



Spreading Ideas in the New Era

Five steps for discovery and impact, now that Google has told us how it works

A talk for European Christian publishers · Rick Denham

OPENING

They did not invent the press. They understood it first

Mainz

Movable type

Wittenberg

Ninety-five theses

Geneva

A city reformed by print

Antwerp

The printers' port

The printing press did not create the Reformation. But there is no Reformation without it.

Within weeks the theses were read across Germany; within two years, across Europe. Not because Luther was a good marketer, but because a new medium arrived and a handful of people who cared about the truth understood it faster than those who did not.

THE THIRD SHIFT

The press. The internet. And now a system that answers rather than points.



The question this afternoon is not whether that is good. It is whether the answers people receive in French, Czech, Romanian, Dutch, Polish will have been shaped by anyone who holds sound doctrine.

In May 2026, Google stopped the guessing

1

No separate AI index.

AI Overviews and AI Mode are rooted in the core Search ranking systems. Not indexed means invisible.

2

The mechanism is retrieval.

The model retrieves pages from the index, generates an answer with links, and runs a fan-out of related queries.

3

Non-commodity content matters most.

A unique point of view that could not be recycled from elsewhere or produced by a model.

4

No special tricks.

No chunking, no rewriting for machines, no llms.txt, no special markup. Google ignores all of it.

5

Pages made for query variants, not readers, are spam.

Not one page per question, which is a legitimate form; pages that exist because a query exists rather than because an answer exists.

From a link economy to a citation economy

Link economy

You made a page. The search engine sent a visitor. The value was the visit.



Citation economy

The model reads your page, synthesises an answer, and often the reader never arrives.

2.5 billion

people a month see AI Overviews (Google, I/O 2026)

about half

of clicks to top results disappear when an overview appears (independent studies, 2026)

More serious. And more hopeful.

More serious

A model with thin material in your language will not decline to answer. It answers anyway, by generalising from the largest corpus it has: popular American evangelicalism.

A pastor in Cluj asking about assurance in Romanian is not getting silence. He is getting something. The question is whose.

More hopeful

The field is now live in your languages. AI Mode expanded to nearly a hundred languages this year, and the last major European market holding out under press-rights law switched on this month.

Whatever is indexed, original and clear in Hungarian or Czech today is what will be cited.

Romania's problem is not a shortage of sound theology

It is that the sound theology is locked in containers machines cannot open. That is a fixable problem, and this is the whole list:

Not on the open web

Drive links, app walls, logins, scanned images with no text.

Not indexed or crawlable

Blocked crawlers, no sitemap, pages Google has never seen.

Not organised for a reader

One long file with no headings, no sections, no page of its own.

Not original

A translation of something an English site already says better.

A WORD BEFORE THE FIVE STEPS

We borrow the mechanics without borrowing the ending.

Their ending

Every secular framework, and Google's guide, describes means. Their journey terminates at a click or a checkout. Once the transaction closes, the work is finished.

Our ending

A pastor buying a book is roughly the middle of our sequence. What we build is a community of disciples who make disciples, who happen also to sustain a publishing house.

THE FIVE STEPS

Each step rests on something Google said, and on a conviction we already held

1

Be in the index

Containers machines can
open

2

**Write what no
model could write**

What we build; who
publishes

3

**One seed, many
soils**

Sustained across formats

4

**Measure the right
ending**

How long it takes

5

**The handoff to
human community**

Where the strategy
terminates

1 Be in the index

Google's whole eligibility rule: a page must be indexed and eligible to show in Search with a snippet, and the site must be included in generative AI features in Search Console. Nothing exotic.

- Take your best twenty titles. Is each on a public URL a crawler can reach with no login, app or Drive link?
- Prefer real pages over documents. A chapter as a web page with headings beats a chapter as a file.
- Verify the site in Search Console. It is free and now the only official window into AI visibility.
- Check the generative AI inclusion setting. Inclusion is the default. Decide it on purpose (see Step 4).
- Check robots.txt for a Google-Extended block. It does not touch Search, but it removes you from Gemini app grounding.
- Make the site fast on a rural mobile connection, not office fibre.

The trap: assigning this to whoever runs the website. It is a publishing decision about whether the catalogue belongs on the open web.

Write what no model could have written

"Commodity" is not a verdict on translation. It is a technical question: could a model produce this page from what it already knows? A faithful translation, carefully localised, answers that question in your favour.

A bare translation

The model already has the English and can approximate a translation. A page that adds nothing to what the model knows is what Google calls commodity — not a moral verdict, a retrieval one.

Localised illustrations

Replacing a Piper illustration with a French one, drawing on a story your reader recognises, requires being French. A model cannot make that substitution.

A local voice at the top

Two paragraphs from a French pastor explaining why this article matters here, now, in France. That knowledge does not exist in English.

The reader's Bible

Quoting the Second 21 rather than a back-translation of the ESV signals to every system — human and algorithmic — that this page was made for a French reader.

A named French editor

An author page, a French institution, a French context. A model-produced page has no author. Yours does.

French application at the end

Here is what this means for a French evangelical in 2026. This cannot be retrieved from the English corpus because it was never in it.

The test: after localisation, could a model have produced this page from the English original? If no — you have built something Minneapolis cannot replicate.

Answer real questions, under real names

Your best material is a 3,000-word article or a book that few are reading. A pastor or average church member who wants a straight answer is served by neither. A concise answer is a **different product doing a different job**. It is the entrance; the article and the book are the depth.

Real questions

Harvested from conference Q&A, pastoral WhatsApp groups, students, your authors' inboxes. Not from keyword tools. Real people don't ask the same question five ways.

One signed answer each

Two to four hundred words, the answer a pastor gives across a table, with the verses. Written or signed by a named, qualified author with an author page.

A next rung on every page

The article, the chapter, a neighbouring question, and where to buy the book. The purchase link is the rung, not the point; the answer must help someone who never buys.

Why GotQuestions is not spam

A catalogue of answers, not of queries. Distinct question, substantive answer, nothing restated under five phrasings. Copy the substance, not just the format.

AI finds the questions. Your editor owns the answers.

You have a translated book nobody can find online. Here is how to give it a door — and where the line falls between a tool that helps and a shortcut that hurts.

Step 1 — AI reads the book

Feed the book to a model. Ask it: what questions does this book answer? It returns a list of thirty or forty. This is a research job. This is what AI is good for.



Step 2 — Your editor takes over

A qualified person takes each draft answer. Checks every verse. Corrects what the model oversimplified. Adds local application. Signs it with his name. Links to the book.



Step 3 — The book has a door

Each question becomes a page. The answer helps someone who never buys. The page ends with: read more in this book, and here is where to buy it.

The trap:


Skipping Step 2. A model drafts, a quick read-through approves, four hundred books go live. Many pages, little human judgment — **Google calls that scaled content abuse**. Twenty books with a real editor beat four hundred without one.

STEP 3

3

One seed, many soils

You already do this. A strong address becomes a transcript, a chapter, a book, the theme of the next conference. That instinct is right — and it stops halfway. What is missing is the digital continuation of the same motion.



Every piece links to two others, and at least one links to a gathering.

Wikipedia in your language

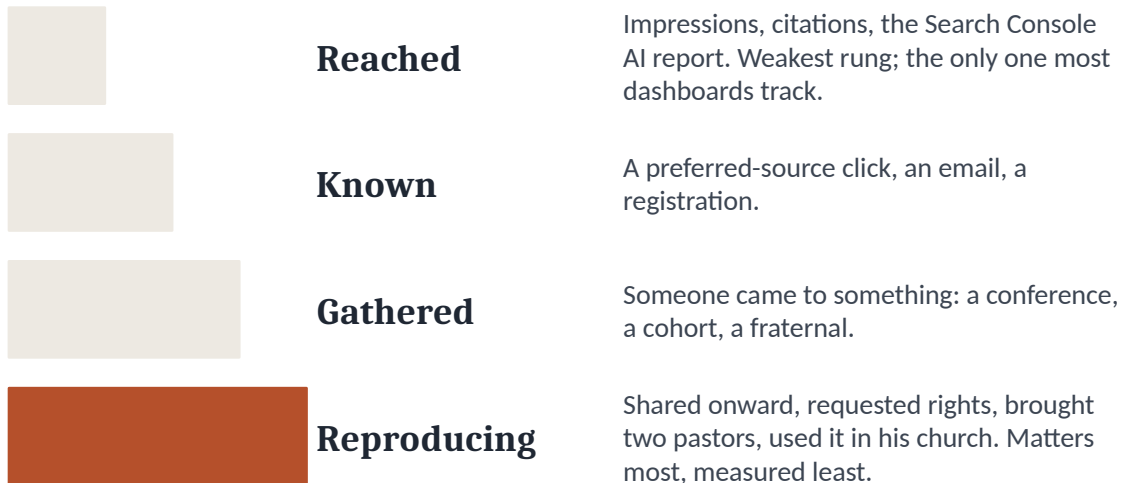
Wikipedia is the reference layer every AI model learns from first. One well-sourced article on justification in Romanian — citing your books — shapes what every system and every reader finds on that subject for years. Almost nobody is doing this. It is the highest-leverage hour your team can spend.

Video

Google says it draws on video and that ordinary video best practice already optimises for AI. A pastor with a phone and a tripod, in his own language, is now a primary source.

Measure the right ending

Adapt quickly in tactics; measure slowly in fruit. Import a conversion dashboard and you build an efficient machine for selling books to people you will never meet.



Two instruments

Search Console's Generative AI performance report: impressions, which pages appear in AI answers, in which countries. Official, free.

The ninety-day test: ten real reader questions, asked in your language of two or three AI systems, on a schedule. Are you cited? Is your terminology intact?

The handoff from digital discovery to human community

Someone finds a clip. Then what? In most of our ministries the honest answer is: nothing. He watches, he nods, he scrolls. There is no path from the reel to the room.

- Every digital asset has a next rung, and the rung is eventually a person.
- Every conference exists online months before and after it exists in a hall.
- Every free resource carries an invitation to a gathering, a cohort, a fraternal.
- Every reader who reproduces is treated as a colleague, not an audience.

Sow abundantly.

His word does not return void.

Soli Deo Gloria

Monday morning, for the person who runs your website

Do, in this order

- Verify the site in Search Console; confirm the generative AI inclusion setting is what the elders decided.
- Check robots.txt for Google-Extended. Decide on purpose.
- Move the twenty most important titles from PDFs onto real pages with headings.
- Add Organization, Person and Book structured data once. Then stop.
- Start the questions section: real questions, one signed answer each, each ending at the article, the book, and the shop. Licence first for invisible books.
- Add the Preferred Sources button where loyal readers land.
- Open the Generative AI performance report monthly, beside the ninety-day test.

Do not

- Chunk pages or write “for the machine.” Google says it does not help.
- Build a page per query variation, or publish drafts nobody edited. That is scaled content abuse.
- Create llms.txt or special AI markup for Google. It is ignored.
- Buy “AI visibility” scores. No third party has Google's data.
- Publish a draft nobody qualified has read.