



METHOD NOTE · BA-MN-3

Survivorship bias in visibility case studies

Published case studies of visibility gains — an entity's mentions in AI answers rising over some period after an intervention — are drawn from a survivor sample.

DOCUMENT	BA-MN-3
KIND	Method note
DATE	2026
SUBJECTS	Measurement & statistics
LICENSE	CC-BY-4.0
AUTHOR	Barkhausen AI

ABSTRACT

Published case studies of visibility gains — an entity's mentions in AI answers rising over some period after an intervention — are drawn from a survivor sample. The cases that regressed or went nowhere are not written up, so a file of success stories carries almost no information about whether the intervention works. Two mechanisms produce winners regardless of any real effect: a portfolio of noisy visibility series will always contain some that rose by chance, and any entity engaged just after an unusually poor measurement will tend to improve on the next one by regression to the mean alone. This note frames the case study as a selection problem, separates it from evidence, and states what would count as evidence instead: pre-registered cohorts, all-entity aggregates, or at minimum a disclosed denominator — of how many tracked entities is this the best case?

CONTENTS

1	The survivor sample	3
2	Chance alone manufactures winners	3
3	Regression to the mean	3
4	What would count as evidence	3
5	Limitations	4

A case study reports one outcome that was chosen for reporting because it was good. That selection step — publish the wins, quietly retire the rest — is what separates a case study from evidence, and it is easy to underrate how much it changes the inference a reader may draw. A file of visibility success stories is consistent with an intervention that works, an intervention that does nothing, and an intervention that on average makes things worse. The format cannot distinguish among these, because in all three worlds the published cases are the ones that happened to rise. This note sets out the two mechanisms that manufacture such cases without any real effect, and what a reader can ask for instead.

1 The survivor sample

Survivorship bias is the error of studying the outcomes that remained visible and treating them as representative of all the outcomes that were attempted. A published visibility case study is a survivor by construction: it exists because its number went up. The measurements that stayed flat, moved with the market, or fell are not fabricated — they are simply not selected for a write-up. A reader who sees only the survivors sees a sample whose selection was determined by the very quantity under study — the condition under which a sample tells you least about the population it came from.

2 Chance alone manufactures winners

Consider a set of entities whose visibility in AI answers is tracked over time. Each series wanders on its own, without any intervention, and by a large margin — day-to-day turnover in cited sources and mentioned brands is documented elsewhere in this publication. Across a portfolio of enough entities, some series will rise substantially over any given window purely by chance, in the same way that some coins in a large enough batch will show a run of heads. Selecting those risers after the fact and presenting them as outcomes demonstrates only that the portfolio was large enough to contain a lucky tail. It says nothing about whether anything caused the rise.

This is why the most useful question to put to a case study concerns its denominator: of how many tracked entities is this the best case? A striking gain selected from three entities is a different claim from the same gain selected from three hundred, and a case study that does not disclose which is not yet a result. The general principle — a percentage is uninterpretable without the base it was drawn from — is the subject of BA-MN-1.

3 Regression to the mean

A second mechanism operates even for a single entity, with no portfolio at all. Parties tend to seek out an intervention precisely when their measurements look unusually bad, and an unusually bad measurement is, in part, the product of transient bad luck rather than a stable level. The next measurement will tend to fall back toward the entity's typical level whether or not anything was done, because the transient component that made the first reading extreme is not repeated. The recovery is then easy to miscredit to the intervention introduced in between. Any before-and-after comparison in which the “before” measurement is what triggered the engagement is exposed to this effect, and a case study is almost always structured exactly that way.

4 What would count as evidence

To move from anecdote to evidence, the selection step has to be removed or, failing that, disclosed. Three things do this, in increasing order of rigor. A pre-registered cohort names the entities and the outcome measure in advance and then reports what happened to all of them, so that failures and non-movers

cannot be dropped after the fact. An all-entity aggregate reports the average outcome across every entity a provider worked on over a defined period, not a hand-picked subset, so that the wins are diluted by the losses exactly as they occur. At minimum — and this is the floor, not the goal — the denominator is disclosed: how many entities were tracked, over what window, and where in that distribution the featured case sits. The window and the sampling behind any such figure still have to be stated for it to be interpretable at all; that minimum disclosure is set out in BA-C-4.

5 Limitations

None of the above says an intervention does not work; it says a survivor sample cannot establish that it does, in either direction. A published case can be perfectly accurate about the entity it describes and still support no general conclusion. The mechanisms here also do not exhaust the ways selection enters a result — outcome measures and comparison windows chosen after the fact compound the same problem — but survivorship and regression to the mean are enough on their own to make an undenominated case study uninformative.

HOW TO CITE

Barkhausen AI (2026). Survivorship bias in visibility case studies. <https://barkhausen.ai/notes/survivorship-bias-in-visibility-case-studies/>

Published under the Creative Commons Attribution 4.0 International (CC-BY-4.0).



Barkhausen AI · Est. MMXXVI · *Every leap begins as a crackle.*