should emphasize more diverse new writing rather than over-polishing, and Tyler Cowen is probably fine in his omnivorous niche of writing about anything with no polish.

So I *should* be writing a lot more on more topics, instead of falling into the artisanal/hobbyist/indie-creator trap of spending as much time as I do on polish and infrastructure. What can I say? I don't write to be popular, and the polish is for me a personal esthetic and an ethos, not a calculated strategy. (I try not to pretend otherwise, and I tell any aspiring writer who asks that they definitely should not try to imitate Gwern.net but focus on writing. They have my permis-

sion to fiddle around with typography after they have, let's say, a million words—at which point they may have some writing worth polishing.)

But other writers might find this a helpful frame: do your readers have a lot of appetite each, or a little? And are they homogenous or heterogeneous? This can help guide you in deciding how much to invest in copyediting, illustrations, marketing, etc.

SIMILAR LINKS

- ✦ [Spend Writing Weirdness Points Well](#)
- ✦ [You should write more online](#)
- ✦ [Thoughts while watching myself be automated](#)
- ✦ [Towards Better LLM Creative Writing](#)
- ✦ [Virtual comments: LLM idea](#)
- ✦ [‘Tools For Thought’ Failure](#)
- ✦ [Proposal: Hybridizing Forums/Wikis](#)
- ✦ [Towards better RSS feeds](#)
- ✦ [Design Of This Website § Returns To Design?](#)
- ✦ [Miscellaneous](#)
- ✦ [Miscellaneous § How Would You Prove You Are a Time-Traveler From the Past?](#)
- ✦ [You Can't Countersignal Online Writing](#)
- ✦ [\[Grok, Could you please name 20 great bloggers that you like?\]](#)
- ✦ [Adding Bits Beats AI Slop](#)
- ✦ [Why do writers still underestimate LLMs?](#)
- ✦ [LLM-optimized research articles](#)
- ✦ [So you want to go to next year's Inkhaven: How *Not* to write 30 blog posts in 30 days](#)
- ✦ [30 Days of Writing: My Inkhaven Experience; Who knew monk mode is at odds with socializing.](#)
- ✦ [The LessWrong 2022 Review § Cost of Book Production](#)
- ✦ [First, Make Me Care](#)