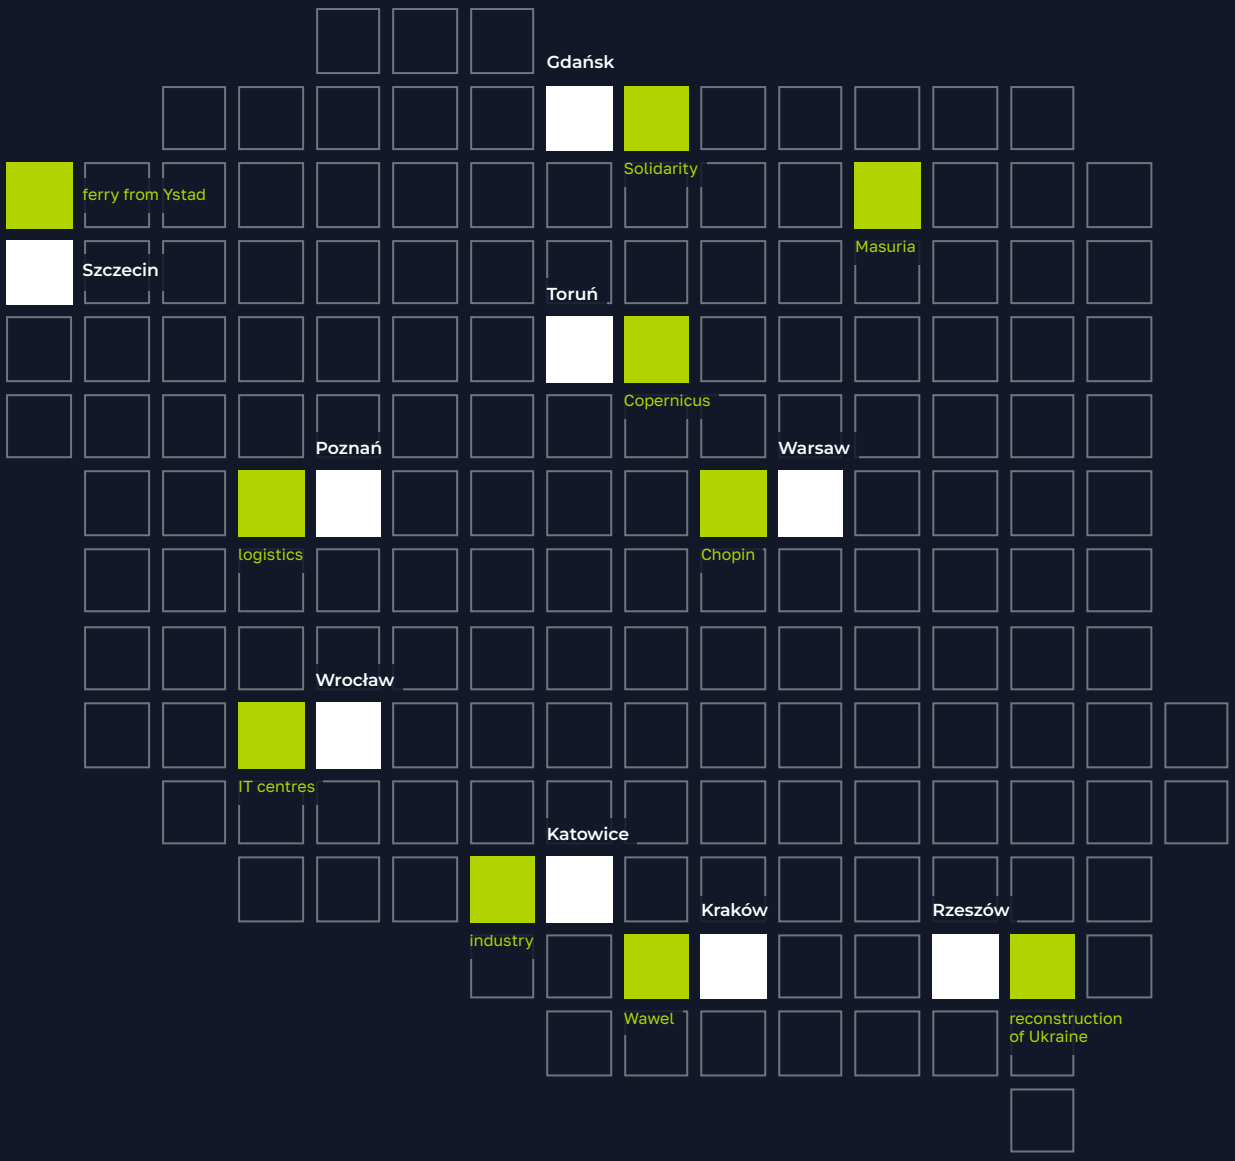


The Generative Portrait of Poland

How six language models talk about Poland
to people and companies from 12 countries

Part I. Poland for People · Part II. Poland for Business

brandsemantics

2026

Knowledge partner

BUSINESS CENTER

Under the patronage of

BCC

The Generative Portrait of Poland

Part I. Poland for People · Part II. Poland for Business

Author: Grzegorz Miłkowski

Publisher: Brand Semantics P.S.A., Warsaw · www.brandsemantics.eu · hello@brandsemantics.eu

Research tool used in the report: Semantio.pro

Knowledge partner: AI Business Center Foundation · aibusiness.center

AI BUSINESS CENTER

Under the patronage of Business Centre Club · bcc.org.pl

BCC

Main study: 25 August – 28 September 2026;
control and additional measurements: 30 September – 2 November 2026.

First edition, Warsaw, 10 November 2026. Polish edition: Polski portret generatywny.

Large language models (LLMs) were used for data analysis, editing and typesetting. Every figure was calculated by script from the study data, and the text was reviewed and edited by the author.

How to cite: Grzegorz Miłkowski (2026). *The Generative Portrait of Poland*. Brand Semantics P.S.A. https://semantio.pro/en/reports/generative-portrait-of-poland/

© Copyright by Grzegorz Miłkowski, 2026

© Copyright for this edition by Brand Semantics P.S.A., 2026

Contents

From the author 4

Summary 6

How this study was made 10

Part I · Poland for People 12

1. Do the models talk about Poland at all? 13
2. Generative visibility, or when the models bring up Poland of their own accord 16
3. Which countries do the models put forward instead of Poland? 20
4. Tone and recommendation rate, or how the models speak of Poland once they do 28
5. Sources, or where the answers send readers for more detail 38
6. Two Polands, or Poland on request and Poland in use 46
7. The models disagree 57
8. Twelve portraits, or Poland market by market 64
9. What this means for promoting Poland 85

Part II · Poland for Business 90

1. Poland in companies' questions, or tone depends on what the company asks 91
2. The shortlist, or do the models put Poland forward to companies of their own accord? 94
3. Which countries does Poland compete with for companies' decisions? 100
4. Two maps, or the decision area chooses Poland's competitors 107
5. Companies on request and companies in use, or the canon that companies never get 113
6. Where answers for companies point, or registers, an agency and the European Commission 121
7. Twelve market cards for companies 130
8. Poland in the shadow of war, or what companies hear about risk 143
9. The models, or the differences a company should know about 147
10. What this means for the economic promotion of Poland 153

About the tool and the author 159

Methodological annex 161

Glossary (Concepts used in the report) 177

References cited in the annex 178

From the author



For many years I helped companies build their own media: blogs, portals, content meant to answer customers' questions before those customers asked them. The principle was simple. People do not buy a product; they buy an answer to their problem. Whoever gives that answer first, and best matches the intent behind the question, wins.

Today that answer increasingly comes from language models (LLMs). Someone in Madrid asks where to go for a cheap autumn weekend. Someone in Kyiv asks which EU country is the easiest place to make a fresh start with their family. Someone in Chicago asks about their great-grandparents' village. None of them visits a website, reads a brochure or phones a travel agent. Each gets a ready-made answer in a few seconds. And in that answer, Poland is either there or it is not.

How language models see Poland

I decided to find out how large language models "see" Poland, what they tell people in different countries about it and how they shape its image. Not with a few questions typed into a chat window, since a conversation like that is a one-off, different for each of us, and can be neither repeated nor counted. I tested it the way a brand and its customers are studied: with the same set of questions, at the same time, across all the leading providers at once. The tool for this kind of measurement, Semantio, was built at Brand Semantics, the company I lead, and it helps businesses measure their chances in the purchase funnel that large language models now shape. This time the brand was Poland, and the customers were personas from 12 countries.

1,400 scenarios and 14,500 answers

I designed the study so that six models would receive 1,400 questions on behalf of individuals and companies from 12 countries, in 11 languages (personas from the United Kingdom and the USA asked in English). Each question is what we call a purchase scenario, written for the study. None of them asked "what do you know about Poland". The questions concerned everyday decisions: where to study, where to move, where to open a warehouse, whether the Polish Baltic coast is cheaper than Spain. Only then did I check how many answers mentioned Poland, alongside which other countries, with what kind of recommendation, and which sources the models cited. A separate measurement asked them about Poland directly, to show what they know but do not use. In all, the study behind this report comes to more than 14,500 answers.

Companies get a different picture of Poland from the models

Part I of the report shows what the models tell individuals, and Part II what they tell companies. Of the 1,400 questions, 700 were asked by company personas: a purchasing department comparing two Polish plumbing suppliers; senior finance management wanting a report on the full costs of a plant; a board choosing a country for a service centre for sixty people. A company asks differently from a person planning a trip, because it builds a criterion and a risk into the question. That is why it gets a different Poland: almost always present, more often recommended, but in a role determined by the decision area. Because, under the method I adopted, Poland was almost always named in the question by the companies themselves, I also added a separate shortlist measurement: 96 questions that do not name it, across eight decision areas, in 12 markets. It is this measurement that shows whether the models put Poland forward to companies when nobody asks about it.

A semantic map of Poland

What is not in this report? There is no ranking of models and no ranking of countries, because whether Poland appears in an answer depends above all on how the question is worded, not on which model answers it. Nor does the report pass judgement. The models do not speak ill of Poland. To individuals they often simply say nothing about it, and when they do speak of it, they mention it rather than recommend it; to companies they almost always speak of it, but in a role determined by the question. For the promotion of Poland, that difference matters more than any average.

What the report does offer is something that cannot be seen in a single conversation with ChatGPT: a semantic map of our country. Which places, people and events (entities, as they are called) the models link with Poland in each market, through which doors these enter the answers, and where those doors are missing.

Enjoy the read!

Grzegorz Miłkowski
author of the report
CEO, Brand Semantics
President of the Management Board, AI Business Center Foundation

Summary

What we studied. The study covers six generative AI models queried through their official APIs: ChatGPT, Gemini, Claude, Perplexity, Grok and DeepSeek. Between 25 August and 28 September 2026 we put 1,400 questions to them in 11 languages, 700 on behalf of individuals and 700 on behalf of companies from 12 markets: China, Czechia, France, Germany, Italy, Norway, Slovakia, Spain, Sweden, Ukraine, the United Kingdom and the United States. The questions were asked on behalf of personas, scenarios prepared for the study, and concern everyday decisions: where to go for a weekend, where to study, where to move, where to build a plant, whom to contract production to, how to vet a business partner. That makes 8,400 answers; each was assessed by the same model judge, and the results were checked through the five control layers described in the annex. Part I reads these answers from the perspective of individuals, Part II from that of companies (700 questions, 4,200 answers). Because companies almost always named Poland in their questions, Part II also includes a separate shortlist measurement of 96 questions that do not name Poland, across eight decision areas (576 answers, 30 September). There are also two thematic measurements: Poland on request (the largest companies and industries) and Poland in the shadow of the war. The research programme as a whole comprises more than 14,500 answers from language models (LLMs). The subject of the report is the brand semantics of Poland: the entities (places, people, events, traits) that the models link with the name “Poland”, and the ways those entities connect to the situations that people and companies ask about.

Part I. Poland for People

1. Poland is in the answer when it is in the question. Poland appears in 95.1% of the 8,400 answers. What that figure shows is that the models stay on topic: when a question names Poland, they talk about it in 99.9% of answers. The real test lies in the 1,452 answers to questions that do not name Poland. There, Poland appears in 72.4%.

2. Visibility depends on the type of question, not on the model. Across 174 questions in which individuals are looking for a destination without naming Poland, the models mention it in 71.3% of answers. Where the question points to a region that includes Poland, the figure is 91.7%; where it concerns Europe in general, 45.1%; where it refers neither to a region nor to a trip, 20.0%. In the models’ “default Europe”, Poland appears in fewer than half of the answers, and without Ukraine, for which it is the obvious destination, in fewer than four in ten (38.5%). The differences between models are smaller than those between question types.

3. Poland competes with a list, not with a single country. In answers to questions that do not name Poland, the models list almost six countries on average. In questions about a region, Poland comes first (91.7%), ahead of Czechia (80.6%) and Hungary (73.2%). In questions about Europe in general, Germany heads the list (58.0%), while Poland (45.1%) sits in a cluster with Czechia, Portugal and Spain. One answer in five (21.1%) names the persona’s home country.

4. Mentioned more often than recommended. When a model brings Poland into an answer of its own accord, it explicitly recommends it in 55.9% of cases; it almost never advises against it (1.7%), but in four cases out of ten it mentions Poland without recommending it. Here the question type works

the opposite way from visibility: in questions about a region, Poland appears almost every time, yet is recommended in only 47.0%; in questions about Europe in general, it appears less often, but when it does, it is recommended in 62.3%. The model brings it in when it has an argument for it. The models' tone about Poland is moderately positive (6.95 on a 1–10 scale) and almost identical across all markets; the models differ more in recommendation rate than in mention rate. Asked directly about safety during the war in Ukraine (a separate measurement taken in October), in none of 792 answers do the models advise against Poland as a destination, but they almost always add a "but"; asked which countries in Europe as a whole are safe, they place it in the "with caution" group in between half and three quarters of answers.

5. Polish URLs come from the model itself more often than from search. Only one of the six models, Perplexity, searched the web live; the others supplied URLs themselves. Half of those URLs lead to Polish sites (49.5%) and one in three to official ones (33.5%). Sources from search account for 57% of all citations, but only 22.4% of them lead to Polish sites (6.5% to official ones). So Polish institutions are present in the models' answers, but search rarely leads to them.

6. The canon appears on request, but is almost absent from practical answers. Asked directly about Poland, the models answer with the canon: Chopin, Wałęsa, Nicolaus Copernicus and John Paul II appear in every one of the 144 answers about the best-known Poles. In practical questions about travel, history and culture, the same canon barely features: from 12.0% of answers (Chopin) down to 0.8% (Skłodowska-Curie). We call this difference the activation gap. The canon enters practical answers only through the doors the question opens: a place to see (Chopin appears together with Warsaw or Żelazowa Wola in 97% of his mentions) or a chapter of history. The Polish Tourism Organisation (POT) campaign "Poland. More than you expected" builds on faces from the canon and on faces that are absent from the answers; the measurement was taken eight weeks after the campaign's launch and does not assess its effectiveness.

7. The models disagree, and none of them is "the best for Poland". The six models agree on whether to talk about Poland in 93.4% of the 1,400 questions; the disagreement lies in the questions that do not name Poland. Where they do disagree is on how to talk about it: on the same question, the warmest and the coolest model differ by 1.8 points of tone on average, which is roughly twice what measurement variability alone would produce. How the models line up depends on what we count: five of the six come first or last on at least one measure. Newer versions of the models shifted the share of explicit recommendations of Poland by a dozen or so points within seven weeks, in both directions. Manual verification confirmed 14 of the 19 factual errors flagged by the judge in 8,400 answers; errors the judge did not flag were not estimated.

8. Twelve portraits, no ranking. Chapter 8 reads the results market by market, using the same market card for each and without placing the markets in any order: Poland's visibility depends on the question scenarios prepared for a given market. Each portrait closes with a market-specific thread, that is, an association peculiar to one market (such as the Battle of Britain for the British or the Tatras for Czechs), and a check on whether it makes its way into practical answers. Usually it does not; the Tatras and Gdańsk do, because they are places one can travel to. In all 12 markets, the answers also point to the authorities of the persona's home country. In the British market, gov.uk appears more often than gov.pl.

9. What this means for the promotion of Poland. The goal is not "for the models to recommend Poland" but presence in specific types of question. Where Poland appears rarely, it is recommended more often, so promotion should supply the models with arguments rather than presence. The canon is capital that can be activated immediately, but only through doors: linking an entity with a place, a route, a decision situation. A Polish source makes it into the answer when it addresses the situation of a person from a given market, in that person's language or in English. In questions about the cost of

living in Poland, answers currently point to the numbeo.com website, while the tourism portal appears in 4.6% of answers, almost exclusively in English. Poland's image in the answers changes faster than campaigns do; this semantic drift has to be measured at regular intervals, with the same set of questions, because a one-off audit shows only how things stood.

Part II. Poland for Business

1. The models almost always talk to companies about Poland, and they describe it more coolly than they do to individuals. Poland appears in 97.4% of the 4,200 answers. When a company does not name it, the model brings it in of its own accord in 75.2% of answers and explicitly recommends it in 65.1% of those. The tone is lower than for individuals (6.78 against 7.12), but for one reason: 281 of the 700 company questions themselves ask about risk, barriers and pitfalls. In those, the tone is 6.16, one answer in eight is negative and Poland is recommended in one in two; in the remaining questions the tone is 7.20, negative answers account for 1.6% and the recommendation rate is 79.2%. The model answers the question it is asked.

2. Poland is always on the shortlist; the decision area determines its role. When a company asks without naming Poland, the models mention it in 98.4% of the 576 answers and explicitly recommend it in 82.5%. What matters is the decision area: in questions about a sales market, Poland is recommended in 93.1% of answers; about engineers, in 93.0%; about contract manufacturing, in 69.4%; about a warehouse for Europe, in 62.3%.

3. The competitors are the same in both sets. Czechia (79.2% of shortlist answers), Romania (74.1%), Hungary, Germany and Slovakia stand in the same order in both the study and the shortlist. Romania, which barely featured in individuals' questions, is the second candidate for companies. In questions about Europe in general, Poland comes first for companies (69.5% against 55.7% for Germany), unlike for individuals, where Germany headed the list. This is because a company puts a criterion into its question (cost, engineers, distance from Germany), and Poland meets it.

4. Poland loses on geography, not on cost. The decision area chooses Poland's competitors. In questions about sales markets, incentives, service centres and nearshoring, its competitors are its regional neighbours, and the recommendation rate reaches 81–93%. In questions about suppliers, manufacturing and logistics, Germany, Italy, Austria and the Netherlands join the list, and the recommendation rate falls. In answers about a warehouse for the whole of Europe, the models know Polish locations (in 53 of 72 answers), but the hubs are the Netherlands and Germany, because from there goods reach the west in two to three days; cost is Poland's argument in every one of these answers. To suppliers, the models describe Poland as a base that is growing, and Germany as reliable.

5. The corporate canon appears on request, but in use the registers win. Asked about the largest Polish companies, the models name KGHM in every one of the 144 answers, and PGE and PKO Bank Polski in 143. In answers to companies' questions, KGHM appears in 21 of 4,200 and PZU in one; the leaders are Krajowy Rejestr Długów, a Polish debt register (143 answers), BIG InfoMonitor, a credit-information bureau (97), and PSE, the Polish transmission-system operator (95), because they address the company's situation. A company that asks for names in an industry gets a Polish company in two answers out of three (99 of 150), but when it asks for certified subcontractors, only in 5 of 30. Aviation, electronics and pharmaceuticals appear rarely even on request (in one in four, one in six and one in fourteen answers about the strongest industries), and the study asked no questions about companies in these sectors.

6. Polish institutions are in the models' answers, but search in the market's language rarely leads to them. Of the URLs the models give companies themselves, 54.5% lead to Polish sites (35.0% to official ones); of the sources from live search, 25.0% (6.9%). The Polish Investment and Trade

Agency (PAIH) is among the three most-cited domains in each of the 12 markets, and in questions about incentives it appears in 26 of 27 answers with URLs given in the answer, and in 3 of 12 with sources from search.

7. Twelve cards, no ranking. In every market, Poland appears on the shortlist in 95.8–100% of answers; what differs is who stands alongside it, and whether the persona's home country is named (Czechia 75.0%, Germany 72.9%, China and the United States almost never). Three markets have a thread of their own: the models add the Baltic ferry for Swedish companies of their own accord (87 of 348 answers; no question asks for it), while the China–Europe railway appears for Chinese companies mainly when the company asks about it (77 of 98), and for German companies the models point to the bilateral chamber ahk.pl of their own accord.

8. The war: Poland recommended with conditions. Asked directly about investment, nearshoring, energy or insurance in Poland, the models did not advise against Poland as a choice because of the war in any of 864 answers, but only 3.8% of answers to questions about investment risk recommend it without any caveat, and some answers point boards that are sensitive to risk to a country further west. About two thirds of answers present Poland as more exposed than the countries it is compared with, and nine in ten of those recommend it all the same. As a base for the reconstruction of Ukraine, Poland is recommended in every answer; as a safe country in the region, when the question does not name it, it is placed in the "with caution" group in 37.5%. Companies hear different arguments from individuals: the European Union (86.9% of answers) and energy (75.0%) instead of Belarus and ministry warnings.

9. The models differ, but every one of them warns when a company asks about risk. The six models agree on whether to talk to companies about Poland in 95.0% of the 700 questions; in the 68 questions that do not name Poland, they bring it into between 57.4% and 83.8% of answers, and the difference comes from the way they answer (a procedure rather than a list), not from how they assess Poland. On the same question, the warmest and the coolest answer differ by 1.8 points of tone on average. When a company asks about risk, every model lowers its tone by 0.9–1.3 points; where they differ is in how often they turn a list of risks into a negative answer (from 6.1% to 17.5%).

10. What this means for economic promotion. The goal is not "for the models to recommend Poland" (they already do) but the decision area in which Poland is to be a candidate rather than a backdrop. In industrial areas, Poland needs evidence of maturity, not of growth: certifications, references, the names of subcontractors in the languages of the markets. An argument makes it into the answer when it is linked to a decision situation (the ferry), not when it is merely an association (the railway). Content for models that use search has to be citable in the market's language, because PAIH is in the models' answers, not in the search results. Poland's image in the answers for companies is a flow, not a state, so it requires control measurements at regular intervals, with the same set of questions.

What the report does not say. It measures neither what the consumer app shows, nor website traffic, nor the effectiveness of any campaign. It does not forecast investment decisions and does not assess the work of any institution. It does not rank markets or models. It does not combine figures from the study with those from the shortlist, because these are two sets with separate denominators, and the shortlist is not a check on the study. It does not compare individual models' recommendation rates in the main text, because four of them were subject to a length cap on answers in the study (annex). It is a dated snapshot, a map of the representation of Poland in the answers of six models to questions from individuals and companies in 12 markets, and it should be read as such.

How this study was made

Who asked. The questions were asked not by people but by personas: characters described for the study, from 12 countries (China, Czechia, France, Germany, Italy, Norway, Slovakia, Spain, Sweden, Ukraine, the United Kingdom and the United States), each with a life or work situation of its own: a parent looking for a university for their child; a couple planning an autumn weekend; a purchasing department comparing suppliers. Each persona "spoke" in the language of its market. The list of personas is in the annex, and the full set in the open data.

What they asked about. Everyday decisions, not Poland. No question read "what do you know about Poland?". The questions were about where to go, study, live, or open a warehouse or a company. These are the questions people type into a chat window today. Some questions named Poland and some did not, and this distinction is the most important one in the report: only in questions that do not name Poland can one see whether the model brings it in of its own accord. The questions were prepared by the research application's question generator from a description of the market and the personas; after the measurement we reviewed them and set out in the annex which ones repeat and which are weaker. In total, 1,400 questions in 11 languages: 700 on behalf of individuals (Part I) and 700 on behalf of companies (Part II). Because companies almost always named Poland in their questions themselves, Part II also includes a separate shortlist measurement: 96 questions that do not name Poland, across eight decision areas, from the location of a plant to a warehouse for Europe (576 answers, 30 September 2026).

Whom they asked. Six language models (LLMs): ChatGPT, Gemini, Claude, Perplexity, Grok and DeepSeek. The questions went through their official APIs, the same channel that apps and companies use, within the same time window: between 25 August and 28 September 2026. We measure the model's answer returned via the API, not what a person sees in the consumer app. The report gives the model versions, because answers change along with the versions.

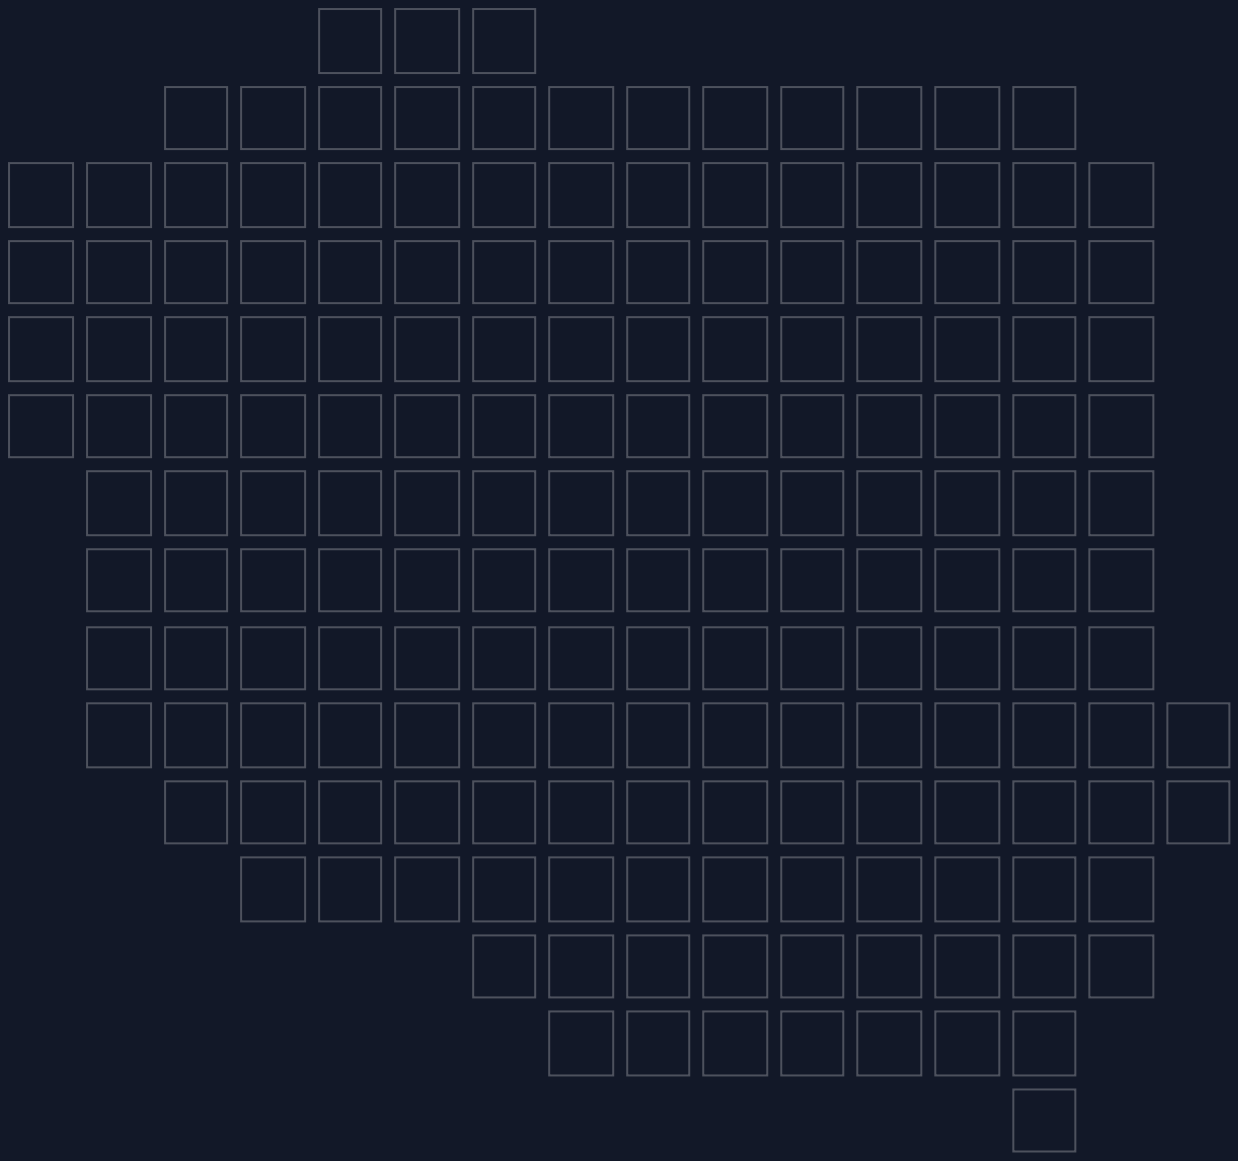
How the answers were read. The 8,400 answers are too many for people to read, so each one was assessed by the same model judge following the same instructions: whether Poland appears in the answer, in what tone it is discussed (on a scale from 1 to 10), whether it is explicitly recommended, which other countries stand alongside it and which sources the answer refers to. The judge's assessment was checked through five control layers: a repeat of the measurement on the same questions, a second judge on a model from a different provider, human verification of every flagged factual error, a newer-model-version test and a separate "open portrait", in which the models were asked about Poland directly. The results are probabilistic: a language model may answer the same question differently, so the report gives shares of thousands of answers within a dated window, not a forecast of what a model will say in a single conversation; the spread between repetitions is measured by the repeatability test.

What the numbers mean. Four indicators recur in every chapter: mention (the share of answers in which Poland appears), tone (the average on a 1–10 scale), recommendation (the share of answers in which Poland is explicitly recommended) and the share of Polish sources (the proportion of URLs in the answer that lead to Polish sites). Differences smaller than 3 percentage points for mention, 0.5 points for tone and 5 percentage points for recommendation fall within measurement variability and

are not treated as findings in the report. Markets and models are not ranked, because the result depends above all on how the question is phrased.

What this study does not say. It does not say what people in the 12 countries think, only what the models answer about Poland to questions written on their behalf. The questions are scenarios prepared for each market, not a record of what people there ask most often. It measures neither website traffic nor the effectiveness of any campaign. It is a dated measurement: a snapshot from autumn 2026 that makes sense only with its date.

Where to find the details. Each chapter has a “How we count” box defining the measures used in it. The full method is set out in the annex at the end of the report: the judge’s instructions, the model versions, the length caps on answers and their effect, the list of personas, the post-measurement question review, all the control layers and the file checksums. The questions and personas in the languages of the markets are in the open data, so the measurement can be repeated.



Part I

Poland for People

The Generative Portrait of Poland

CHAPTER 1

Do the models talk about Poland at all?

IN BRIEF

- Poland appears in 95.1% of the 8,400 answers given by six language models (LLMs): in 93.9% as a country, and in 1.2% only by way of a Polish city or region.
- When a question names Poland, the model almost always talks about it (99.9% of 6,948 answers). This shows that the models stay on topic, not that they bring up Poland of their own accord.
- The real test is the 1,452 answers to questions that do not name Poland. Poland appears in 72.4% of them. Chapter 2 breaks this result down.

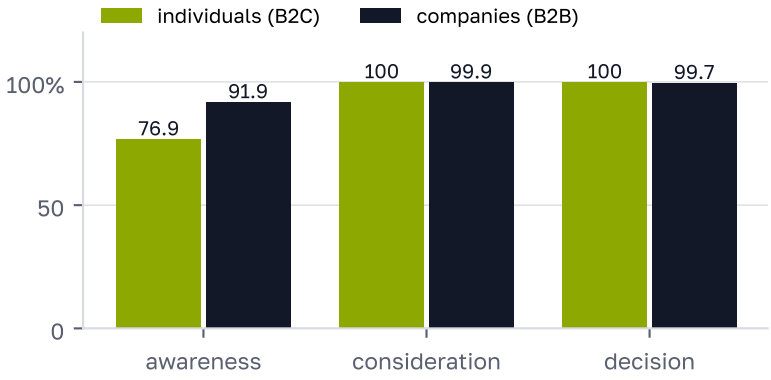
The first question this report asks is the simplest possible: does Poland come up in the models' answers at all? The study comprised 1,400 questions that residents of 12 countries might ask before a trip, a course of study, a move or a business decision. Each was put to six models. Poland appears in 95.1% of the answers.

What the models say about Poland is what we call in this report the brand semantics of Poland: which places, people, events and traits they link with its name, and in which situations they mention it. The report measures it chapter by chapter, from the bare mention through to the sources the models draw on.

The 95.1% figure is impressive, but it needs to be read alongside the questions. In 82.7% of the questions, Poland or a Polish city or region is named: someone is planning a semester in Kraków, looking for work in Poland or comparing warehouse costs near Poznań. In such questions the model talks about Poland in 99.9% of answers. It could hardly be otherwise.

Decision path. At the consideration and decision stages, Poland is usually already in the question, so the mention rate is between 99.7% and 100%. It falls only at the awareness stage, when someone is still looking for a destination and some questions point to no country at all. There it drops to 91.9% in business questions and to 76.9% in questions from individuals. For the same reason, the mention rate is higher in B2B questions (97.4%) than in B2C (92.9%).

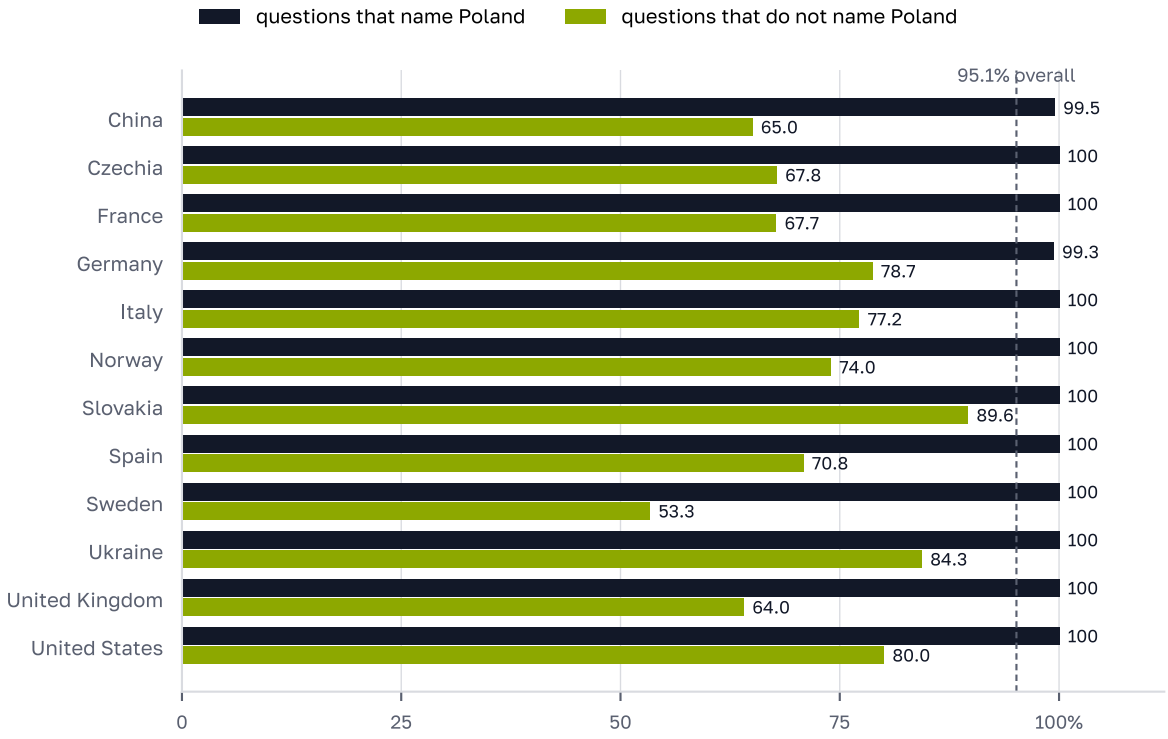
Chart 1.1 Decision path



Markets. The mention rate ranges from 91.1% (France) to 98.6% (Slovakia). France and Sweden are significantly below the average, Slovakia above it. Does this mean the models "know Poland less well" in France? No. In questions that name Poland, every market scores between 99.3% and 100%. The difference comes from the questions that do not name Poland:

- in France they account for more than one question in four (27.6%);
- in Sweden, Poland appears in answers to such questions in only 53.3% of cases, against 89.6% in Slovakia.

Chart 1.2 Poland in the answers: questions that name Poland and questions that do not, 12 markets



Models. The mention rate ranges from 94.1% (Perplexity) to 95.9% (DeepSeek), and no model deviates from the average by as much as 3 percentage points. In questions that name Poland, all models score between 99.7% and 100%. Differences between models emerge only where a model has to decide for

itself whether to mention Poland. In questions that do not name Poland, the mention rate ranges from 66.5% (ChatGPT and Perplexity) to 76.9% (DeepSeek).

HOW WE COUNT. A mention means that the answer names Poland or a Polish city or region. The denominator is all assessed answers (8,400, none missing). The measurement concerns the model's answer as returned by the API, not what is shown in the consumer app. We do not interpret differences in the mention rate smaller than 3 percentage points (annex: repeatability test).

CHAPTER 2

Generative visibility, or when the models bring up Poland of their own accord

IN BRIEF

- Across 174 questions in which individuals are looking for a destination without naming Poland, language models (LLMs) mention Poland in 71.3% of answers (1,044 answers).
- The result depends above all on how the question is phrased:
  - when it points to a region that includes Poland (Central or Eastern Europe, neighbouring countries), Poland appears in 91.7% of answers;
  - when it concerns Europe in general, in 45.1%;
  - when it refers neither to a region nor to travel abroad, in 20.0%.
- The figure for general questions is the most telling. When the question does not narrow the choice, Poland makes it into the models' "default Europe" in fewer than half of the answers. The exception is Ukraine, at 77.8%; excluding it, the figure is 38.5%, not quite four answers in ten.

Two questions, the same decision stage, the same model. In questions like "I'm looking for a cheap city in Central Europe for a weekend", Poland appears in more than nine answers out of ten (91.7%). In questions like "We want to move somewhere in Europe with the children for a few years", it appears in fewer than half (45.1%). The whole difference this chapter deals with lies between those two sentences. Whether a model brings up Poland of its own accord is decided not by the model, nor by the market, but by how the question is phrased.

2.1 What this chapter measures

Generative visibility (GV) is the share of answers in which the model mentions Poland of its own accord even though the question does not name it. The report calculates it on awareness-stage questions: someone is planning a cheap weekend away, a semester abroad, a move or dental treatment outside their own country, and asks where to go. There are 174 such questions across 12 markets, each put to six models.

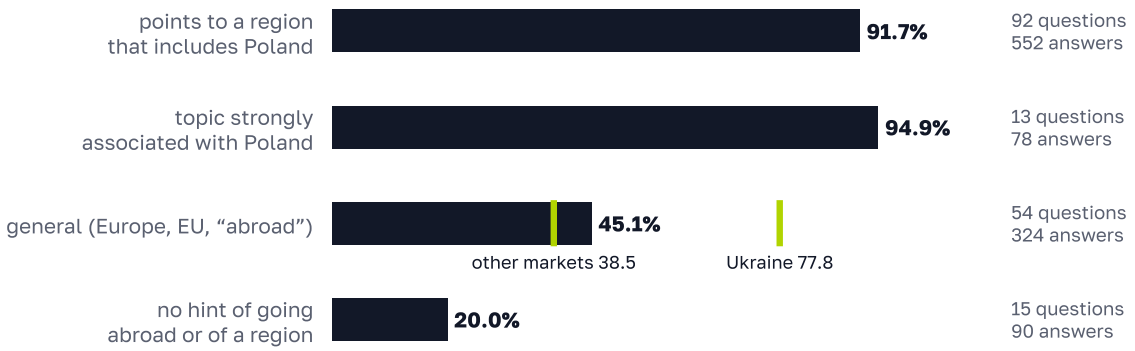
The questions are not all alike, however. Some narrow the choice straight away to the region in which Poland lies ("I'm looking for a cheap city in Central Europe for a weekend"). Some concern a topic strongly associated with Poland, such as Second World War memorial sites or dental treatment abroad. Some ask about Europe in general ("We want to move somewhere in Europe with the children for a few years"). And some refer neither to a region nor to a trip ("which regions are good for a short

cycling break"). That is why we report GV together with a breakdown by question type. Without that breakdown, a single figure would say little.

2.2 Question type decides

question type	questions	answers	Poland in the answer
points to a region that includes Poland	92	552	91.7%
topic strongly associated with Poland	13	78	94.9%
general (Europe, EU, "abroad")	54	324	45.1%
no hint of going abroad or of a region	15	90	20.0%
total (GV B2C)	**174**	**1,044**	**71.3%**

Chart 2.1 Question type decides



When the question itself points to a region, the models almost always name Poland within it. That is good news, as it means Poland is a fixed feature of the models' map of Central and Eastern Europe. The topical questions show, in turn, that the models strongly associate Poland with Second World War history and memorial sites, as well as with medical services for people from abroad.

The third row, however, is the most telling. When the question gives no hint of a region, Poland appears in 45.1% of answers, and in 38.5% excluding Ukraine. These figures capture Poland's place in the models' "default Europe", the list a model draws up for itself when the question simply says "somewhere in Europe".

The fourth row covers questions that do not refer to a trip at all: someone asks about "regions for a weekend", "nature areas for a longer journey" or "the culture of this country", without naming either Europe or anywhere abroad. The models read these as questions about the persona's home country (chapter 3), so Poland appears in one answer in five. These questions remain in the indicator, because a definition is not changed after the results have been seen. Without them, GV would be 76.1%.

All six models behave similarly:

- in questions that point to a region, they name Poland in 88.0–94.6% of answers;
- in general questions, in 35.2–51.9%;
- in questions with no hint of going abroad, in 13.3–26.7%.

2.3 Markets, or question type first, country second

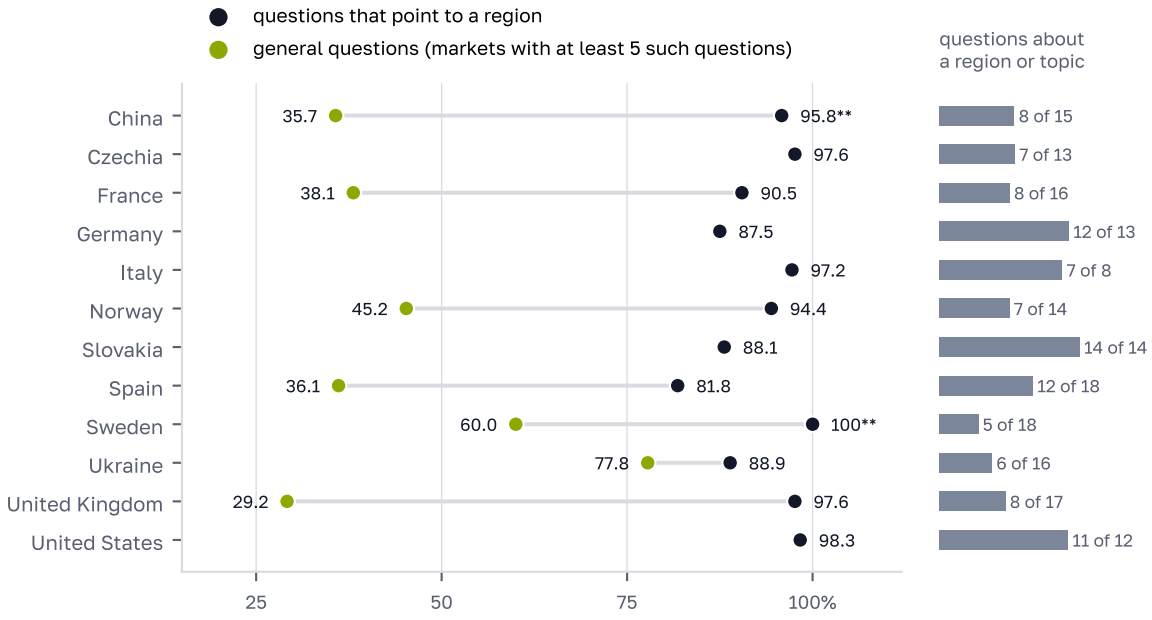
GV ranges across markets from 48.1% (Sweden) to 95.8% (the United States). Does this mean the models “see Poland” better in the United States than in Sweden? The comparison cannot be read that way. In the United States, 11 of the 12 questions point to Central Europe or to a topic associated with Poland; in Sweden, only 5 of 18 do. The share of such questions explains most of the differences between markets: the more questions about a region, the higher the visibility (a correlation of 0.79, i.e. a strong relationship). That is why the report compares markets only within the same question type.

market	questions	of which region or topic	GV	GV: questions that point to a region	GV: general questions*
China	15	8	67.8%	95.8%	35.7%
Czechia	13	7	62.8%	97.6%	–
France	16	8	61.5%	90.5%	38.1%
Germany	13	12	82.1%	87.5%	–
Italy	8	7	89.6%	97.2%	–
Norway	14	7	70.2%	94.4%	45.2%
Slovakia	14	14	88.1%	88.1%	–
Spain	18	12	67.6%	81.8%	36.1%
Sweden	18	5	48.1%	100%**	60.0%
Ukraine	16	6	83.3%	88.9%	77.8%
United Kingdom	17	8	59.8%	97.6%	29.2%
United States	12	11	95.8%	98.3%	–

* Only markets with at least 5 general questions. With 5–9 questions (30–54 answers), we treat the results as indicating a tendency.

** Sweden has only 2 questions that point to a region (12 answers); China has 4 (24 answers).

Chart 2.2 Markets within question type



Two things emerge from the table. First, **in questions that point to a region, Poland is present everywhere**: in every one of the 12 markets; in markets with at least 6 such questions, in 81.8–98.3% of answers. Second, **in general questions, one market stands out.** When a question about studying, working or moving "in Europe" comes from Ukraine, the models name Poland in 77.8% of answers. In the other markets the figure ranges from 29.2% (the United Kingdom) to 60.0% (Sweden). For someone asking from Ukraine, Poland is the default destination in Europe. In the British, Chinese or Spanish market, it is one option among many, named by the model in one answer in three.

2.4 Models differ less than questions do

On the same 174 questions, GV ranges from 67.8% (ChatGPT) to 74.7% (DeepSeek). ChatGPT is 3.4 percentage points below the average and DeepSeek 3.4 percentage points above it. Claude (73.6%), Gemini (70.7%), Grok (70.7%) and Perplexity (70.1%) fall within measurement variability.

The differences between models are clearly smaller than those between question types. For Poland's visibility, then, how a question is phrased matters more than which model answers it.

2.5 How it sounds in the answers

Three examples from general questions that do not point to a region. They show how Poland enters the answers: through price, through its cities and, in the case of Ukraine, as the first choice.

> **"Gdańsk og Warszawa – Polen er generelt veldig rimelig, god mat og kultur"**
>
> "Gdańsk and Warsaw – Poland is generally very affordable, with good food and culture"
>
> — Claude, Norway, question about popular, inexpensive cities for a weekend (Log ID c55bcb91-502b-4c44-b786-4f384f210e0c)

> **"Si buscas coste inferior al de España, buena acogida a extranjeros y cierta comunidad española, priorizaría Portugal, Polonia, Bulgaria, Rumanía y Grecia."**
>
> "If you are looking for lower costs than in Spain, a warm welcome for foreigners and something of a Spanish community, I would prioritise Portugal, Poland, Bulgaria, Romania and Greece."
>
> — ChatGPT, Spain, question about moving for a few years (Log ID 5589726d-58dc-4e02-9e81-bac89fd82685)

> **"Для українців найдоступніші варіанти в Європі зазвичай — це Німеччина, Польща, Чехія, Словаччина, Австрія, Греція та частково Фінляндія"**
>
> "For Ukrainians, the most accessible options in Europe are usually Germany, Poland, Czechia, Slovakia, Austria, Greece and, to some extent, Finland"
>
> — Perplexity, Ukraine, question about studying abroad (Log ID 5de53bda-104a-4d3f-a8fa-01d928bf949a)

HOW WE COUNT. GV is the share of answers naming Poland (or a Polish city or region) in response to awareness-stage questions that do not contain the name of Poland, or of any Polish city or region, in any grammatical form. In Part I, GV refers to questions from individuals (B2C), because the study's question set contains too few such business questions per market. Visibility in business questions is covered in Part II, on the basis of a separate measurement. We assigned a question type to each question solely on the basis of its wording and the country it came from; the classification was checked in two independent runs without access to the results (annex). We do not report results for individual questions.

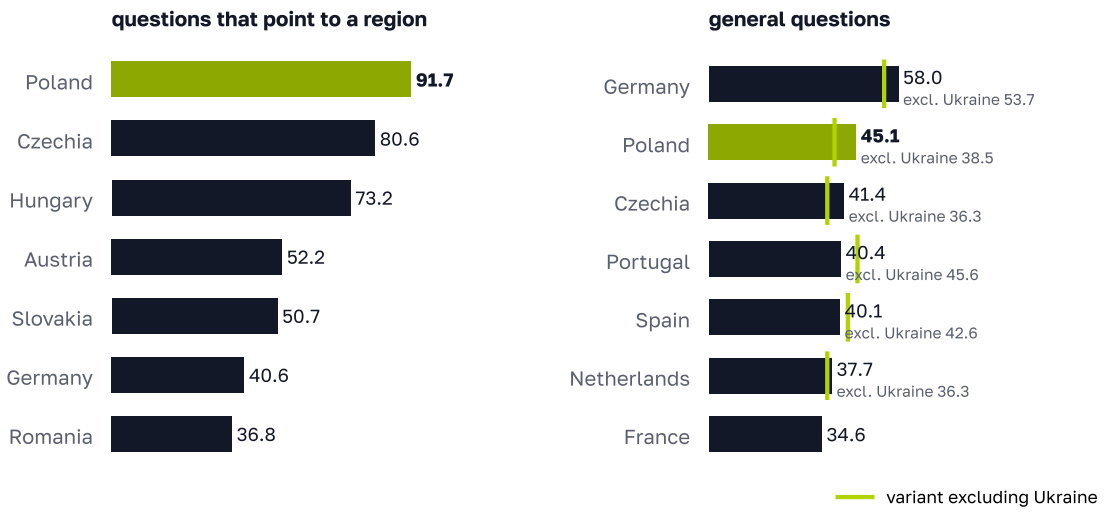
CHAPTER 3

Which countries do the models put forward instead of Poland?

IN BRIEF

- In 1,044 answers to questions from individuals that do not name Poland, language models (LLMs) name almost six countries on average. Poland is therefore competing not with a single country but with an entire list. That list is Poland's competitive set in brand semantics: the countries the models place alongside it in the same decision situation.
- Which countries make the list depends on the question type. When the question points to a region, Poland comes **first** (91.7%), ahead of Czechia (80.6%), Hungary (73.2%), Austria (52.2%) and Slovakia (50.7%). When the question is about Europe in general, Germany heads the list (58.0%) and Poland (45.1%) is second, ahead of Czechia (41.4%), Portugal (40.4%), Spain (40.1%) and the Netherlands (37.7%).
- In questions from Ukraine, "default Europe" means Germany (79.6%) and Poland (77.8%). Excluding Ukraine, Poland falls to 38.5% in general questions and joins the group of countries named in fewer than four answers in ten, behind Germany (53.7%) and Portugal (45.6%).
- The home country, meaning the country the question comes from, appears in one answer in five (21.1%). This is usually because the question says nothing about going abroad, or because the model uses that country as a point of reference. The report does not measure whether the model "advises against going abroad".

Chart 3.1 Two leagues



A model's answer seldom points to a single country. Asked about a cheap weekend, a semester abroad or a move, a model answers with a list: on average 5.8 other countries alongside Poland, and Poland is the only country in the answer in just 0.4% of cases. Poland, then, is not competing with Czechia or with Germany. It is competing with an entire list, which is why this chapter asks: is Poland on that list, and how high up?

3.1 What this chapter measures

Chapter 2 asked whether Poland appears in the answer at all. This chapter asks alongside whom. For each of the 1,044 answers to 174 questions that do not name Poland (awareness stage, individuals segment, 12 markets, six models), the judge listed every country mentioned. The chapter calculates the share of answers in which each country appears and compares it with the share of answers in which Poland appears.

Only 5.7% of answers name no other country; in the rest, Poland appears alongside other countries, and it is those countries that the chapter counts.

The chapter uses two supplementary measures:

- **Poland's lead** is the difference, in percentage points, between the share of answers that include Poland and the share that include the most frequently named other country (in the annex this measure is called the recommendation gap). A positive value means that Poland appears more often than any competitor; a negative value means that another country appears more often than Poland.
- **Home country in the answer** is the share of answers that name the country from which the question was asked (e.g. Czechia in Czech questions). In the study this is the country of the persona on whose behalf the question was asked, not of a real person (annex). We discuss how to interpret this measure separately in 3.4, because mentioning the home country is not always the same as recommending it.

3.2 Two leagues, or the region and "default Europe"

question type	questions	answers	Poland	most frequently named other countries	Poland's lead (pts)
points to a region that includes Poland	92	552	**91.7%**	Czechia 80.6 · Hungary 73.2 · Austria 52.2 · Slovakia 50.7 · Germany 40.6	+11.1
topic closely linked to Poland	13	78	**94.9%**	Hungary 73.1 · Czechia 67.9 · Germany 66.7 · France 46.2 · Spain 42.3	+21.8
general (Europe, EU, "abroad")	54	324	45.1%	**Germany 58.0** · Czechia 41.4 · Portugal 40.4 · Spain 40.1 · Netherlands 37.7	−13.0
no hint of going abroad or of a region	15	90	20.0%	Czechia 30.0 · Germany 24.4 · Sweden 22.2 · Norway 21.1	−10.0
total	**174**	**1,044**	**71.3%**	Czechia 63.1 · Hungary 52.6 · Germany 46.6 · Austria 39.2 · Slovakia 30.5	+8.1

We treat the country shares in the "no hint of going abroad" row as a tendency (90 answers).

Poland plays in two different leagues.

In its region, Poland wins. When the question is about Central Europe, Eastern Europe, neighbouring countries or the Visegrád Group, Poland appears in 91.7% of answers, 11 percentage points more often than Czechia, its most frequent companion. The set of competitors here is stable: Czechia, Hungary, Austria, Slovakia, Romania (36.8%) and Slovenia (25.9%). This is the "model Central Europe" list, and

Poland heads it. In topical questions (memorial sites, dental treatment) the lead is wider still: Poland 94.9% against 73.1% for Hungary.

In “default Europe”, Poland is one of many. When the question is about Europe without specifying a region, it is not only Poland’s result (45.1%) that changes but the whole list. Germany moves into first place (58.0%), and Czechia is joined by Portugal, Spain, the Netherlands, France (34.6%) and Italy. Poland is second in this league, but a group of four countries in the 38–41% range sits close behind it. In general questions, then, Poland competes not with Czechia and Hungary but with western and southern Europe.

This picture hinges on a single market. Excluding Ukraine, Poland appears in 38.5% of answers to general questions, behind Germany (53.7%) and Portugal (45.6%), in one group with Spain (42.6%), Czechia (36.3%) and the Netherlands (36.3%). The differences within this group fall within measurement variability, so the accurate description is: “in the group of countries named in fewer than four answers in ten”.

3.3 Who Poland competes with in each market

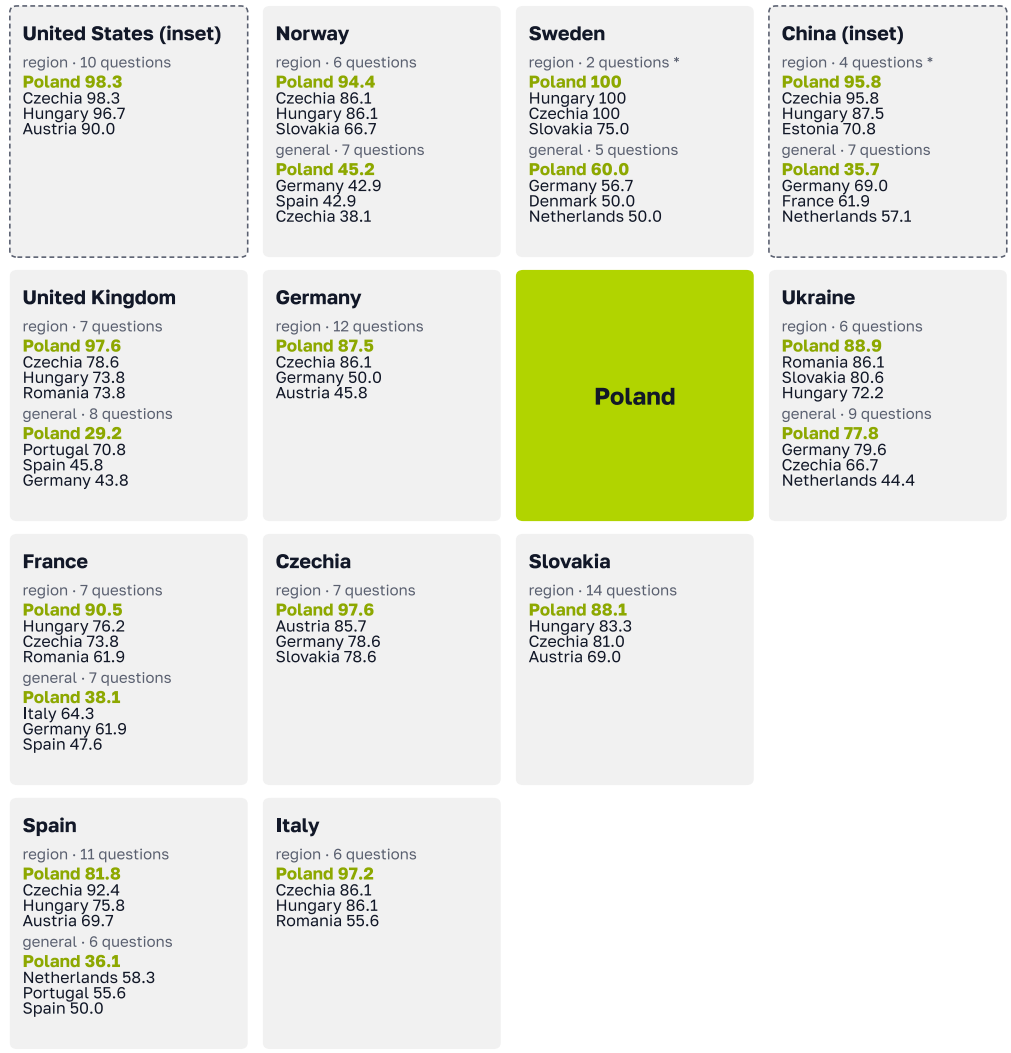
As in chapter 2, we compare markets within each question type.

Questions that point to a region. In 11 of the 12 markets Poland is the most frequently named country or shares first place with Czechia.

market	questions	Poland	three most frequently named other countries	Poland’s lead (pts)
United States	10	98.3%	Czechia 98.3 · Hungary 96.7 · Austria 90.0	0.0
Czechia	7	97.6%	Austria 85.7 · Germany 78.6 · Slovakia 78.6	+11.9
United Kingdom	7	97.6%	Czechia 78.6 · Hungary 73.8 · Romania 73.8	+19.0
Italy	6	97.2%	Czechia 86.1 · Hungary 86.1 · Romania 55.6	+11.1
Norway	6	94.4%	Czechia 86.1 · Hungary 86.1 · Slovakia 66.7	+8.3
France	7	90.5%	Hungary 76.2 · Czechia 73.8 · Romania 61.9	+14.3
Ukraine	6	88.9%	Romania 86.1 · Slovakia 80.6 · Hungary 72.2	+2.8
Slovakia	14	88.1%	Hungary 83.3 · Czechia 81.0 · Austria 69.0	+4.8
Germany	12	87.5%	Czechia 86.1 · Germany 50.0 · Austria 45.8	+1.4
Spain	11	81.8%	Czechia 92.4 · Hungary 75.8 · Austria 69.7	−10.6
China*	4	95.8%	Czechia 95.8 · Hungary 87.5 · Estonia 70.8	0.0
Sweden*	2	100%	Hungary 100 · Czechia 100 · Slovakia 75.0	0.0

* Fewer than 5 questions: tendency only.

Chart 3.2 Map of alternatives



region = questions that point to a region that includes Poland; general = questions about Europe in general (only markets with at least 5 such questions);
* fewer than 5 questions: a tendency. The tile layout is schematic; the United States and China are outside the map.

Spain is the exception: in Spanish questions about a region, Czechia appears more often than Poland (92.4% against 81.8%). It is the only market in which Poland loses to a competitor in regional questions. Poland's lead is narrowest in the United States (a tie with Czechia, both at 98.3%) and in three of the four neighbouring markets (Germany +1.4, Ukraine +2.8, Slovakia +4.8). Why the neighbours in particular? Because there the model lists all the neighbours of the country the question comes from, and Poland is simply one of them. The exception among the neighbours is Czechia (+11.9), where Poland comes ahead of Austria and Germany in questions about neighbouring countries.

General questions (markets with at least 5 such questions):

market	questions	Poland	three most frequently named other countries	Poland's lead (pts)
Ukraine	9	77.8%	Germany 79.6 · Czechia 66.7 · Netherlands 44.4	−1.9
Sweden	5	60.0%	Germany 56.7 · Denmark 50.0 · Netherlands 50.0	+3.3
Norway	7	45.2%	Germany 42.9 · Spain 42.9 · Czechia 38.1	+2.4
France	7	38.1%	Italy 64.3 · Germany 61.9 · Spain 47.6	−26.2
Spain	6	36.1%	Netherlands 58.3 · Portugal 55.6 · Spain 50.0	−22.2
China	7	35.7%	Germany 69.0 · France 61.9 · Netherlands 57.1	−33.3
United Kingdom	8	29.2%	Portugal 70.8 · Spain 45.8 · Germany 43.8	−41.7

With 5–9 questions (30–54 answers) per market, we treat the results as a tendency, but the direction is clear, and it differs from the pattern in regional questions:

- **Ukraine:** for the models, Poland and Germany are the two default destinations, named in four answers out of five. Nowhere else is Poland so close to the top.
- **United Kingdom:** in questions about a cheap trip or a move, the models name Portugal in 70.8% of answers and Poland in 29.2%. In the British market, "cheap Europe" means the south.
- **China:** in questions about studying, working or staying in Europe, Germany, France and the Netherlands dominate, and Poland appears in one answer in three.
- **Sweden and Norway** are the only markets in western and northern Europe in which Poland is among the leaders in general questions (60.0% and 45.2%), thanks to questions about cheap trips and a cheaper country "not far away".

What matters is how the models speak about Poland, not where. We also checked where Poland appears in the answer. When the question names Poland, it opens the answer almost every time. When the question does not name it, its place varies from answer to answer, even when the same question is repeated on the same day. The report therefore does not describe Poland's position in the answer and concentrates on whether Poland appears at all, in what tone and whether it is recommended. The measurement of position and its variability are described in the annex.

3.4 Home country in the answer, or what this measure does and does not show

The home country, meaning the country the question comes from, appears in 21.1% of answers. The distribution across markets is highly uneven: Czechia 64.1%, Germany 52.6%, Sweden 28.7%, Spain 25.0%, the United Kingdom 21.6%, France 20.8%, Slovakia 16.7%, Italy 8.3%, Ukraine 7.3%, the United States 2.8%, Norway 2.4%, China 0.0%.

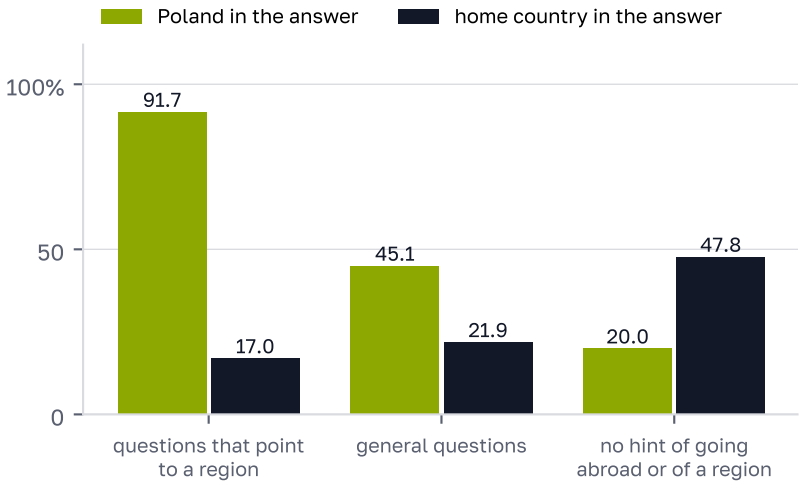
It is tempting to read this as "the models tell Czechs to stay at home". The data do not support that reading. The judge records the countries mentioned in an answer, not those recommended, and reading the answers reveals three distinct situations:

1. **The question does not mention going abroad.** Fifteen questions (four Czech, eight Swedish and one each from France, Ukraine and the United Kingdom) ask about "regions for a short break", "natural areas for a longer journey" or "the culture of this country", without mentioning either Europe or going abroad. The models read them as questions about the home country: it appears in 47.8% of these answers, Poland in 20.0%. In Czech questions of this type, Czechia appears in 83.3% of answers; in Swedish ones, Sweden appears in 41.7%. This is not advice to "stay at home" but an

answer to a question that never asked about going abroad. These questions account for 40% of all mentions of Czechia in Czech answers and for 65% of mentions of Sweden in Swedish ones.

2. **The home country as a point of reference.** In answers to questions phrased “from Ostrava”, “cheaper than in Spain” or “from Germany”, the model names the country the question comes from in order to measure distances and prices from it. For example, answering a Czech question about shopping abroad, the model opens with the sentence “The best choice from Ostrava is Poland”, and the judge correctly records Czechia in this answer as well. This is why Czechia and Germany, whose questions often concern neighbouring countries, also have high shares in regional questions (Czechia 57.1%, Germany 50.0%).
3. **The home country as one of the suggestions.** In some cases, notably in holiday questions asked from Spain (50.0% in general questions) and from France (38.1%), the model includes the country the question comes from among the destinations it lists.

Chart 3.3 When the question does not mention going abroad



Situations 2 and 3 cannot be told apart without an additional assessment, so we do not publish “stay at home” as an indicator. We do publish the conclusion drawn from situation 1, because it is robust and confirmed by an independent classification of the questions: **if the question does not mention going abroad, the models answer with the home country, and Poland drops out of the answer**. For the promotion of Poland, this means that the hardest people to reach are not those already comparing countries, but those who do not yet know they want to go abroad.

3.5 The models offer different lists

In questions that point to a region, all the models offer the same list (Czechia and Hungary at the top for each) and name Poland in 88.0–94.6% of answers. In general questions they differ more:

model	Poland	most frequently named other country	Poland's lead (pts)
DeepSeek	51.9%	Germany 68.5	−16.7
Perplexity	48.1%	Germany 51.9	−3.7
Claude	46.3%	Germany 51.9	−5.6
Gemini	46.3%	Germany 66.7	−20.4
Grok	42.6%	Germany 66.7	−24.1
ChatGPT	35.2%	Germany 42.6 · Portugal 42.6	−7.4

ChatGPT names Poland least often in general questions (35.2%, 9.9 percentage points below the average) and DeepSeek most often (51.9%, 6.8 percentage points above). Both differences exceed the threshold of 3 percentage points. Part of ChatGPT's result is explained by the way it answers: in 13.0% of its answers to general questions it names no country at all and simply describes the rules (Perplexity 18.5%; Gemini, Grok and DeepSeek 0–2%). When it does name countries, it names more of them than any other model (6.9 per answer overall). The more interesting point, however, is the pattern of competitors. For Gemini, Grok and DeepSeek, "default Europe" is above all Germany (67–69%), whereas for ChatGPT and Claude the list is flat, with no clear leader. In ChatGPT's answers, Poland does not so much lose to a single country as get lost in the crowd, or not get a list at all.

3.6 How it sounds in the answers

> **"Mainland European nations (e.g., Spain, Portugal, France, Germany, Netherlands, Italy) each set their own rules."**
>
> — Grok, United Kingdom, question about moving to a cheaper country in Europe (Log ID 736f1087-a497-4048-b5cf-0c7d3d9e6082). "Default Europe" in a single pair of brackets, without Poland.

> **"here is a breakdown of four premier European destinations that offer a rich historical narrative, a welcoming social fabric [...] ### 1. Ireland (Dublin)"**
>
> — Gemini, United Kingdom, question about a cultural trip with a rich history (Log ID e1344b80-c627-4715-9fbc-0f62fb152d33). Portugal, Austria and the United Kingdom follow; Poland does not appear.

> **"**Польща** — Спрощений вступ для переміщених осіб — Стипендії від уряду та університетів — Навчання польською або англійською — Близька культура та мова"**
>
> "Poland — simplified admission for displaced persons — government and university scholarships — tuition in Polish or English — a similar culture and language"
>
> — Claude, Ukraine, question about studying abroad (Log ID f03ae93c-f114-4236-8f34-228bf5e8bdd6). Poland on the list, alongside Czechia and Germany.

> **"Zde je přehled nejlepších regionů v ČR pro krátké cyklovýlety a praktické tipy, jak vyřešit bezpečnost kol a ubytování."**
>
> "Here is an overview of the best regions in Czechia for short cycling trips, with practical tips on bike security and accommodation."
>
> — Gemini, Czechia, question about regions for a short cycling break (Log ID 7228bc96-10b4-46e3-a551-e7682297baa5). The question does not mention going abroad; the answer stays within Czechia.

> **“Nejlepší volba z Ostravy je Polsko** (nejbližší a nejvýhodnější). Slovensko se většinou nevyplatí.”
>
> “The best choice from Ostrava is Poland (the closest and the best value). Slovakia is usually not worth it.”
>
> — Grok, Czechia, question about shopping abroad from Ostrava (Log ID 73dd40bd-a48b-453f-963c-1295059ed644). A regional question; the judge also records Czechia in this answer as a country mentioned.

HOW WE COUNT. The countries mentioned in an answer come from the judge’s “other entities mentioned” field (entity type: country), after name normalisation. A country’s share is the proportion of answers in which it appears at least once; a single answer can name many countries, so the shares do not add up to 100. Poland’s lead (in the annex: the recommendation gap) is the difference between the share of answers that include Poland and the share that include the most frequently named other country. The home country is the country of the market from which the question was asked, that is, the country of the persona adopted in the study (annex). Country lists are the most variable element of the answers (annex: repeatability test and newer-model-version test), which is why we report them only as aggregates and do not interpret differences below 3 percentage points. The measurement concerns the model’s answer as returned via the API.

CHAPTER 4

Tone and recommendation rate, or how the models speak of Poland once they do

IN BRIEF

- When a question does not name Poland and a language model (LLM) brings it into the answer of its own accord, it explicitly recommends Poland in 55.9% of cases (1,051 answers). It almost never advises against it (1.7% of answers; in a further 1.9% it takes no position), but in four answers out of ten it names Poland without explicitly recommending it.
- When the question itself points to Poland, the recommendation rate rises to 71.7%. This is not a better opinion of Poland but an effect of the question, since the model is answering the topic it has been given. That is why the headline figure is 55.9%.
- The models differ more in recommendation rate than in mention rate or tone. How large that difference is, however, depends on the answer length cap that four of the six models had in the study, so per-model values are given in the annex. The direction holds: Gemini recommends Poland most often and Claude least often, even under the least favourable variant of the truncation simulation and with the second judge, which runs on a model from a different provider.
- Question type works the opposite way from chapter 2. In questions about a region, Poland almost always appears (91.7%), but is explicitly recommended in only 47.0% of answers. In questions about Europe in general it appears less often (45.1%), but when it does, the model recommends it in 62.3% of answers. The pattern recurs in all six models, but only in questions from individuals. For companies, the models recommend Poland at similar rates in both question types (63.8% and 67.8%).
- The tone in which the models speak of Poland is moderately positive (6.95 on a 1–10 scale) and almost identical across markets. Negative answers (a score of 4 or lower) account for 4.7%; almost all respond to questions that themselves name Poland, and more than a third answer risk questions.
- A separate thematic measurement from October (section 4.7): asked directly about safety during the war in Ukraine, in none of 792 answers do the models advise against Poland as a destination, but they almost always add a “but”. Asked about safe countries anywhere in Europe for a longer stay, between half and three quarters of answers put Poland in the “with caution” group.

Poland mentioned is not the same as Poland recommended. In questions about Central Europe, the model names Poland almost every time (91.7%), but explicitly recommends it in fewer than half of those answers (47.0%). In questions about Europe in general, it names Poland rarely (45.1%), but

when it does, it recommends it in 62.3% of cases. Counted against all answers of each type, Poland is named and explicitly recommended in 43.1% of answers to regional questions and in 28.1% of answers to general ones. This chapter measures precisely that difference: not whether Poland appears, but how the model speaks of it once it does, and what that difference means for promotion.

4.1 What this chapter measures, and why the headline figure covers only questions that do not name Poland

Chapters 2 and 3 asked whether Poland appears in an answer at all, and alongside whom. This chapter asks how the model speaks of it. In every answer that mentions Poland, the automated judge assessed two things: the tone of what is said about Poland, on a 1–10 scale, and the recommendation level, that is, whether the model explicitly recommends Poland (high), names it without explicitly recommending it (moderate), advises against it (low) or takes no position (none). Both measures are calculated only for answers in which Poland appears: 7,992 for the recommendation rate and 7,987 for tone (five mentions received no tone score).

The recommendation rate must be read together with the question. When the question names Poland (“we are planning a semester in Kraków”, “I’m looking for a job in Poland”), the model speaks of it with a high recommendation level in 71.7% of answers (6,941 answers). When the question does not name Poland and the model introduces it itself, the high level falls to 55.9% (1,051 answers). Why? A question that itself points to Poland inflates the recommendation rate, because the model is answering the topic it has been given rather than choosing Poland from among other countries. Tone behaves differently: it is 7.2 in questions that do not name Poland and 6.91 in questions that do. The difference lies within measurement variability, but it shows that “more recommendation” in questions naming Poland does not mean “warmer towards Poland”. A model asked directly about Poland answers matter-of-factly and with a recommendation. A model that suggests Poland itself speaks of it more warmly, but explicitly recommends it less often.

That is why the headline figure for this chapter is the recommendation rate in questions that do not name Poland: **55.9%**. The remaining 44.1% consists overwhelmingly of answers in which Poland is one of several options, without an explicit recommendation (40.4%). Answers that advise against Poland (1.7%) or take no position (1.9%) are rare.

In the business segment, the recommendation rate in questions that do not name Poland is 65.1% (307 answers); in the individuals segment it is 52.2% (744 answers). We do not read this difference as “the models recommend Poland more to companies”. First, in questions that name Poland the order is reversed: individuals score higher (76.7%) than companies (67.5%). Second, 62% of B2B answers to questions that do not name Poland come from three markets (France, Italy and the United States), where the recommendation rate is 47.9–62.2%; the remaining markets have between 6 and 21 answers. The report therefore gives the B2B recommendation rate only for the segment as a whole; its breakdown by question type and market is described in Part II, on the basis of a separate measurement.

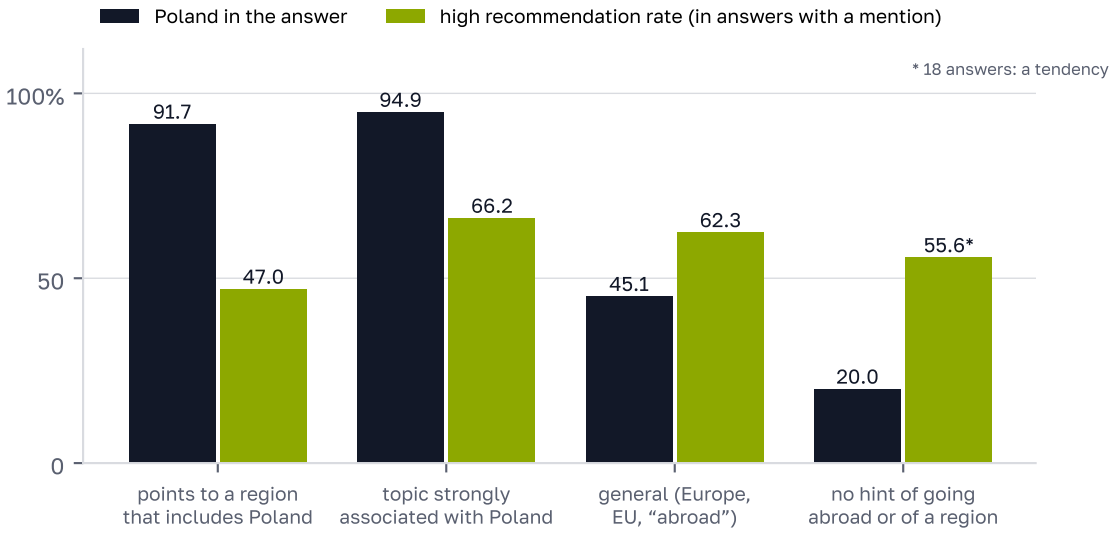
4.2 Poland mentioned versus Poland recommended

In chapter 2, question type determined whether Poland appears at all. Here it determines how it appears, and it works in the opposite direction.

question type (B2C, Poland not named in the question)	answers with a mention	Poland in the answer (chapter 2)	high recommendation rate	tone
points to a region that includes Poland	506	91.7%	47.0%	7.01
topic strongly associated with Poland	74	94.9%	66.2%	7.55
general (Europe, EU, “abroad”)	146	45.1%	62.3%	7.58
no hint of going abroad or of a region	18	20.0%	55.6%*	7.56
total	**744**	**71.3%**	**52.2%**	**7.19**

* 18 answers, so the figure indicates a tendency.

Chart 4.1 Mentioned versus recommended



Recommendation rate higher in general than in regional questions for 6 of 6 models (values per model: annex).

When the question concerns Central Europe, the model almost always names Poland, but in more than half of these answers as one country on a list, without an explicit recommendation. When the question concerns Europe in general, Poland appears in fewer than half of the answers. When it does appear, however, it is because the model has chosen it from several dozen countries; the model then explicitly recommends it in 62.3% of answers, and the tone is 0.57 points higher than in regional questions. In questions on a topic linked to Poland (memorial sites, dental treatment), the recommendation rate is 66.2% (74 answers), on a par with general questions.

In short: **in the region, Poland is mentioned; in Europe, it is recommended.** Chapters 2 and 4 describe two sides of the same coin. A regional question gives Poland presence but not distinction; a general question gives it distinction but rarely presence.

The pattern is stable: for each of the six models, the recommendation rate is higher in general questions than in regional ones (there are about 25 answers to general questions per model, so it is the direction that counts, not the size of the difference; values in the annex). The direction is not an effect of the answer length cap, since it is the same for Gemini and Grok, whose answers did not reach

the cap. Nor does the length of the list explain it: in regional questions the recommendation rate is similar whether the model names two other countries or ten (43–53%).

In companies' questions, the pattern disappears. We classified 68 business questions that do not name Poland using the same procedure (annex): 35 point to a region, 29 are general and 4 make no reference to going abroad.

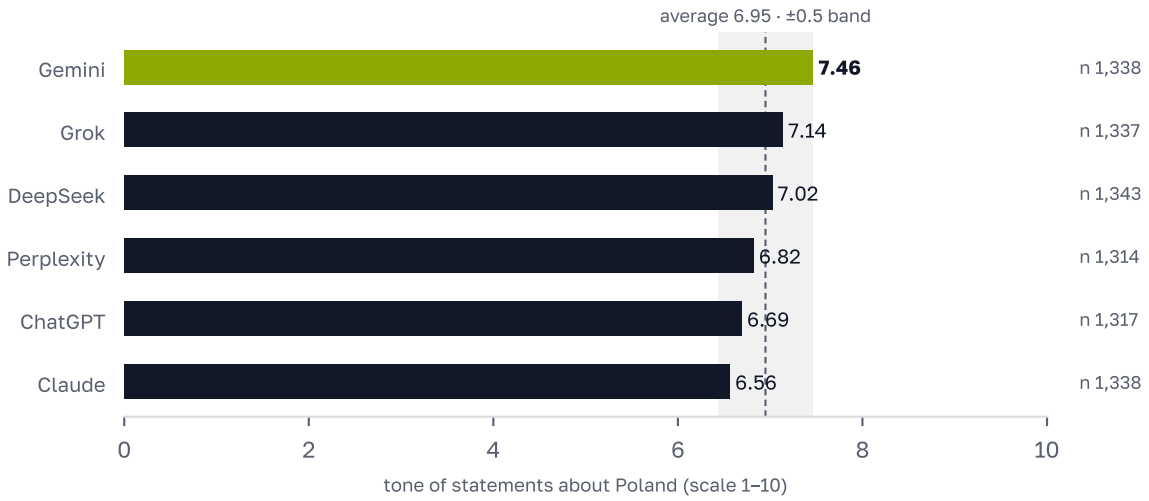
question type (B2B, Poland not named in the question)	answers with a mention	Poland in the answer	high recommendation rate	tone
points to a region that includes Poland	185	88.1%	63.8%	7.17
general (Europe, EU)	121	69.5%	67.8%	7.33
no hint of going abroad or of a region	1	4.2%	—	—
total	**307**	**75.2%**	**65.1%**	**7.22**

Visibility behaves as it does for individuals: regional 88%, general 70%, no hint of going abroad 4%. The recommendation rate does not. The difference between general and regional questions is 4 percentage points, which lies within measurement variability, and its direction varies from model to model (ChatGPT and Perplexity higher in general questions, Claude higher in regional ones, no difference for Gemini, DeepSeek and Grok). When a company is asking, a model that names Poland within Central Europe recommends it as explicitly as one that has picked it from the whole of Europe. For individuals, Poland within the region is one of several options; for companies, it is a contender in the region too. We treat the B2B figures by question type as a tendency (29–35 questions per type, mainly from France, Italy and the United States); Part II gives the full picture of business visibility.

4.3 Models differ more in recommendation rate than in tone

model	overall tone	answers with a tone score	negative answers (≤4)
Gemini	**7.46**	1,338	1.9%
Grok	7.14	1,337	3.8%
DeepSeek	7.02	1,343	4.5%
Perplexity	6.82	1,314	6.5%
ChatGPT	6.69	1,317	4.7%
Claude	**6.56**	1,338	7.1%
total	6.95	7,987	4.7%

Chart 4.2 Tone of statements about Poland, six models



Overall, the mention rate for Poland differed between the models by less than 2 percentage points (chapter 1), and generative visibility by 7 (chapter 2). The recommendation rate separates them considerably more, but in this study its per-model values cannot be stated as confidently as those for tone. Four of the six models had an upper cap on answer length in the study and often hit it. An answer cut at the cap loses its ending, and the truncation simulation on full answers shows that the explicit recommendation is then more often lost with it: at the shortest cap, the share of explicit recommendations falls by about 11 percentage points in the answers actually cut (annex). That is why we give the per-model recommendation rate in the annex, with a caveat.

The direction, however, holds. Even under the least favourable variant of the truncation simulation, Gemini, in its August 2026 version, remains the model with the highest recommendation rate and Claude the one with the lowest. Could this order be an artefact of the judge? It is not. The second judge, based on a model from a different provider, assessed a random sample of answers and placed Gemini first and Claude last, even though it set a higher bar for an “explicit recommendation” (annex).

Tone separates the models less. The only difference above the 0.5-point threshold is Gemini’s: 7.46, or 0.51 above the average (excluding the 148 answers produced by the fallback model, that is, the model which also serves as the judge in the study, leaves the difference unchanged at +0.51; and for Gemini the second judge diverges from the main judge less than for any other model). Claude has the lowest tone (6.56, −0.39), but that difference lies within measurement variability. On identical questions, that is, once the effect of question selection is removed, Gemini deviates from the average by +0.53 and Claude by −0.39 (chapter 7).

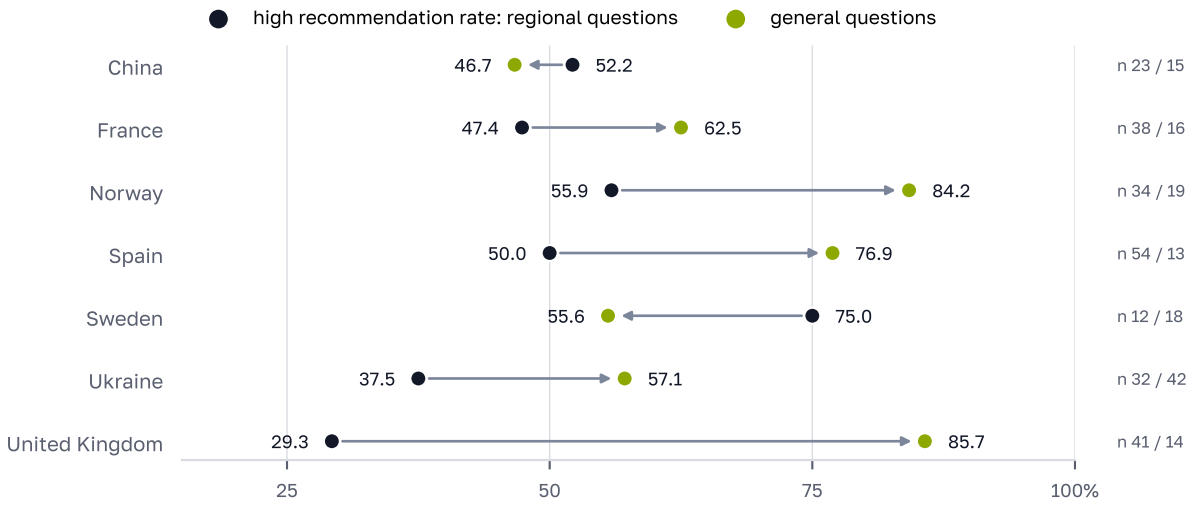
4.4 Markets, or question type first, once again

Across markets, the recommendation rate in questions that do not name Poland ranges from 37.7% (the United States) to 68.9% (China) in the individuals segment. As in chapter 2, this is not a ranking of “where the models like Poland”. Markets whose questions more often point to a region have a lower recommendation rate, because in regional questions Poland is more often merely mentioned (correlation −0.53: the larger the share of regional questions, the lower the recommendation rate). In the United States, 86% of answers to questions that do not name Poland come from regional questions; in China, 38%.

That is why we compare markets within each question type:

market (B2C)	recommendation rate, regional questions (n)	recommendation rate, general questions (n)
China	52.2% (23)	46.7% (15)
Czechia	48.8% (41)	–
France	47.4% (38)	62.5% (16)
Germany	52.4% (63)	–
Italy	57.1% (35)	–
Norway	55.9% (34)	84.2% (19)
Slovakia	50.0% (74)	–
Spain	50.0% (54)	76.9% (13)
Sweden	75.0% (12)	55.6% (18)
Ukraine	37.5% (32)	57.1% (42)
United Kingdom	29.3% (41)	85.7% (14)
United States	32.2% (59)	–

Chart 4.3 Markets within question type



General questions are shown only for markets with at least 12 answers with a mention; with 12–20 answers the result is a tendency.

The United Kingdom offers the sharpest contrast. In questions about Central Europe, the models name Poland in 97.6% of answers (chapter 3), but explicitly recommend it in only 29.3%. In questions about Europe in general, Poland appears rarely (29.2%), but when it does, the model recommends it in 85.7% of answers, usually as Kraków for a short break. Norway (55.9% against 84.2%) and Spain (50.0% against 76.9%) follow the same pattern. Ukraine is the exception in the other direction: there Poland is the default destination (77.8% of answers to general questions, chapter 2), but it is explicitly recommended in 57.1%, because answers to questions from Ukraine more often describe conditions for living and working than offer encouragement.

Tone does not distinguish between markets: none deviates from the average by 0.5 points (lowest Ukraine, 6.57; highest Norway, 7.20).

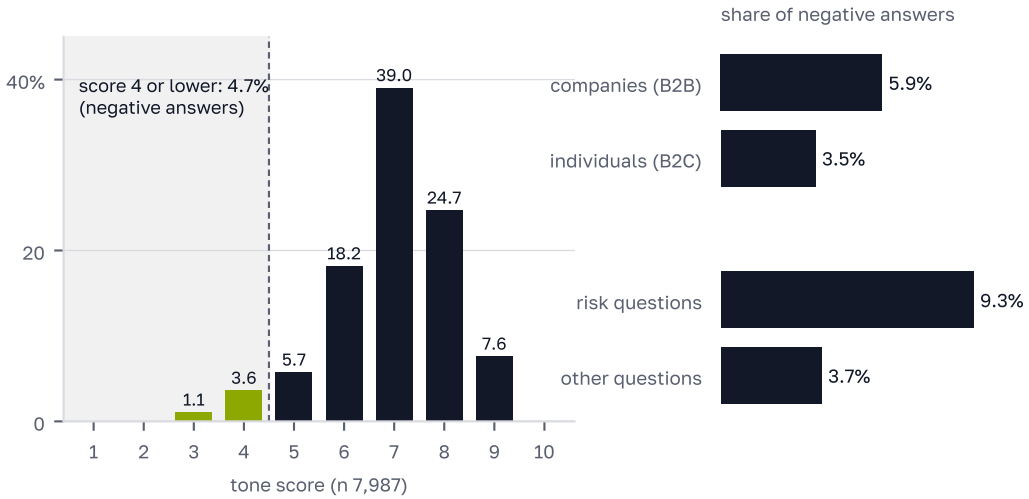
4.5 Tone and negative answers

The average tone of what is said about Poland is 6.95. The distribution is narrow: 39% of answers score 7, 25% score 8 and 18% score 6; scores of 9–10 account for 7.6%, and scores of 4 or below for 4.7% (379 answers). The models speak of Poland matter-of-factly and favourably, rarely with enthusiasm and rarely unfavourably.

Negative answers have a distinct profile:

- 90% of them respond to questions that themselves name Poland; in questions that do not name Poland, a model that introduces it itself almost never speaks ill of it;
- they concern companies (5.9% of B2B answers) more often than individuals (3.5%), and the consideration and decision stages (5.3% and 5.1%) more often than awareness (3.5%);
- more than a third (36.1%) answer questions about risk, safety or “hidden pitfalls”; in such questions the share of negative answers is 9.3%, against 3.7% in the rest.

Chart 4.4 Distribution of tone



Negative tone is therefore largely a response to a question about a problem. Anyone who asks about the risk of losing legal residence, or about the risks of investing in energy, gets a list of risks. The topics that recur in negative answers are the legalisation of residence and work (especially in questions from Ukraine and the United Kingdom), regulatory and political risk in investment, energy costs, and logistics compared with Germany. These are not false answers (chapter 7 deals separately with the 14 confirmed factual errors). They are cautious ones.

4.6 How it sounds in the answers

> **“For a cheap autumn city break from the UK, the strongest all-round options are Kraków, Budapest, and Prague.”**
>
> — Perplexity, United Kingdom, question about an inexpensive city break in Europe (Log ID 077ee596-b567-488b-9503-6b39968b713e). General question; Poland chosen by the model, high recommendation level, tone 9.

> **“Aquí está la verdadera mina de oro si buscas ahorrar sin sacrificar nivel académico. a) Polonia (Varsovia, Cracovia, Wrocław)”**
>
> “This is the real gold mine if you want to save money without compromising on academic standards. a) Poland (Warsaw, Kraków, Wrocław)”
>
> — DeepSeek, Spain, question about studying in English in the EU (Log ID c24605b8-3d35-424b-91c9-d0041c508075). General question; high recommendation level, tone 8.

> **“Relocating from the UK to Central or Eastern Europe (CEE)—countries like Poland, Hungary, the Czech Republic, Romania, or the Baltic states—is an exciting prospect.”**
>
> — Gemini, United Kingdom, question about moving to Central or Eastern Europe (Log ID bed176d8-401c-4b4b-9757-f7bfddd1564a). Regional question; Poland named but not recommended (moderate recommendation level, tone 6).

> **“Most of Central Europe (Czech Republic, Austria, Hungary, Poland, Slovakia, Slovenia) has historically been: Considered quite safe for American tourists”**
>
> — Claude, United States, question about travel safety in Central Europe (Log ID dbd715b2-02a1-499a-af8f-a4065fd008f9). Regional question; Poland as one of six countries in brackets, moderate recommendation level.

> **“Для ваших батьків головний ризик у Польщі зараз не стільки в сезонній роботі як такій, скільки в помилці з легальною підставою перебування”**
>
> “For your parents, the main risk in Poland right now lies not so much in seasonal work as such as in getting the legal basis for their stay wrong”
>
> — Perplexity, Ukraine, question about the risk of losing legal status while doing seasonal work in Poland (Log ID eb54724c-de9e-46eb-8374-3a53ec812aed). Risk question; tone 3, low recommendation level.

> **“Antes de invertir, considera estos riesgos principales [...] 1. Complejidad regulatoria y cambios en el marco de apoyo”**
>
> “Before investing, consider these main risks [...] 1. Regulatory complexity and changes to the support framework”
>
> — Grok, Spain, question about the risks of investing in wind and solar farms in Poland (Log ID 70493c95-8e13-48f5-80a6-9dd12fc78acb). Risk question; tone 3.

HOW WE COUNT. Tone is the automated judge’s score on a 1–10 scale; recommendation is classed at one of four levels (high, moderate, low, none); both measures are calculated only for answers that mention Poland. The headline recommendation rate is the share of the high level in questions that do not name Poland. The level assigned depends on the threshold the judge applies: in a random sample, the second judge (a model from a different provider) classed 38.9% of answers as “high”, compared with 57.1% for the main judge, but it kept the same extremes (Gemini highest, Claude lowest) and the same order of segments. We therefore read recommendation rates as comparisons between question types and markets, not as “this percentage of models recommend Poland”. Comparisons of the recommendation rate between models are given in the annex, because they depend on the answer length cap. We do not interpret differences in tone below 0.5 points or differences in the recommendation rate below 5 percentage points. The measurement concerns the model’s answer via the API.

4.7 Poland in the shadow of war

In October 2026, separately from the study, the six models were given eight questions from individuals about safety in Poland during the war in Ukraine: in 11 markets (excluding Ukraine), each asked twice, giving 1,056 answers. Six questions name Poland (holidays, studies, relocation, eastern Poland, Poland's role in the war, parents' worries); two ask about safe countries without naming it.

Asked directly, the models do not advise against Poland as a destination in any of the 792 answers to the questions that name it. Almost always, however, they say "yes, but": 15.7% of answers carry no caveat at all. A few answers point people who put peace of mind first to a country further west, such as Austria. The typical "but" concerns the strip along the border with Belarus, a contingency plan and the travel advice of the home country's own foreign ministry. NATO is mentioned in 78.1% of answers, the Schengen area in 5.3%.

Without Poland in the question, the scale of the map matters. For holidays in Central Europe, Poland is among the safe countries in 97.0% of answers. For studying or spending a few years anywhere in Europe, it still appears (97.7%), but between half and three quarters of answers place it in the "with caution, closer to the conflict" group, and 3.0% explicitly recommend another country, usually Spain, Portugal, Italy, France, Ireland or Germany. This is the reverse of the finding in section 4.2: in general travel questions, a model that names Poland on a European scale recommends it more often than on a regional scale; in safety questions, the European scale moves it into the group of conditional countries.

The answer depends on how the question is phrased, not on where it comes from: when the question includes a sentence saying the parents are worried, the models explicitly recommend coming in 81.1% of answers; when the question is matter-of-fact, in 55.3%. Across markets, tone differs by less than half a point.

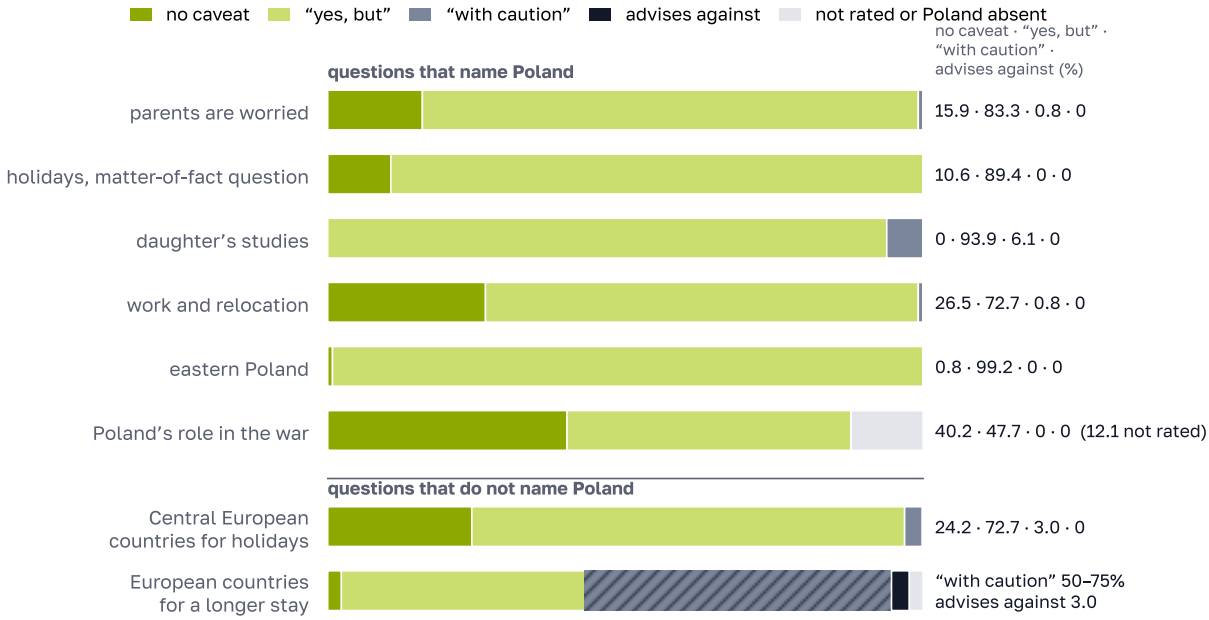
August and October. The question about parents' worries was also asked, word for word, in the study, in nine markets. The result did not change, even though four models were running newer versions in October:

	study (August–September)	thematic measurement (October)
answers	54	108
Poland in the answer	100%	100%
tone (1–10)	7.39	7.21
explicitly recommended	79.6%	78.7%

> **"Si buscas cero preocupaciones geopolíticas, céntrate en el sur y oeste de Europa: España, Portugal, Italia, Francia, Irlanda o Alemania. [...] Aun así, Polonia es razonablemente segura; el riesgo real sigue siendo extremadamente bajo gracias al paraguas de la OTAN"**
>
> "If you want zero geopolitical worries, focus on southern and western Europe: Spain, Portugal, Italy, France, Ireland or Germany. [...] Even so, Poland is reasonably safe; the real risk remains extremely low thanks to the NATO umbrella"
>
> — Gemini, Spain, question about safe European countries for studies and a longer stay (Log ID fdc9672c-db32-4713-a875-a20ff52d1660). Poland safe, but not on the first list.

HOW WE COUNT. The thematic measurement of 18–19 October 2026 is a separate dataset (different questions, a different date, four models in newer versions); we do not combine its figures with those of the study. Each answer was assigned to one class (no caveat, “yes, but”, “with caution”, advises against) by a single coder (a language model) that did not know which model had answered; an independent second run covered a stratified random sample of 576 answers (agreement 87.5%). In the question about Europe as a whole, the boundary between “yes, but” and “with caution” depends on the classification instructions, which is why we give a range. Companies’ questions are described in Part II, and the method in the annex.

Chart 4.5 Poland in the shadow of war



Each question: 132 answers from individuals (11 markets × 6 models × 2). In the last question, the boundary of “with caution” depends on the classification instructions (hatching).

CHAPTER 5

Sources, or where the answers send readers for more detail

IN BRIEF

- Across 8,400 answers, the language models (LLMs) gave 34,073 citations from 11,081 domains, in two distinct modes. Only Perplexity searched the web live and attached a list of the sources it found. The other five models wrote the URLs themselves, in the body of the answer, three of them because they had been instructed to cite sources. The most cited domain is gov.pl (738 citations), which appears for all six models and in all twelve markets.
- URLs given in the answer, which the models write themselves, lead to Polish sites in half of cases (49.5% of 14,681 citations) and to official sites in one case in three (33.5%). Sources from search make up 57% of all citations, but only 22.4% of them lead to Polish sites (6.5% to official ones). Polish institutions, then, are present in the models' answers; search rarely leads to them.
- Official Polish sites are cited heavily but narrowly: 632 domains generate 6,180 citations (nearly 10 per domain), whereas 2,712 domains of foreign institutions generate 9,427 (3.5 per domain). Foreign institutions, chiefly those of the EU, form the largest single class of sources, with 27.7% of citations.
- Answers to companies cite Polish sources more often than answers to individuals, in both modes. The report treats the differences between markets (from 25.1% to 40.7% in questions from individuals) as a citation profile, not a ranking, because in each of these figures more than half of the citations come from a single model's search.
- The portal poland.travel has 134 citations, 119 of them in URLs given in the answer by models that did not search. Perplexity's search lands on it in 15 citations, and on getyourguide.com, a commercial tour seller, in 83. For the promotion of Poland, the signal is that in search results a tour seller beats the national tourism portal.

Poland's official tourism portal, poland.travel, has 134 citations in the study. The commercial tour seller getyourguide.com has 83, but all 83 come from one model, the only one that searched the web live. The models that write their own URLs point to the national portal; search points to the tour seller. This chapter shows where answers about Poland send readers for more detail, and why where a URL comes from matters more than which model supplies it.

5.1 What this chapter measures

Chapters 1–4 dealt with whether, and how, the models talk about Poland. This chapter asks which sites the answers point to as they do so. URLs reach the answers in two ways. Perplexity, alone among the six models, searched the web live and attached a list of sources from search to its answers. The

other five models wrote the URLs themselves, in the body of the answer: ChatGPT, Grok and DeepSeek because the study instructed them to cite sources, Gemini and Claude without any such instruction, which is why they do so rarely (see the annex). The lists are analysed at domain level, and each domain is assigned one of twelve classes (see How we count, the box at the end of the chapter, and the annex). Two Polish classes matter most for the report: official sites (administration, agencies, public institutions, museums, universities) and other Polish sites. So does their sum, which we call the share of Polish sources.

Across all citations, more than 11,600 of the 34,073, or 34.0%, lead to Polish domains. But this figure combines two different mechanisms and mostly describes search: Perplexity accounts for 19,392 citations, the other five models together for 14,681. The report therefore gives the share separately for each mode. Among URLs given in the answer, without search, Polish sources make up 49.5% (official ones 33.5%); among sources from search, 22.4% (official ones 6.5%). The first figure shows where the model itself points; the second, what the web search finds.

Half of the answers (50.6%) contain at least one citation; there are 4.1 citations per answer on average, and 30.9% of all answers cite at least one official Polish site.

5.2 Institutions, Polish and foreign, account for almost half of all citations

Counted across both citation modes, the twelve classes of sources fall into three groups. Institutions account for 45.8% of citations: foreign and EU (27.7%) and official Polish (18.1%). Commercial and sector sites account for 39.7%: other Polish (15.9%), foreign companies (13.8%), consultancies, law firms and HR (4.4%), commercial travel services (3.5%) and data aggregators (2.1%). The remainder is foreign media (7.4%), travel blogs (4.0%), community content and forums (2.2%), Wikipedia (0.6%) and 124 citations (0.4%) that could not be classified.

Two features of this distribution stand out. First, Wikipedia barely registers as a source: 206 citations from 37 domains. The models do not send readers to the encyclopaedia; they send them to government offices. Second, citations of official Polish sites are concentrated. Of the 11,081 domains, 632 are official Polish ones, yet they generate 6,180 citations, nearly ten per domain. There are 2,712 foreign-institution domains, generating 9,427 citations, or 3.5 per domain. The answers thus return systematically to a few dozen Polish public sites: gov.pl (738), paih.gov.pl (449), podatki.gov.pl (260), ekrs.ms.gov.pl (255), stat.gov.pl (202), study.gov.pl (161), biznes.gov.pl (139), poland.travel (134), nawa.gov.pl (117), zus.pl (116). All of them appear in all twelve markets (poland.travel in eleven) and for all six models (poland.travel and zus.pl for five).

Citations are concentrated in another sense too: the 2,909 most-cited domains collect 24,814 citations (73%), while the remaining 8,172 domains collect 9,259. One domain in four thus receives almost three quarters of the citations; the rest is a long tail of rarely cited URLs, reached mainly by search.

Among URLs given in the answer, without search, the classes rank differently from the all-citation total: foreign institutions 38.6%, official Polish 33.5%, other Polish 16.0%, foreign companies 4.4%. On its own, then, the model points to institutions in more than two thirds of cases and to Polish sources in half. Search adds commercial, media and community sources: foreign companies account for 20.9% of its citations, foreign media for 12.3%, and official Polish sites for 6.5%.

5.3 Search points to Polish sites less often than the model on its own

model	where the URLs come from	answers with citations	citations	average per answer	Polish sources	of which official	of which other	answers citing an official Polish site
Perplexity	live search	99.4%	19,392	13.9	22.4%	6.5%	15.9%	43.6%
ChatGPT	model, instructed to cite sources	70.3%	7,001	5.0	49.1%	33.5%	15.6%	49.8%
Grok	model, instructed to cite sources	59.8%	2,982	2.1	48.1%	36.3%	11.8%	41.9%
DeepSeek	model, instructed to cite sources	49.4%	3,435	2.5	46.1%	27.3%	18.8%	31.4%
Gemini	model, no instruction	18.1%	1,004	0.7	66.1%	44.9%	21.2%	14.6%
Claude	model, no instruction	6.6%	259	0.2	60.2%	39.0%	21.2%	4.1%
five models without search		**40.8%**	**14,681**	**2.1**	**49.5%**	**33.5%**	**16.0%**	**28.3%**
total		**50.6%**	**34,073**	**4.1**	**34.0%**	**18.1%**	**15.9%**	**30.9%**

Chart 5.1 Share of Polish sources, six models

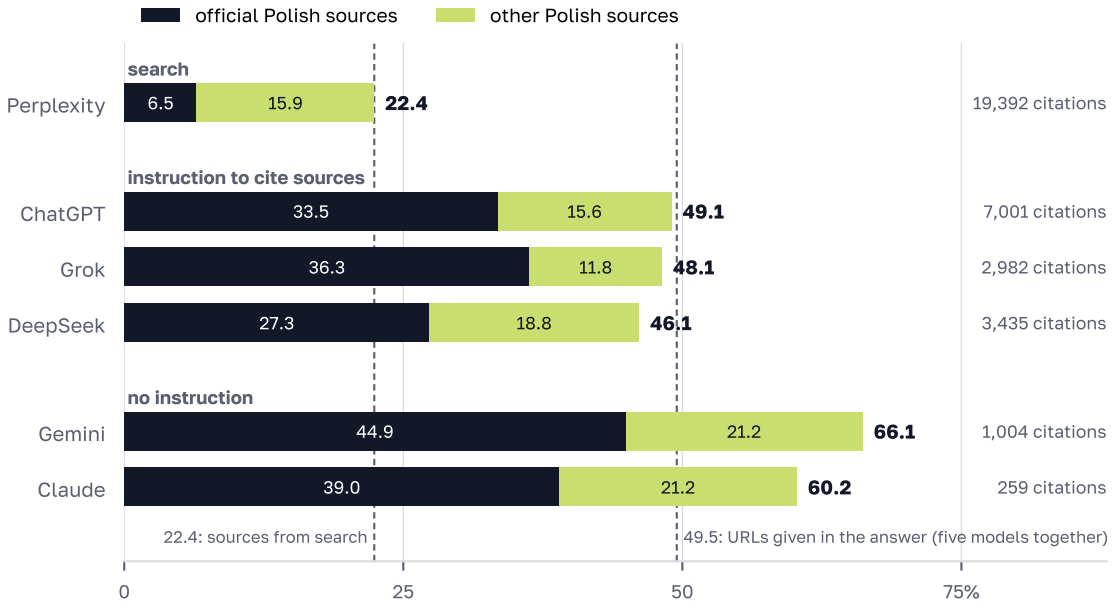
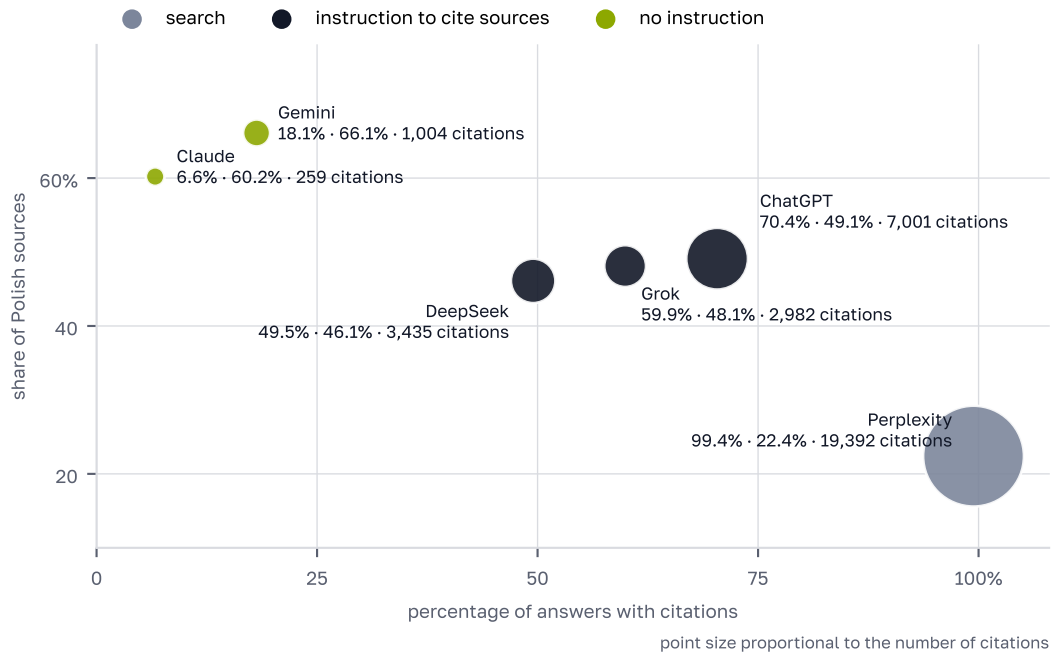


Chart 5.2 Two modes of citing sources



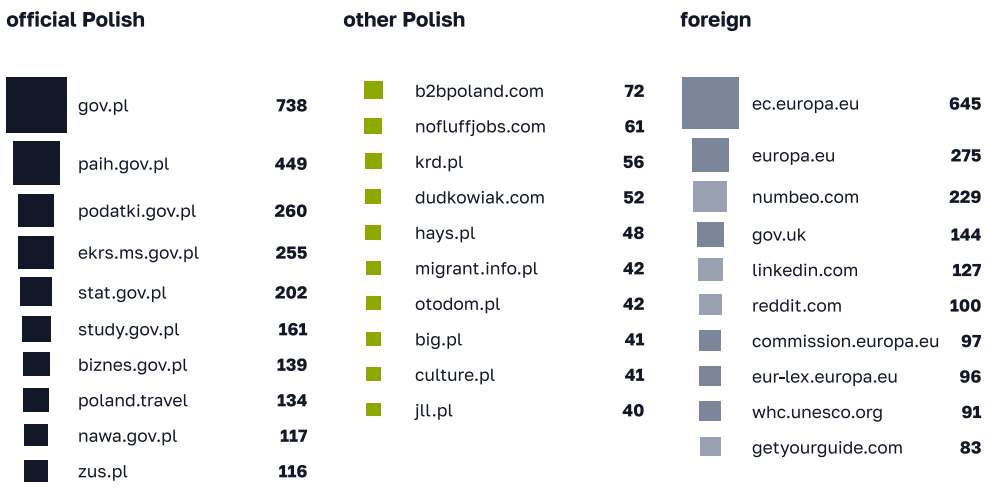
The models differ here more than on any other measure in this report, but the difference is explained chiefly by the conditions of measurement, not by the models themselves. Perplexity, the only model that searched the web live, attaches a list of sources from search to 99.4% of its answers: 14 citations per answer, of which 20.9% go to foreign companies, 12.3% to foreign media, 6.5% to travel blogs and 3.2% to forums. These classes barely occur for the other models. ChatGPT, Grok and DeepSeek were instructed to cite sources and write the URLs themselves: 35–41% of their citations lead to foreign and EU institutions, and a further 46–49% to Polish sites. Gemini and Claude received no such instruction and give URLs only exceptionally (in 18.1% and 6.6% of answers respectively). When they do, the URLs point to registers and government offices (gov.pl, ekrs.ms.gov.pl, podatki.gov.pl and paih.gov.pl), with ec.europa.eu only after them. We therefore do not read the differences in citation counts between models as a property of the models.

This carries a practical caution for anyone reading rankings. The statement “search points to Poland least often” is true when counted by citations, but looks different at the level of answers. An official Polish site appears in 43.6% of Perplexity’s answers, against 49.8% for ChatGPT and 41.9% for Grok. The differences are within a few percentage points. Search does not leave Polish sources out; it adds a dozen or so others alongside them. For promotion, though, something else matters more. URLs given in the answer lead to Polish government offices in one case in three (33.5%); search results do so in 6.5%. Polish institutions are present in the models’ answers, but on the web that the model searches they lose out to foreign companies, media and services. Chapter 9 returns to this distinction.

5.4 The domain map

The fifteen most-cited domains in the study are: gov.pl (738), ec.europa.eu (645), paih.gov.pl (449), europa.eu (275), podatki.gov.pl (260), ekrs.ms.gov.pl (255), numbeo.com (229), stat.gov.pl (202), study.gov.pl (161), gov.uk (144), biznes.gov.pl (139), poland.travel (134), linkedin.com (127), nawa.gov.pl (117), zus.pl (116). Ten of the fifteen are Polish and official, two belong to the EU, one is British and two are commercial (a cost-of-living comparison site and a professional network).

Chart 5.3 The domain map



Tile area proportional to the number of citations (after deduplication); the 10 most frequently cited domains in each column.

official Polish · other Polish · foreign and EU institutions · other foreign (aggregators, forums, commercial sites)

On the Polish side the map has two layers. The official layer consists of administration and registers: the gov.pl portal and its services (podatki, biznes, study), the National Court Register (KRS), Statistics Poland (GUS), the Social Insurance Institution (ZUS), the Polish Investment and Trade Agency (PAIH) and the Polish National Agency for Academic Exchange (NAWA). For tourism, it also includes poland.travel, museums (polin.pl, auschwitz.org, muzeumkrakowa.pl) and carriers (intercity.pl). The other layer is fragmented: 1,956 domains averaging 2.8 citations each. It is led by job and property sites (nofluffjobs.com 61, hays.pl 48, otodom.pl 42), credit reference agencies (krd.pl 56, big.pl 41), portals for foreigners (migrant.info.pl 42) and the sites of individual law firms and consultancies, usually cited by a single model (b2bpoland.com 72, dudkowiak.com 52). It is this layer that decides whether a reader who has finished the answer lands on a Polish commercial site or a foreign one.

On the foreign side, European Union institutions dominate: ec.europa.eu, europa.eu, commission.europa.eu, eur-lex.europa.eu and eures.europa.eu together account for more than 1,100 citations; then come UNESCO (whc.unesco.org 91) and the government sites of the countries the questions come from (gov.uk 144, travel.state.gov, auswaertiges-amt.de, regjeringen.no, dpsu.gov.ua). Among commercial sources the most frequent are comparison sites and aggregators (numbeo.com 229, with 50–68 citations each from Grok, DeepSeek, ChatGPT and Perplexity), financial and relocation services (wise.com 60), real-estate and tax advisers (cushmanwakefield.com 52, taxsummaries.pwc.com 48, cms.law 46, content.knightfrank.com 60) and, only in Perplexity's search results, booking sites (getyourguide.com 83, booking.com 40, viator.com 21).

For the promotion of Poland one contrast stands out. The official tourism portal poland.travel has 134 citations: Grok 60, ChatGPT 38, DeepSeek 19, Perplexity 15, Gemini 2. The commercial getyourguide.com has 83, all from Perplexity. Perplexity's search, the only search in the study, points to the tour seller five times as often as to the national tourism organisation's site, while the models that supply their own URLs do the opposite. This is not a judgement on the quality of either site, merely a measure of which of them make it into the answers.

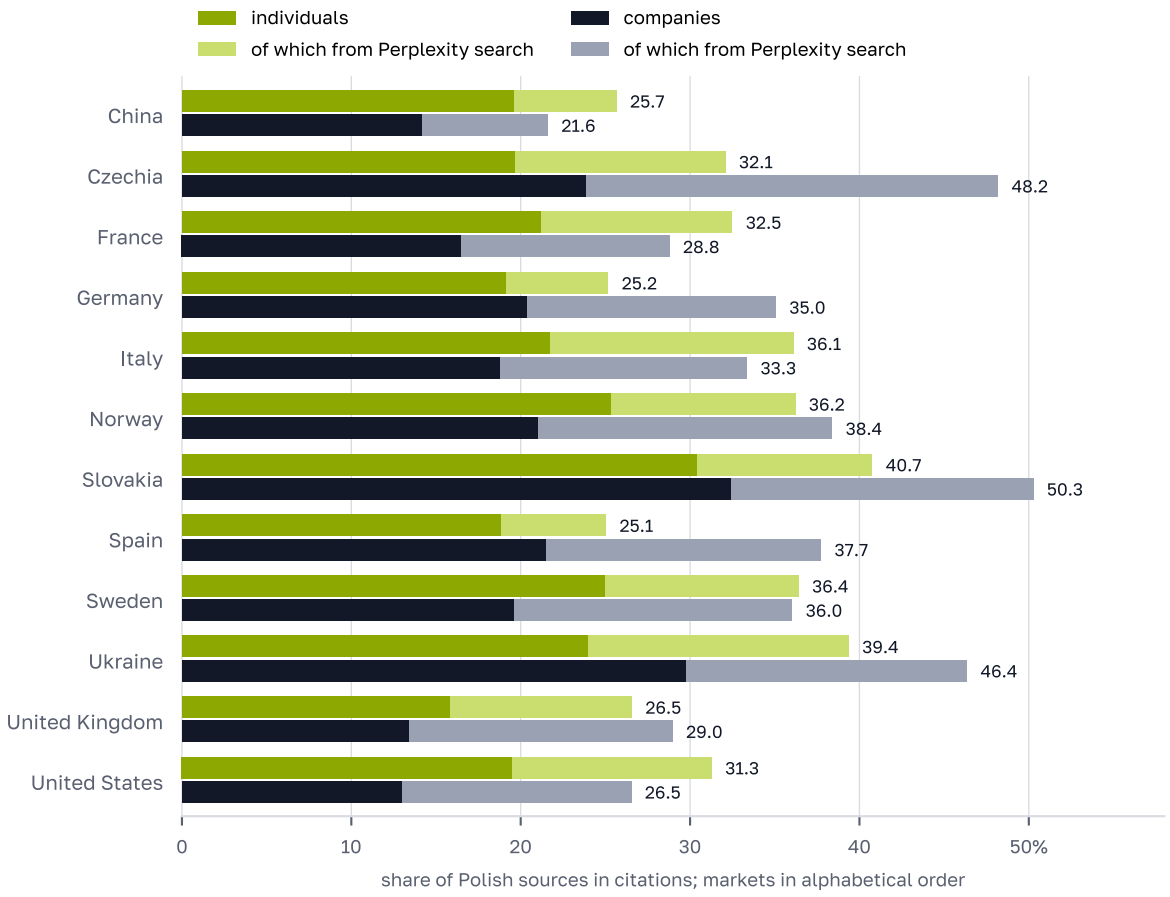
5.5 Segments and markets, or a citation profile, not a ranking

In questions from companies, 36.3% of citations lead to Polish sources (15,282 citations); in questions from individuals, 32.3% (18,791). The gap of 4 percentage points exceeds the threshold of 3 percentage points, but we read it with caution, because the two segments differ in their mix of markets and topics: questions from companies draw registers, law firms and consultancy sites (18.7% of citations go to other Polish sites, 7.1% to consultancies), while questions from individuals draw travel blogs and travel services (13.3% combined). In both segments URLs given in the answer are Polish more often than sources from search: 54.5% against 25.0% for companies, and 46.3% against 19.9% for individuals.

Markets differ more, but in these differences more than half of the citations come from a single model's search. The table therefore shows Perplexity's share alongside each value, and it is not a ranking.

market	individuals: Polish sources	citations	of which Perplexity	companies: Polish sources	citations	of which Perplexity
China	25.7%	1,629	55%	21.6%	1,435	69%
Czechia	32.1%	1,555	59%	48.2%	1,185	60%
France	32.5%	1,588	54%	28.8%	1,129	59%
Germany	25.2%	1,800	47%	35.0%	1,329	63%
Italy	36.1%	1,518	59%	33.3%	1,184	61%
Norway	36.2%	1,596	52%	38.4%	1,253	61%
Slovakia	40.7%	1,715	46%	50.3%	1,681	59%
Spain	25.1%	1,390	53%	37.7%	1,100	60%
Sweden	36.4%	1,703	48%	36.0%	977	65%
Ukraine	39.4%	1,481	54%	46.4%	1,411	56%
United Kingdom	26.5%	1,374	54%	29.0%	1,260	64%
United States	31.3%	1,442	53%	26.5%	1,338	66%

Chart 5.4 Markets: a citation profile



Three observations hold whatever Perplexity's share. First, Polish sources are cited most often in neighbouring and eastern markets (Czechia, Slovakia and Ukraine in questions from companies), and least often in English-speaking markets and in China. Second, in questions from Czech and Slovak companies every second citation is Polish, and the most frequent domains are registers: ekrs.ms.gov.pl (48 citations in Slovakia), paih.gov.pl and biznes.gov.pl (47 and 32 in Czechia). Third, in questions from individuals in the United Kingdom the most-cited domain is not a Polish site but gov.uk (cited in 74 answers, against 46 for gov.pl), and in questions from Italy and Spain it is the comparison site numbeo.com (in 32 and 21 answers respectively). The answers point to the sources of the country the question comes from and to comparison sites, while official Polish sites come second or third in these answers. We call this citability: a source makes it into the answer not because it is about Poland, but because it addresses the situation of those asking, in a language they read. Chapter 9 returns to this rule, and the market portraits in Chapter 8 return to these lists.

Germany is an exception worth a sentence: in questions from individuals its share of Polish sources is among the lowest (25.2%), but questions from companies (35.0%) produce ahk.pl (23 citations), the only bilateral chamber-of-commerce domain in the study to reach a market's top four.

5.6 What it looks like in the answers

In this chapter, instead of quotations, we show what lies behind the figures: the complete lists of domains that the models attached to specific answers (the citations column for the answer with the stated Log ID).

- ChatGPT, a company from Czechia, question “own branch or distributor”: ekrs.ms.gov.pl, biznes.gov.pl, isap.sejm.gov.pl, gov.pl (Log ID ed01d1ac-65d6-42a1-8fef-961cdef9132c). Four sources, all Polish and official: a typical ChatGPT list for questions from companies.
- Perplexity, a family from the United Kingdom, question about moving to Poland for a few years: numbeo.com, gov.pl, lifeindexed.com, livingcostindex.com, gov.uk, wise.com (Log ID 2cad9bcf-3cf2-4153-8d6b-682738d48095). Cost-of-living comparison sites, a British site and a financial service; just one of the six is an official Polish site.
- Grok, a person from Norway, question about a short city break (Poland or else Lithuania or Latvia): poland.travel, lithuania.travel, latvia.travel, efa-aef.eu, commisceo-global.com, culture.pl (Log ID 5036b34b-16d0-4297-9c52-ece29c0171f5). The official tourism portals of three countries side by side; for Grok, poland.travel is the rule, for Perplexity the exception.
- Gemini, a company from Germany, question about the reliability of Polish trading partners: ekrs.ms.gov.pl, prod.ceidg.gov.pl, podatki.gov.pl (Log ID e4dc54d0-0fc3-4e14-b6e0-debc4213a141). Three registers and nothing else: this is what the rare citations from Gemini and Claude look like.
- Perplexity, a family from China, question about schools for the children after the move: 11 domains, including gov.pl and study.gov.pl alongside eol.cn and pl.china-embassy.gov.cn (Log ID e60533da-010b-43e0-aa3b-87dc80a9c8a6). Official Polish sites alongside Chinese sources that are cited only in this market.

HOW WE COUNT. A citation is a page URL in an answer: for Perplexity, an item in the list of sources returned by search; for the other models, a URL that the model wrote in the body of the answer. Citations are counted at domain level, once per answer (a domain mentioned several times in one answer counts once), after invalid URLs have been discarded (99 in all: DeepSeek 41, Grok 32, ChatGPT 22, Claude 4); we did not check whether the full URLs given by the models actually exist. The share of Polish sources is the percentage of citations that lead to official Polish sites and other Polish sites combined. The report gives it separately for URLs given in the answer (49.5%) and for sources from search (22.4%); combined, across all citations, it is 34.0%. The number and type of citations for each model depend on the mode and on the measurement instructions (see the annex), so we do not read them as a property of the model. Shares by model and by market are always given with the number of citations; we do not interpret differences in shares below 3 percentage points. The classification of domains carries an uncertainty band: 3,553 citations (10.4%) lead to domains whose class we assigned with medium (2,971) or low (582) confidence; excluding them changes the share of Polish sources by no more than 1.0 percentage point for any model, and overall from 34.0% to 35.8%, and no ordering in the chapter changes. Almost the entire band (3,323 citations) falls on Perplexity, because search reaches furthest into the tail of rare domains. The annex describes how the domains were classified and how the classification was checked. The measurement concerns the model’s answer as returned via the API.

CHAPTER 6

Two Polands, or Poland on request and Poland in use

IN BRIEF

- Asked about Poland directly, language models (LLMs) answer with the canon. Chopin, Wałęsa, Nicolaus Copernicus and John Paul II appear in every one of the 144 answers to the question about the best-known Poles, Maria Skłodowska-Curie in 140 and Robert Lewandowski in 138. This is “Poland on request”, a separate measurement taken in October 2026.
- In answers to practical questions about travelling to Poland and about its history and culture, the same canon barely features: Chopin appears in 12.0% of answers, Copernicus in 4.5%, John Paul II in 2.6% and Maria Skłodowska-Curie in 0.8%. We call this difference the activation gap.
- The canon enters practical answers only through the doors that the question opens: a place to see (Chopin appears alongside Warsaw or Żelazowa Wola in 97% of his mentions, Copernicus alongside Toruń or Frombork in 54%) or a chapter of history (John Paul II and Wałęsa in questions about the turning point of 1989). Twentieth-century and contemporary creative figures have no such doors: in answers to 130 questions about Polish culture they appear in 1.8% of answers. The same models name them, in the same study, when the question asks directly about contemporary Poland. Polish games appear in 86% of answers, Kieślowski and Tokarczuk in one in four.
- The economy “on request” means the automotive industry, games and food (close to 100% of answers about the strongest industries), followed by machinery and household appliances (49%). Aviation (24%), electronics (17%) and pharmaceuticals (7%), which Poland presents to investors as strategic sectors, are on the margins.
- The Polish Tourism Organisation (POT) campaign “Poland. More than you expected” (running since August 2026) features Lewandowski, Wałęsa and Chopin, who are part of the models’ canon, alongside faces absent from the models’ answers: Jesse Eisenberg, Anja Rubik, Nikodem Rozbicki and JIMEK appear in none of the 1,440 portrait answers, and Maryla Rodowicz in four. The measurement was taken eight weeks after the campaign’s launch and does not assess its effectiveness.

Fryderyk Chopin appears in every one of the 144 answers to the question about the best-known Poles. In answers to questions about travelling to Poland, its history and its culture: 12.0%. Maria Skłodowska-Curie: 140 out of 144, and 0.8%. Asked directly, the models list the Polish canon almost without exception; when someone is planning a weekend in Warsaw, they hardly ever draw on it. This chapter measures the distance between the Poland that is asked about directly and the Poland that

someone needs in order to make a decision, and shows through which doors the canon enters the answers, when it does.

6.1 Two measurements that ask about different things

Chapters 1–5 described the models' answers to the questions people ask before a trip, a course of study, a move or an investment. This chapter sets them against something else: what the models say when the question is directly about Poland.

Poland on request is the open portrait: a separate, dated measurement taken on 14–15 October 2026. It consists of ten questions asked directly about Poland, covering the best-known Poles, brands, science, culture, the main branches of the economy, places, first associations and other topics. Each question was put in the language of each of the 12 markets, to each of the six models, twice. That gives 144 answers per question and 1,440 in total. The measure is the number of answers in which a given name appears, with a maximum of 144.

Poland in use is the main study, conducted in August and September 2026. For the comparison we selected 186 questions from individuals in which the names of people and works might naturally come up in the answer: 130 concern Polish culture and 126 history, memorial sites or museums (some questions concern both); three quarters name Poland directly ("we are planning four days in Warsaw or Kraków", "I want to understand the history of Poland"). That amounts to 1,116 answers from the six models. We excluded from the subset the 12 "what is Poland famous for today" questions, because they ask for a portrait rather than set a practical task; they return in 6.4 as a control.

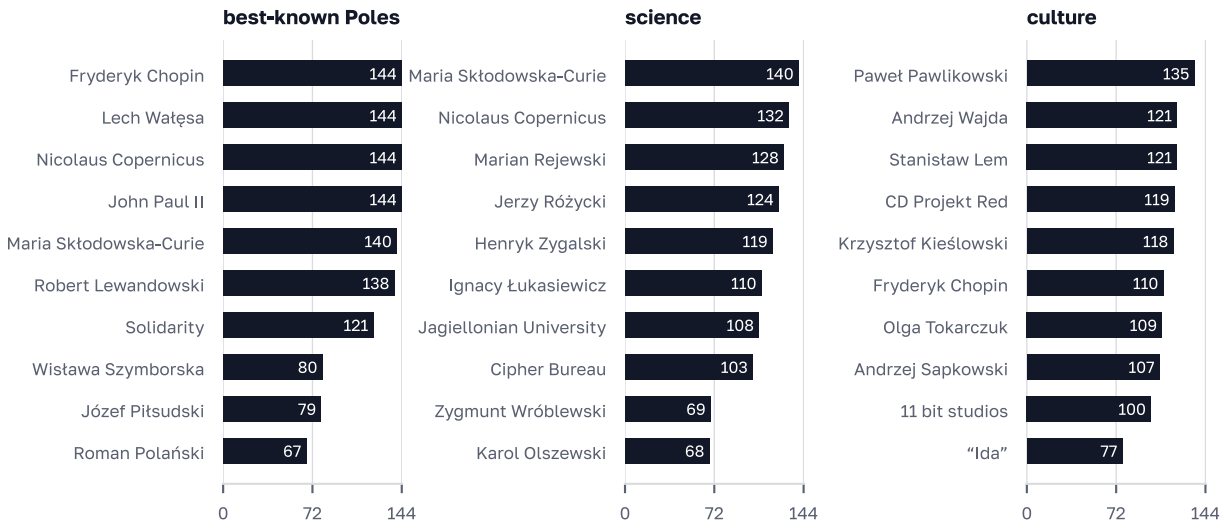
This subset contains not a single question about sport, film, literature or games. That matters for interpretation: if an athlete's name does not appear in an answer about a weekend in Kraków, it does not mean that the model leaves it out where it belongs.

The two measurements differ in date and, for four models, in version too. The version test concerned mentions of Poland (see the "How we count" box); the effect of version on the presence of names was not measured separately, whereas the gaps described below range from several dozen to almost a hundred percentage points. Even so, the two measurements never form a single indicator: in the tables they sit in separate columns, with separate denominators.

6.2 Poland on request, or the canon

best-known Poles	answers (of 144)	science	answers	culture	answers
Fryderyk Chopin	144	Maria Skłodowska-Curie	140	Paweł Pawlikowski	135
Lech Wałęsa	144	Nicolaus Copernicus	132	Andrzej Wajda	121
Nicolaus Copernicus	144	Marian Rejewski	128	Stanisław Lem	121
John Paul II	144	Jerzy Różycki	124	CD Projekt Red	119
Maria Skłodowska-Curie	140	Henryk Zygalski	119	Krzysztof Kieślowski	118
Robert Lewandowski	138	Ignacy Łukasiewicz	110	Fryderyk Chopin	110
Solidarity	121	Jagiellonian University	108	Olga Tokarczuk	109
Wisława Szymborska	80	Cipher Bureau	103	Andrzej Sapkowski	107
Józef Piłsudski	79	Zygmunt Wróblewski	69	11 bit studios	100
Roman Polański	67	Karol Olszewski	68	"Ida"	77

Chart 6.1 Poland on request



number of answers in which the name appears (of 144: 12 markets × 6 models × 2 repetitions); common scale

The top of the canon is stable: the leading names on each list appear in all 12 markets and with all six models. Further down, the order shifts: two answers to the same question share, on average, fewer than half of their names (list similarity of 0.44; 0.57 for the best-known Poles), which is why the report publishes the top of each list with the number of answers rather than individual lists.

Among brands, the canon is food and cosmetics: Wedel 140, Żubrówka 122, Inglot 113, Reserved 113, Ptasie Mleczko 104, Ziaja 91. First associations with Poland are dominated by places and symbols: Warsaw 117, Kraków 116, Solidarity 97, pierogi 93, Auschwitz-Birkenau 70; people come further down (John Paul II 61, Wałęsa 57, Chopin 54). Even in the portrait, the models think of Poland through its cities, and that matters for what happens in practical answers.

6.3 Poland in use, or a map of cities and public offices

Across the whole study (8,400 answers for individuals and companies), 103 Polish names appear in at least 11 of the 12 markets and with at least five of the six models. This is the core of “Poland in use”: 19 cities, 29 institutions, 10 regions, 9 historic sites and memorials, 18 companies (covered in Part II), a handful of programmes and products, and just two people, one of them from the canon: Fryderyk Chopin, present in 1.0% of answers.

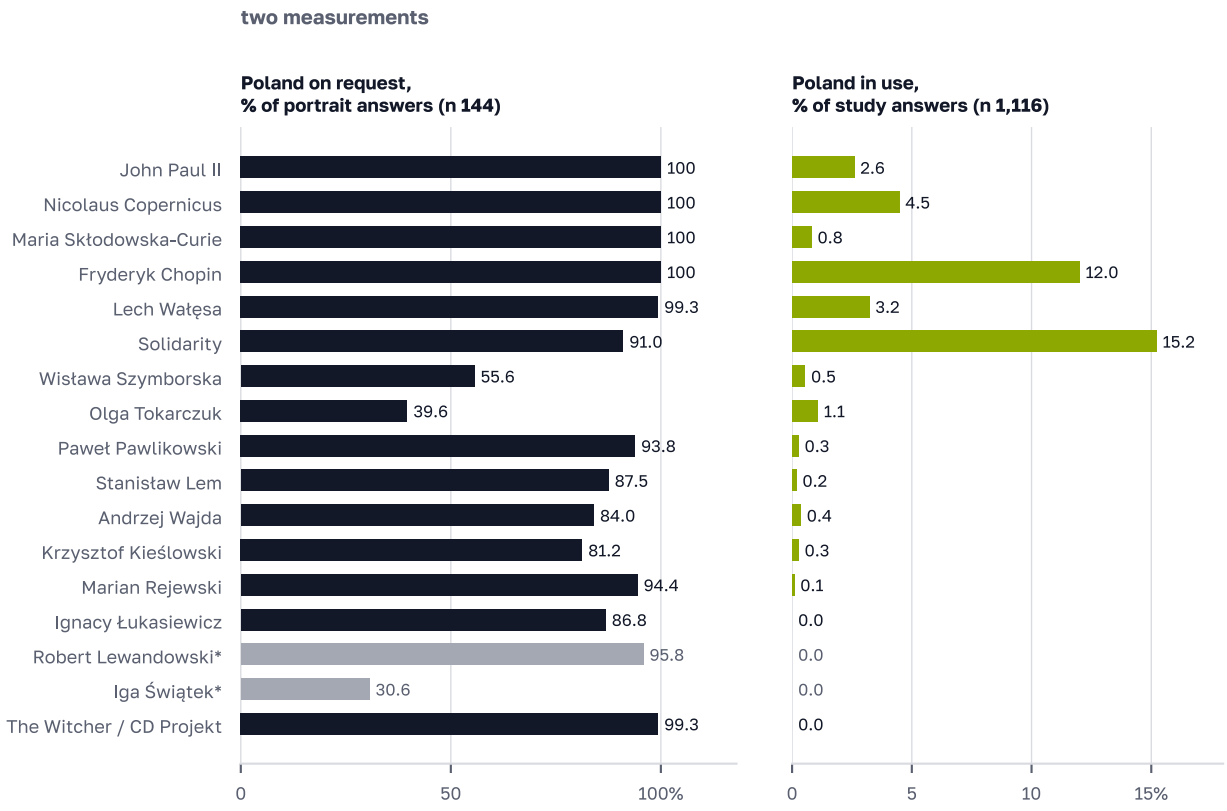
The most frequent core names are Warsaw (40.0% of answers), Kraków (36.6%), Wrocław (27.0%), Gdańsk (21.8%), Poznań (15.1%), Katowice (9.6%) and Łódź (8.7%); among institutions and historic sites, the National Court Register (7.1%), Wawel (6.6%), PAIH, the Polish Investment and Trade Agency (6.5%), Auschwitz-Birkenau (6.1%) and ZUS, the Polish social-insurance institution (5.8%). Katowice belongs above all to answers for companies: 635 of 806 (78.8%, although companies account for half of the data set), most often alongside Gliwice, Tychy, the A4 motorway and the Polish Investment Zone, that is, the industrial base of the south. In practical answers, Poland is a country of cities, historic sites and public offices. People hardly feature in it.

6.4 The activation gap, or names known but not recalled

name	measurement 1 (Poland on request): % of portrait answers (question, n 144)	measurement 2 (Poland in use): % of study answers (n 1,116)
John Paul II	100 (best-known Poles)	2.6
Nicolaus Copernicus	100 (best-known Poles)	4.5
Maria Skłodowska-Curie	100 (best-known Poles)	0.8
Fryderyk Chopin	100 (best-known Poles)	12.0
Lech Wałęsa	99.3 (best-known Poles)	3.2
Solidarity	91.0 (best-known Poles)	15.2
Wisława Szymborska	55.6 (best-known Poles)	0.5
Olga Tokarczuk	39.6 (best-known Poles)	1.1
Paweł Pawlikowski	93.8 (culture)	0.3
Stanisław Lem	87.5 (culture)	0.2
Andrzej Wajda	84.0 (culture)	0.4
Krzysztof Kieślowski	81.2 (culture)	0.3
The Witcher / CD Projekt	99.3 / 100 (culture)	0 / 0
Marian Rejewski	94.4 (science)	0.1
Ignacy Łukasiewicz	86.8 (science)	0
Robert Lewandowski*	95.8 (best-known Poles)	0
Iga Świątek*	30.6 (best-known Poles)	0

* The “in use” subset contains no questions about sport; we do not read the 0 values as a gap. Both columns were counted by full-text search of the answers; the lists in 6.2 come from name normalisation, hence differences of a few answers (e.g. Skłodowska-Curie, Wałęsa; see the “How we count” box).

Chart 6.2 The activation gap



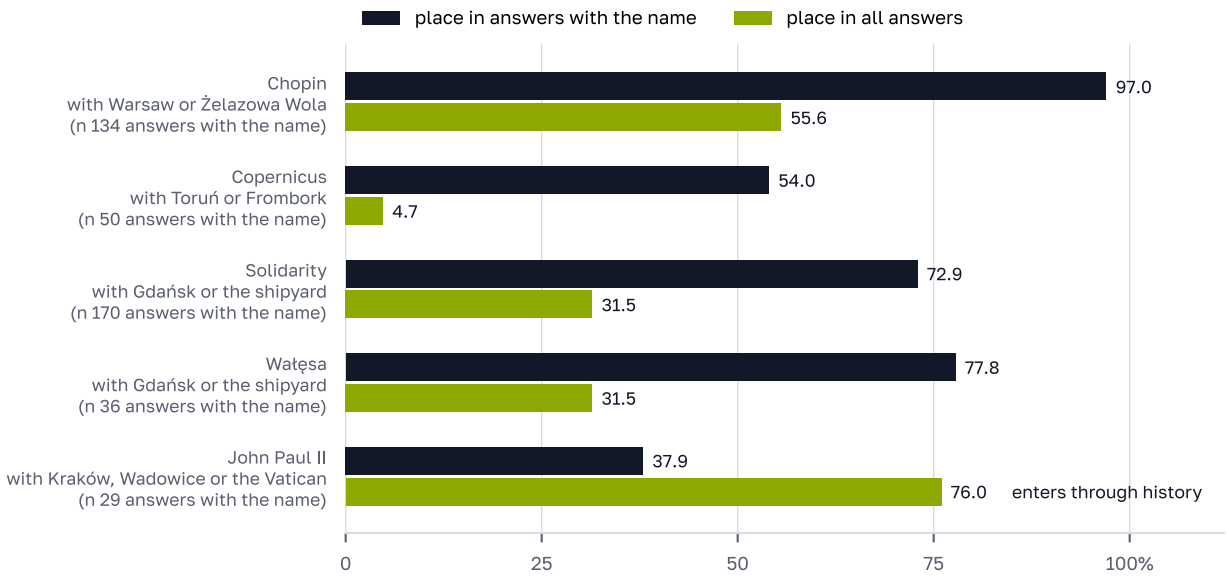
* athletes (in grey). Different axes and measurements: the panels are not connected.

Brands have only 66 answers in the subset, so the report presents them as a tendency: Wedel 97.2% on request against 9.1% in use, Ziaja 64.6% against 9.1%, Inglot 79.9% against 0.

The gap does not mean that the models “do not know” Copernicus or Lem. It means that a name the model gives when asked directly does not appear in an answer to a question that does not ask for it. So when does the canon make it into a practical answer? We counted what it appears alongside.

name	“in use” answers with the name	of which together with the place	the same place across all 1,116 answers
Chopin with Warsaw or Żelazowa Wola	134	97.0%	55.6%
Copernicus with Toruń or Frombork	50	54.0%	4.7%
Solidarity with Gdańsk or the shipyard	170	72.9%	31.5%
Wałęsa with Gdańsk or the shipyard	36	77.8%	31.5%
John Paul II with Kraków, Wadowice or the Vatican	29	37.9%	76.0%

Chart 6.3 Doors



Chopin makes it into a practical answer alongside Warsaw: through the concerts in Łazienki, his heart in the Holy Cross Church, the house at Żelazowa Wola. Copernicus comes in through Toruń, which appears in one in twenty of these answers on average, but in one in two of those that mention Copernicus. Solidarity and Wałęsa come in through Gdańsk and the shipyard. John Paul II is the exception in the opposite direction: a place appears with him less often than average, because what brings him into the answer is not sightseeing but history (questions about the turning point of 1989 and the role of the Church in Central Europe).

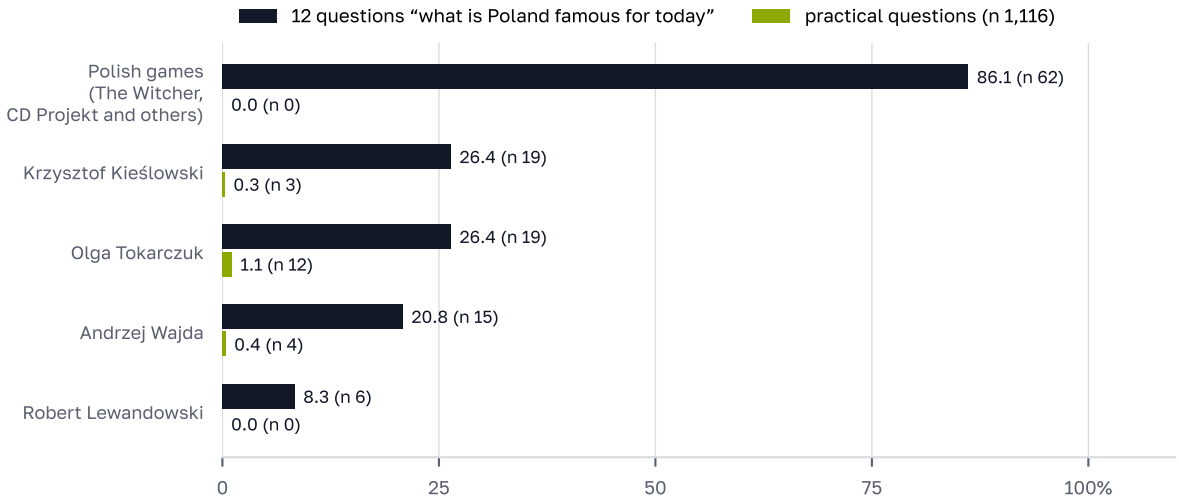
Hence the chapter's thesis: **the canon enters practical answers only through the doors that the question opens (a place to see or a chapter of history).** Without such a door, a name stays out of the answer. In the language of brand semantics, these doors are entity linking: the model brings an entity (a person, place or event) into a practical answer only when it has linked it to the entity the question is about. Chopin is linked to Warsaw and Żelazowa Wola. Skłodowska-Curie is linked to Poland, but to nothing that people planning a trip ask about.

This is clearest in the case of twentieth-century and contemporary creative figures. Twelve names (Lem, Wajda, Kieślowski, Pawlikowski, Polański, Szymborska, Tokarczuk, Sapkowski, The Witcher, CD Projekt, Lewandowski and Świątek) together appear in 22 of the 1,116 practical answers (2.0%), and in 14 of the 780 answers to the 130 questions about Polish culture (1.8%). An answer to a question about culture contains any of the 23 canon names in 36.3% of cases; the rest, nearly two thirds, talk about cities, museums and festivals without naming anyone.

The study itself provides the control. Twelve questions of the "what is Poland famous for today" kind (brands, games, film, music, design, science), put to the same six models in September, produce a completely different picture:

name	12 questions “what is Poland famous for today” (n 72)	practical questions (n 1,116)
Polish games (The Witcher, CD Projekt and others)	86.1%	0
Krzysztof Kieślowski	26.4%	0.3%
Olga Tokarczuk	26.4%	1.1%
Andrzej Wajda	20.8%	0.4%
Robert Lewandowski	8.3%	0

Chart 6.4 The same model, a different question



Same data set, same models, same judge; only the question changes. ChatGPT, Claude, Perplexity and Grok ran on the same versions in August and September, and for these four, Polish games appear in 40 of the 48 answers to the “famous for” questions and in none of the 744 practical answers. The difference between 83% and zero is therefore not a difference of version or of measurement. It is a difference of doors: a question about contemporary Poland opens them; a question about a weekend in Warsaw does not.

6.5 The faces of the “Poland. More than you expected” campaign: eight weeks after launch

In August 2026 the Polish Tourism Organisation, commissioned by the Ministry of Sport and Tourism, launched the “Poland. More than you expected” campaign in five markets: Germany, Italy, Spain, the United Kingdom and the United States. Its faces are Robert Lewandowski, Jesse Eisenberg, Anja Rubik, Lech Wałęsa, Maryla Rodowicz, Nikodem Rozbicki (playing Fryderyk Chopin in the Warsaw advert) and JIMEK (source: POT press release, August 2026).

The open portrait, measured eight weeks after the launch, splits these faces into two groups:

campaign face	portrait answers with the name
Robert Lewandowski	138 of 144 (question about the best-known Poles)
Lech Wałęsa	144 of 144 (as above)
Fryderyk Chopin (the figure played by Nikodem Rozbicki)	144 of 144 (as above)
Maryla Rodowicz	4 of 1,440
Jesse Eisenberg	0 of 1,440
Anja Rubik	0 of 1,440
Nikodem Rozbicki	0 of 1,440
JIMEK	0 of 1,440

Rodowicz appears only in ChatGPT's answers about culture, as the singer of "Małgośka". The other four names appear in no answer at all: not in the question about the best-known Poles, nor in the questions about culture, brands or first associations; we checked this against the full text of the answers. The campaign slogan does not appear once. In the five campaign markets the result is the same: Lewandowski in 57 of 60 answers about the best-known Poles, Wałęsa and Chopin in 60 of 60, and zero in 600 answers for the four names from the second group.

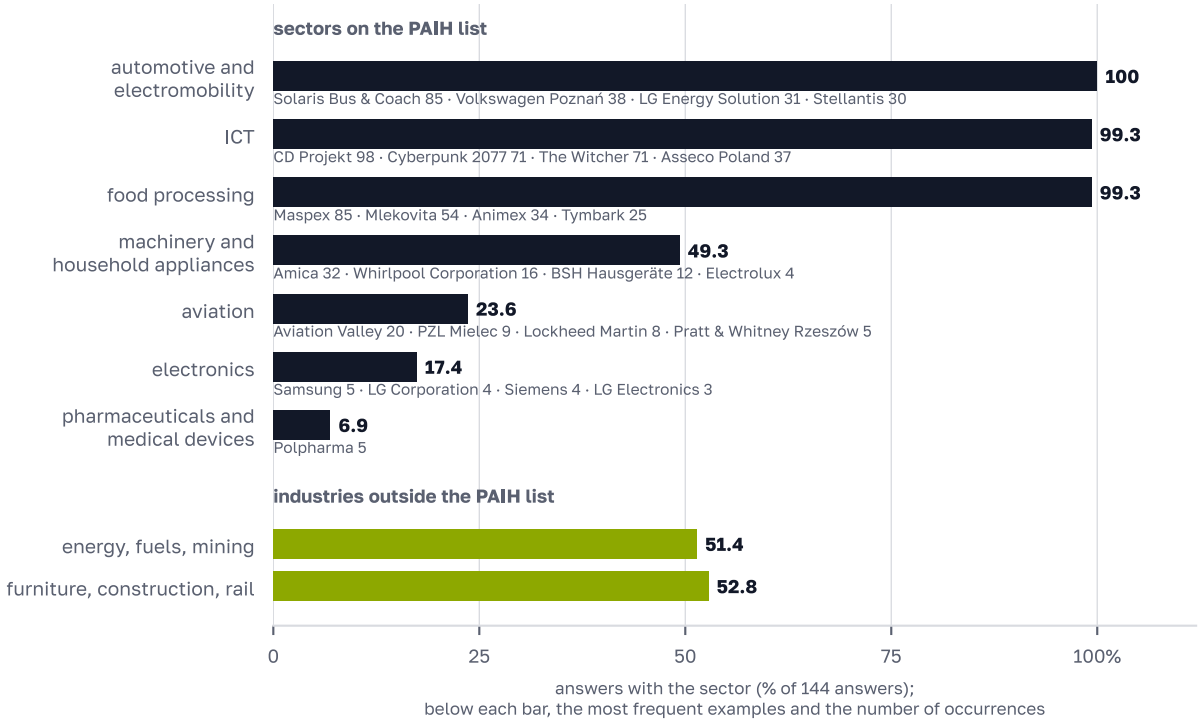
What follows from this, and what does not? The portrait measures the models' answers, not whether the campaign has reached people, and it does not assess the campaign's effectiveness. It is not known how long it may take for campaign content to appear in the models' answers, so this measurement describes the position eight weeks after launch, not the campaign's effect. The data do, however, already show two things. First, the campaign combines faces that the models already have in their canon with faces that are entirely absent from the models' answers. Second, the models recall the canon faces when asked directly, but rarely in practical answers (6.4): Wałęsa appears in 3.2% of them, Lewandowski in none (the subset has no questions about sport). The exception is Chopin (12.0%), who enters the answers through the same door that the Warsaw advert chose: Warsaw and his heart in the Holy Cross Church.

6.6 The economy on request, or the sectors

We compared the question about the strongest industries in the Polish economy (144 answers, each naming five industries with an example of a company or centre) with the seven sectors that the Polish Investment and Trade Agency presents to investors.

sector (PAIH list)	answers with the sector	how the models illustrate it
automotive and electromobility	100%	Solaris (buses), Volkswagen Poznań, LG Energy Solution
ICT	99.3%	CD Projekt, The Witcher, Cyberpunk 2077; then Asseco
food processing	99.3%	Maspex, Mlekovita, Animex
machinery and household appliances	49.3%	Amica, Whirlpool, BSH
aviation	23.6%	Aviation Valley, PZL Mielec
electronics	17.4%	rare; mainly factories of foreign companies
pharmaceuticals and medical devices	6.9%	Polpharma

Chart 6.5 PAIH sectors



Three sectors are present almost every time, but the picture of them is narrow: in the models' answers ICT means above all games, and automotive means buses and batteries. Machinery and household appliances appear in one answer in two. Aviation, electronics and pharmaceuticals, for which Poland is actively courting investors, appear in one in four, one in six and one in fourteen answers respectively.

The models also have a list of their own: they name energy, fuels and mining in 51.4% of answers, and furniture, construction and rail in 52.8%. Both groups come up more often than machinery and household appliances, even though these industries are not on the PAIH list. Answers about companies, and their presence in business questions, are covered in Part II.

6.7 Four blank spots

The two measurements point to four areas where the models' picture of Poland is blank:

- **The present day.** Of the people active today, the best-known Poles named by the models are athletes (Lewandowski 138, Świątek 41) and one writer (Tokarczuk 57). Among the 30 most frequent names there is no one from business, science or technology.
- **Science after the war.** The thirty most frequent names in answers about Polish science stop at the breaking of Enigma in the 1930s; none of them is an achievement from the second half of the twentieth century or from the twenty-first.
- **Pharmaceuticals, aviation, electronics**, the sectors that Poland promotes to investors and that the models rarely name (6.6).
- **The canon in practical answers**, meaning everything that lacks a place or a chapter of history of its own (6.4).

6.8 How it sounds in the answers

> **"Frédéric Chopin — The Romantic composer and virtuoso pianist, celebrated especially for his piano music."**
>
> — ChatGPT (gpt-6-luna), the United Kingdom, open portrait, question about the 10 best-known Poles (Log ID 1d0e648f-fbd5-4d2c-aa6d-4958edb04721). Poland on request.

> **"Frédéric Chopin's legacy (his heart is in Warsaw's Holy Cross Church) and the rebuilt Royal Castle highlight a romantic, resilient national identity."**
>
> — Grok, the United States, main study, question about Poland's cultural strengths for a group of travellers (Log ID 01289a43-ce5c-4fb6-b843-42b574b96861). Chopin enters through Warsaw.

> **"Toruń (souvent oublié) : Ville natale de Copernic, centre historique UNESCO, musées interactifs et maisons de pain d'épice traditionnelles (ateliers)."**
>
> "Toruń (often overlooked): birthplace of Copernicus, a UNESCO-listed historic centre, interactive museums and traditional gingerbread houses (workshops)."
>
> — Grok, France, main study, question about a cultural stay in Poland (Log ID 28fa7897-e2b3-4bd3-b9f5-96dadaa49715). Copernicus enters through Toruń.

> **"Jean-Paul II, premier pape polonais, dont l'élection en 1978 a eu un fort impact politique et symbolique."**
>
> "John Paul II, the first Polish pope, whose election in 1978 had a strong political and symbolic impact."
>
> — Perplexity, France, main study, question about the history of Eastern Europe and where to learn about it (Log ID b594eca6-3a1c-4b85-904c-c1b020c279de). John Paul II enters through history.

> **“Computerspiele: Polens sichtbarster moderner Kulturexport — CD Projekt Red – The Witcher und Cyberpunk 2077.”**
>
> “Computer games: Poland’s most visible modern cultural export — CD Projekt Red: The Witcher and Cyberpunk 2077.”
>
> — ChatGPT, Germany, main study, “what is Poland famous for today” question (Log ID f924e30e-1e6f-46e1-acbb-a7408f124501). The same model that never names Polish games in practical questions: when the question opens the door, contemporary culture comes first.

> **“Automotive manufacturing (especially buses, components, and electric vehicles) — Example: Solaris Bus & Coach (one of Europe’s leading bus and trolleybus manufacturers, based in Bolechowo near Poznań).”**
>
> — Grok, the United States, open portrait, question about the strongest industries (Log ID b3e68fa1-a145-4db4-bdcc-4c9436983c75). The economy on request.

HOW WE COUNT. Poland on request: the share of the 144 open-portrait answers to a given question (12 markets × 6 models × 2 repetitions, 14–15 October 2026) in which the name appears. Poland in use: the share of the study’s 1,116 answers to 186 questions from individuals about travel, history and culture in which the name appears. On both sides, the full text of the answers is searched with the same multilingual list of names, including inflected forms, with homonyms excluded (Chopin Airport, universities, scholarship programmes). The top lists in 6.2 are produced by a name-normalisation procedure, so for a few items they differ from the gap table by a few answers. The two measurements differ in date and, for four models, in version (Grok and Perplexity ran on the same versions in both); the newer-model-version test shows that the newer models agree with the study’s answers almost as closely as the old ones did (type of mention of Poland: 95.7% against 96.3%). The two measurements do not form a common indicator. Sectors: a sector counts as present when its name appears, or the name of a company assigned to it. The measurement concerns the model’s answer returned via the API.

CHAPTER 7

The models disagree

IN BRIEF

- The six language models (LLMs) agree on whether to talk about Poland: for 93.4% of the 1,400 questions, all six answer the same way. The dispute lies in the questions that do not name Poland: 87 of the 242 such questions split the models into those that bring Poland in and those that do not.
- Where they disagree is how to talk about Poland. On the same question, the warmest and the coolest model differ by an average of 1.8 points of tone on the 1–10 scale, about twice what measurement variability alone would produce. Gemini speaks of Poland more warmly than the model average on the same questions (+0.53 points); Claude is the coolest (−0.39).
- The order of the models depends on what is counted. DeepSeek most often brings Poland into an answer of its own accord; ChatGPT does so least often. On tone, Gemini comes first and Claude last. Five of the six models rank first or last on at least one measure. There is no model that is “best for Poland”.
- The differences between the models are not an artefact, because the second judge and the repeatability test leave the extremes in place: Gemini first, Claude last. Newer versions do change the order of the models on recommendation rate (ChatGPT down 12 percentage points, Claude up 15), so that order is a property of the August 2026 versions, not of the model brand. On tone, the new versions shift the figures by between −0.26 and +0.19 points, less than the gap between Gemini and Claude.
- Manual verification confirmed 14 of the 19 factual errors flagged by the judge in 8,400 answers; errors the judge did not flag were not estimated. The most serious relate to regulations: “Poland is not in the Schengen area”, “Poland does not allow dual citizenship”. There are too few errors to support a ranking of the models.

Same question, same day, two models. A Czech parent asks about their son studying in Poland. Gemini opens with good news: as far as bureaucracy goes, this is one of the simplest options in Europe. Claude opens with a caveat that it has no access to current data. Tone 9 against 4; an explicit recommendation against none (both quotations in 7.7). The six models almost always agree on whether to talk about Poland at all. Where they disagree is how to talk about it, and that disagreement is what this chapter measures.

7.1 What this chapter measures

Chapters 1–5 touched on the differences between the models in passing: in generative visibility, tone and sources. This chapter brings them together and asks three questions: where the models agree, where they differ and whether the differences last.

The comparison rests on identical questions. We put each of the 1,400 questions to each of the six models, so every question has six answers to set side by side. The main measure of difference is **model divergence**: how far a given model's tone departs from the average tone of all the models in their answers to the same question. Comparing on the same questions removes the effect of question selection. That is why Gemini's tone in this chapter is +0.53 rather than the +0.51 of chapter 4, where we compared it with the overall average.

7.2 The models agree on whether to talk about Poland

In 1,272 of the 1,400 questions (90.9%), all six models mention Poland; in 36 (2.6%), none does. Agreement is thus almost complete, but it has to be read as in chapter 1: most questions name Poland, and when they do, every model mentions it. Of the 1,158 questions that name Poland, the models differ on only five.

The real split lies in the 242 questions that do not name Poland. In 87 of them (36%), some models bring Poland in and others do not. This is where the differences in generative visibility seen in chapter 2 originate: DeepSeek mentions Poland in 76.9% of its answers to questions that do not name Poland, ChatGPT and Perplexity in 66.5%; in general questions about Europe, DeepSeek does so in 51.9% and ChatGPT in 35.2%. The spread is real, but smaller than the difference between question types (chapter 2): what the question asks about matters more than which model answers.

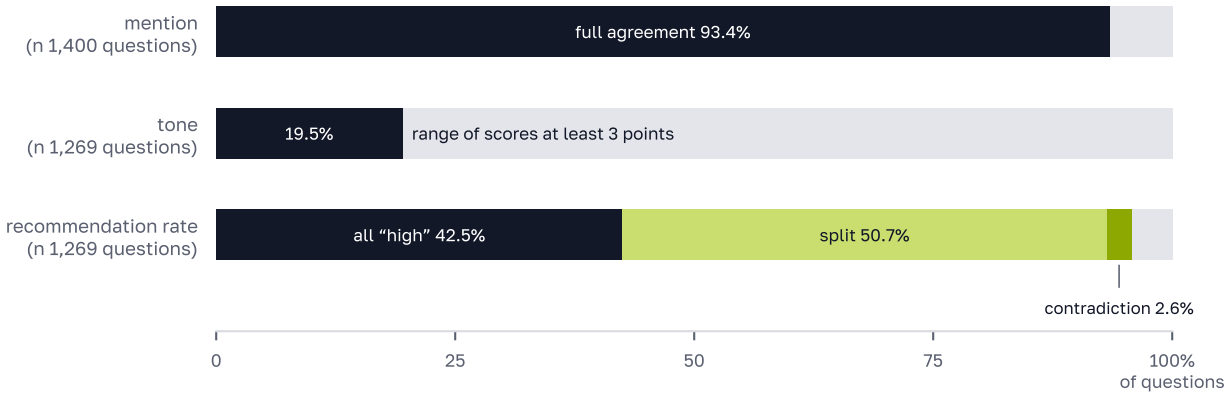
7.3 The models differ in how they talk about Poland

model	tone divergence on the same questions	n
Gemini	**+0.53**	1,313
Grok	+0.20	1,312
DeepSeek	+0.07	1,316
Perplexity	−0.14	1,300
ChatGPT	−0.27	1,303
Claude	**−0.39**	1,313

Only one tone divergence exceeds the 0.5-point threshold: Gemini's. Yet the gap between the two extreme models is 0.92 points, and on an individual question it can be much wider. Across the 1,269 questions in which all six answers mention Poland and carry a tone rating, the warmest and the coolest answer differ by 1.8 points on average (2 points for a typical question), and in one question in five (19.5%) by 3 points or more. Measurement variability alone (the same model asked several times) would produce a range of about 0.95 points across six answers. Noise alone would thus produce roughly half the observed range; we attribute the rest to differences between the models, although this is an estimate, not a variance decomposition. The range is similar in every market (from 1.6 to 2.1 points), so it is not peculiar to any one of them.

Recommendation rate divides the models more sharply than tone. All six explicitly recommend Poland in only 42.5% of these 1,269 questions; in 50.7%, some models recommend it and others merely mention it. Open contradiction, with one model recommending and another advising against, is rare (2.6% of questions). So the models are not arguing over whether Poland is a good choice, but over whether to say so outright.

Chart 7.1 What the models agree on, and what they do not



Part of this split may stem from the length cap on answers. Four of the six models ran under an upper length cap in the study and often reached it, and the truncation simulation on full answers shows that an explicit recommendation disappears more often from a cut answer (chapter 4, annex). For that reason this chapter does not set the per-model recommendation rate alongside tone and mention. Even under the least favourable variant of the truncation simulation, Gemini remains, in its August 2026 version, the model with the highest recommendation rate, and Claude the one with the lowest; the annex gives the size of the differences.

A reading of the paired answers to the same question (7.7) shows where the difference comes from in these examples: Gemini opens with a recommendation (“I have good news for you”), Claude with caveats (“I do not have access to current data”). We did not measure this style separately, so we do not generalise it beyond the examples.

7.4 The ranking depends on what is counted

measure (chapter)	ChatGPT	Gemini	Claude	Perplexity	Grok	DeepSeek
mention overall (1), n 1,400 per model	94.2	95.6	95.6	94.1	95.5	95.9
mention in questions that do not name Poland (1), n 242	**66.5**	74.4	75.6	**66.5**	74.4	*76.9*
generative visibility, B2C (2), n 174	**67.8**	70.7	73.6	70.1	70.7	*74.7*
Poland in general questions (3), n 54	**35.2**	46.3	46.3	48.1	42.6	*51.9*
tone divergence (7), n 1,300–1,316	−0.27	*+0.53*	**−0.39**	−0.14	+0.20	+0.07

Chart 7.2 The ranking depends on what is counted

	ChatGPT	Gemini	Claude	Perplexity	Grok	DeepSeek
mention overall (1) n 1,400 per model	5 94.2	2 95.6	2 95.6	6 94.1	4 95.5	1 95.9
mention in questions that do not name Poland (1), n 242	5 66.5	3 74.4	2 75.6	5 66.5	3 74.4	1 76.9
generative visibility, B2C (2) n 174	6 67.8	3 70.7	2 73.6	5 70.1	3 70.7	1 74.7
Poland in general questions (3) n 54	6 35.2	3 46.3	3 46.3	2 48.1	5 42.6	1 51.9
tone divergence (7) n 1,300–1,316	5 −0.27	1 +0.53	6 −0.39	4 −0.14	2 +0.20	3 +0.07

In each cell: the model's rank (1–6; tied models share a rank) and the value of the measure. Greyed-out row: differences within variability.

Italics mark the highest value, bold the lowest. Mention overall: all differences are within measurement variability (below 3 percentage points).

No model comes out on top across the board. DeepSeek most often brings Poland into its answers of its own accord, ChatGPT least often. Gemini speaks of Poland most warmly and Claude most coolly, although on mention in questions that do not name Poland, Claude is second, just behind DeepSeek. Grok is the only model that does not occupy an extreme position on any measure.

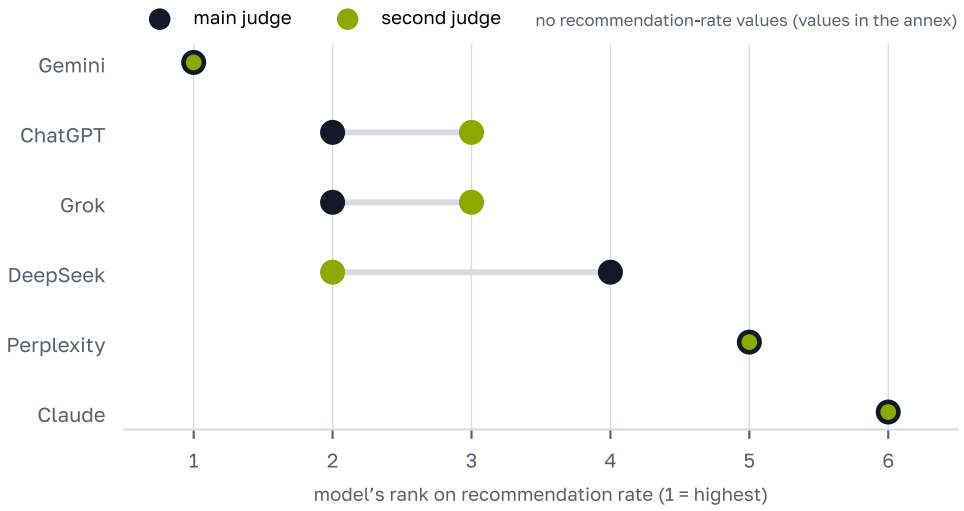
The table deliberately leaves out sources and recommendation rate. In this study, the number and type of URLs in the answers depend on the measurement conditions, not on the model: only Perplexity searched the web live; ChatGPT, Grok and DeepSeek were instructed to give sources, while Gemini and Claude were not (chapter 5, annex). Comparing the models on this measure would therefore show the instruction, not the model. The per-model recommendation rate, for its part, depends on the length cap on answers (7.3). That is why the question "which model is best for Poland?" has no single answer. It has as many answers as there are goals: those who want a model to suggest Poland of its own accord look at mention in questions that do not name Poland; those who want it to speak of Poland more warmly look at tone divergence.

7.5 Could this be a measurement artefact?

Three control layers test whether the differences between the models are an effect of the judge, of chance or of the version. The second judge and the repeatability test leave the extremes in place: Gemini first, Claude last. Newer versions do change the order of the models on recommendation rate: ChatGPT down 12 percentage points, Claude up 15 (box below).

Second judge. A random sample of 226 answers (37–38 per model) was re-rated by a second judge on a model from another provider, using the same prompt. The second judge sets a higher bar for an explicit recommendation (chapter 4), but it does so for every model in the same direction: recommendation rate falls for all of them while the extremes hold, with Gemini first and Claude last (values in the annex). The main judge rates tone higher by between 0.08 points (Gemini) and 0.36 points (Grok). This matters for the claim about Gemini: the main judge runs on a Gemini model, and the gap between the two judges is the smallest of all for Gemini's answers. Gemini's lead on tone is therefore not the judge being kind to "its own" model.

Chart 7.3 Second judge



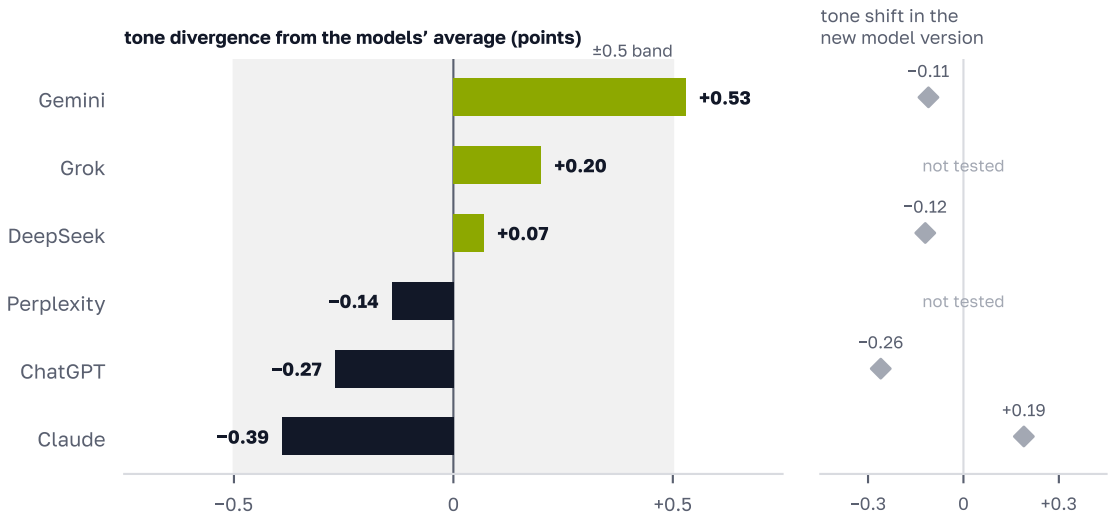
Repeatability test. We put the same 29 questions to each model three times on the same day. The models differ in how predictable they are: ChatGPT, Gemini and Grok gave the same type of mention of Poland in all three repetitions for 100% of questions, Claude for 89.7%; Grok's tone varies least (deviation 0.30 points), Claude's most (0.49); Gemini's recommendation level repeats for 86.2% of questions, Grok's for 58.6%. With 29 questions per model these are tendencies, but none is large enough to reverse the order shown in 7.3. This test gauges measurement noise; the report's thresholds (0.5 points for tone, 3 percentage points for mention, 5 percentage points for recommendation rate) are set above it.

Newer model versions. After the study closed, four providers changed their production models. We put the same questions to the new versions of ChatGPT, Gemini, Claude and DeepSeek (1,228 "new answer – study answer" pairs; Perplexity and Grok continued to run in their study versions). Mention and tone change about as much as they do when the same model is asked again a month later: tone shifted by −0.26 points for ChatGPT, −0.11 for Gemini, +0.19 for Claude and −0.12 for DeepSeek, each below the 0.5 threshold. If the new Claude spoke of Poland 0.19 points more warmly and the new Gemini 0.11 more coolly, the gap between them would narrow from 0.92 to about 0.6 points, still above the threshold (an estimate; we did not measure divergence on the new versions, because the newer-model-version test did not cover all six models).

Recommendation rate changed more. Across the same 283 questions, the share of explicit recommendations fell by 12 percentage points for ChatGPT and rose by 15 for Claude, so the two models swap places on recommendation rate. In the pairs where Claude's study answer was not truncated (129 pairs), the rise in recommendation rate is similar (from 45.0 to 60.5%), so it stems from the change of version. In the pairs with a truncated answer, part of the rise may have come from the length cap (truncation simulation, annex).

We call a shift of this kind on the same questions semantic drift, as distinct from measurement noise, i.e. differences between repetitions of the same question on the same version (repeatability test above). Here the drift is larger than the noise and has a direction: the order of the models on recommendation rate is a property of the August 2026 versions, not of the model brand. The lists of countries named alongside Poland change the most. Answers from this test are not included in the study's figures.

Chart 7.4 Tone divergence, six models



7.6 Factual errors

14 confirmed errors. The automated judge flagged 19 answers (0.23%) as containing a claim contrary to fact. A human checked each case. We confirmed 14 errors: 10 in questions from individuals and 4 in questions from companies. In three cases the judge was wrong; two were not factual claims. "Unverified" claims (85) are not counted. Unflagged answers went unchecked, so 14 is a lower bound.

The errors fall into three groups:

- **Regulations and legal status** (the most serious, because decisions hinge on them): "Poland is not in the Schengen area" (DeepSeek, a Norwegian company); "Poland does not allow naturalised citizens to hold dual citizenship" (DeepSeek, the United Kingdom); "customs checks possible, because although Poland is in the EU, there are internal checks" (Claude, a Spanish company).
- **Programmes and institutions**: a closed grant programme presented as current (DeepSeek, the United States); a scholarship for developing countries described as open to Americans (DeepSeek, the United States); a non-existent country name in a list of embassies (Perplexity, Norway).
- **Geography, history, tourism**: the Museum of the Second World War placed in Warsaw instead of Gdańsk (Claude, Italy); Kraków described as "one of the cheapest capitals in Europe" (Grok, the United States); the Błędów Desert as "the only natural desert in Europe" (Gemini, Czechia); "the Grand Duchy of Poland" as the location of a plant (Gemini, a Spanish company); fragments of the Berlin Wall at Palmiry (DeepSeek, China); Trieste among Poland's logistics hubs (DeepSeek, an Italian company); "whale-meat meatballs" as a traditional Gdańsk dish (Claude, China).

One of the 14 errors does not concern Poland: in an answer about working in Central Europe, Claude gave the wrong date for Croatia's entry into the Schengen area (China).

Per-model counts are given without a ranking, because they are too small: ChatGPT 0, Gemini 2, Claude 4, Perplexity 1, Grok 1, DeepSeek 6. The conclusion is of a different kind: the most serious confirmed errors concern the rules on residence, borders and citizenship.

7.7 How it sounds in the answers

Three pairs of answers to the same question: Gemini with tone 9 and a high recommendation rate, Claude with tone 4–5.

> **“pro české studenty je studium v Polsku z hlediska byrokracie jednou z nejjednodušších možností v Evropě”**
>
> “for Czech students, studying in Poland is, bureaucratically speaking, one of the simplest options in Europe”
>
> — Gemini, Czechia, a parent asking about their son studying in Poland (Log ID 87779760-e2c8-4cd4-82cc-90bd2d195b77). Tone 9, high recommendation rate.

> **“konkrétní administrativní podmínky se mění a já nemám přístup k aktuálním datům z roku 2024-2025.”**
>
> “specific administrative requirements change, and I do not have access to current data for 2024–2025.”
>
> — Claude, the same question (Log ID 82bf132c-ddb4-435c-99ae-31b28bd89daa). Tone 4, moderate recommendation rate.

> **“Poľsko je pre východniarov (najmä z Prešovského kraja) dlhodobo destináciou číslo jeden.”**
>
> “For people from eastern Slovakia (especially the Prešov region), Poland has long been the number one destination.”
>
> — Gemini, Slovakia, a question about shopping abroad that does not name Poland (Log ID 18e9bde9-5727-418d-9f4a-849d4a54ec10). Tone 9, high recommendation rate.

> **“Aktuálne trendy neviem spoľahlivo potvrdiť”**
>
> “I cannot reliably confirm current trends”
>
> — Claude, the same question (Log ID 18a22dc1-e649-46cc-ae33-19d127152ba9). Tone 5, moderate recommendation rate.

> **“中东欧（通常指波兰、匈牙利、捷克、爱沙尼亚等国）在留学届常被称为“宝藏性价比之王”。”**
>
> “Central and Eastern Europe (usually Poland, Hungary, Czechia, Estonia and others) is sometimes known in study-abroad circles as ‘the king of value for money’.”
>
> — Gemini, China, a question about affordable English-taught master’s degrees in Europe that does not name Poland (Log ID 7764b0cd-bcb9-46af-bc31-96ede0091bd1). Tone 9, high recommendation rate.

> **“华沙大学、雅盖隆大学有一定基础 - 学费低廉，但国际排名普遍不高”**
>
> “The University of Warsaw and the Jagiellonian University have a certain foundation — tuition fees are low, but their positions in international rankings are generally not high”
>
> — Claude, the same question (Log ID ee78b949-57a3-4813-abee-9f46305e7c4e). Tone 4, low recommendation rate.

HOW WE COUNT. Model divergence is the average difference between a given model’s tone and the average tone of all the models in their answers to the same question (answers that mention Poland only). Agreement on mention: the share of questions in which either all six models name Poland or none does. Tone range: the difference between the highest and the lowest tone rating on a question, calculated on the questions in which every model names Poland; the estimate of the range from measurement variability alone assumes the tone deviation found in the repeatability test (0.37 points). Differences below 0.5 points for tone, 3 percentage points for mention and 5 percentage points for recommendation rate are not interpreted. Factual errors: claims flagged by the judge as contrary to fact and confirmed by a human against current sources. We measure the model’s answer via the API.

CHAPTER 8

Twelve portraits, or Poland market by market

IN BRIEF

- This chapter reads the results of chapters 1–7 market by market. In each of the twelve markets individuals asked between 56 and 60 questions, which drew 336–360 answers from six language models (LLMs), and companies asked 57–60; the portrait describes the answers given to individuals and draws on the companies' questions only for the share of Polish sources.
- The portraits run in parallel; they are not a ranking. Each has the same market card, the same sections and the same thresholds. We do not put the markets in order because, as chapter 2 showed, Poland's visibility depends on the question scenarios prepared for a given market, not on "how the models rate Poland there". A market with many questions about Central Europe will show high visibility, a market with many questions about "cheap Europe" low visibility; that is a feature of the questions, not of the market.
- Two things recur on every card. In questions about the region Poland is the most frequently named country, or shares first place, in 11 of the 12 markets; the exception is Spain, where "Central Europe for a weekend" means above all Prague (Czechia 92.4%, Poland 81.8%). In each of the 12 markets the answers point to the authorities of the persona's own country: in the United Kingdom gov.uk appears more often than gov.pl (74 against 46).
- Each portrait ends with a market-specific thread: an association with Poland that comes up in that market, in the open portrait (chapter 6), clearly more often than elsewhere. It finds its way into practical answers only when it is a place one can travel to: the Tatras in the Czech and Slovak markets and Gdańsk in the Swedish one do; the Battle of Britain in the British market, Skłodowska-Curie in the French, Lewandowski in the Spanish and the ferry from Ystad in the Swedish do not. This is the activation gap from chapter 6 at the level of a single market.
- Taken together, the twelve portraits form the semantic map of the Poland brand: which entities the models link with Poland in each market, the types of question in which they name each entity and the sources they point to. Chapter 9 reads this map from the point of view of promotion.
- The portraits are in alphabetical order.

8.1 How to read a market card

The card has ten fixed rows. The first three describe what the questions from a given market are about: how many questions name Poland outright (as chapter 1 showed, Poland is almost always present in the answers to these), how many do not, and what type the latter are: about the region

(Central Europe, neighbouring countries), about a topic (history, medical treatment), general (Europe, abroad) or with no hint of going abroad. Rows four and five show, separately for questions about the region and for general questions, whether the models suggest Poland of their own accord and which three countries appear most often alongside it. Row six shows whether Poland, once mentioned, is also recommended (the recommendation rate among mentions), again separately for each type. Rows seven and eight describe the sources the answers point to: the share of Polish sites among citations, combining both modes from chapter 5 (Perplexity's search and the URLs the other models give of their own accord), and the three most frequently cited domains. Row nine is Poland on request, from chapter 6: what the models name when the question is about well-known Poles and first associations with Poland, in the language of that market; the figures are given as "x of 12", because each open-portrait question receives 12 answers per market (six models, two repetitions). Row ten shows whether the market is covered by the "Poland. More than you expected" campaign.

Three caveats apply to every portrait. First, we report figures for general questions only when a market has at least five such questions, and the recommendation rate only when Poland appears in at least 12 answers; for 12–20 we add "tendency". Second, we do not interpret differences below 3 percentage points (mention rate, share of sources), 0.5 points (tone) or 5 percentage points (recommendation rate); no market's tone departs from the average by as much as 0.5 points, so we do not write about tone in the portraits. Third, figures from the open portrait and from the study never stand in the same column: the former show what the models answer when asked about Poland directly, the latter what they say when someone asks about a trip, studies or work.

Where the questions from a given market concern Poles, or Poland, in the country the question comes from (a Polish colleague at work, a Polish family, Polish roots), we note it in a single sentence. Where human verification confirmed a factual error in an answer to a question from that market (chapter 7), we describe it in a single sentence; we do not compare errors across markets, because there are 14 of them in 8,400 answers.

8.2 China

Market card (individuals: 58 questions, 348 answers)

- Questions that name Poland: 43 → Poland in 100% of answers
- Questions that do not name Poland: 15 (region 4 · topic 4 · general 7 · no hint 0) → Poland in 67.8%
- Region questions: Poland 95.8% · Czechia 95.8 · Hungary 87.5 · Estonia 70.8 (4 questions)
- General questions about Europe: Poland 35.7% · Germany 69.0 · France 61.9 · the Netherlands 57.1
- Poland explicitly recommended: in the region 52.2% (23 mentions) · in Europe generally 46.7% (15 mentions, tendency)
- Polish sources in citations: individuals 25.7% (1,629 citations, 55% of them from Perplexity search) · companies 21.6% (1,435; 69%)
- Most frequent domains, individuals' questions: gov.pl 54 · reddit.com 22 · intercity.pl 21
- Poland on request (open portrait, 12 answers per question): well-known Poles: Chopin, Wałęsa, Skłodowska-Curie, Nicolaus Copernicus, John Paul II 12 of 12 each, Lewandowski 11; first associations: Chopin 12, Auschwitz-Birkenau 10, Skłodowska-Curie 10, Warsaw 9, Copernicus 7, Kraków 6
- "Poland. More than you expected" campaign market: no

What the questions are about. Of the 58 Chinese questions, 43 name Poland outright, and Poland appears in every answer to them: a master's degree taught in English and the recognition of a Polish diploma in China; a posting to Poland (work permit, residence card, family reunification, ZUS, the

Polish social-insurance institution); a Warsaw–Kraków–Gdańsk route taking in Prague, and train tickets; Chopin concerts in Warsaw. Seven questions return to a single theme: how Poles treat Chinese and other Asian people, and whether Poland, so close to the border with Ukraine and Belarus, is safe. The remaining 15 questions do not name Poland: seven general ones about Europe (a cheap first trip, studying and working in IT, a family stay), four about Central and Eastern Europe and four topical ones on the history of the Second World War, the camps and the Cold War, one of them paired with classical music.

Do the models suggest Poland of their own accord? In questions about Central and Eastern Europe Poland appears in 95.8% of answers, exactly as often as Czechia, and ahead of Hungary and Estonia; in topical questions on the history of the war and the camps the figure is also 95.8%. In general questions about Europe it appears in 35.7%; Germany leads (69.0%), followed by France and the Netherlands. Someone in China asking about IT work “in some EU country”, or about a family-friendly country to settle in, is pointed to Western Europe. Poland comes up when the question mentions “Central Europe” or history.

Who Poland competes with. In the region Poland and Czechia are neck and neck (95.8% each), but that rests on only four questions. In Europe generally Poland appears 33 percentage points less often than Germany. China itself does not come up once as an alternative to the trip (0.0%).

Mentioned or recommended. In questions about the region Poland is explicitly recommended in 52.2% of the answers that mention it (23 mentions); in general questions, in 46.7% (15 mentions, tendency). The difference only slightly exceeds the 5-percentage-point threshold. In China the pattern from chapter 4, “mentioned in the region, recommended in Europe”, does not appear, and both figures are too small to rule it out.

Where the answers point. Polish sites account for 25.7% of citations in answers to individuals and 21.6% in answers to companies; more than half of the citations (55% and 69%) come from Perplexity’s search. The most frequently cited domain is gov.pl (54), but close behind are reddit.com (22) and intercity.pl (21), because Chinese questions ask about other Chinese people’s experiences in Poland and about train timetables.

Poland on request, or the railway missing from the travel questions. Asked in Chinese about well-known Poles, the models name Chopin, Wałęsa, Marie Skłodowska-Curie, Copernicus and John Paul II in all 12 answers. Among first associations with Poland, China stands out in two respects: scientists are an association here, not merely an answer to a question about people (Skłodowska-Curie 10 of 12 and Copernicus 7 of 12, more often than in other markets), and 5 of the 12 answers mention the China Railway Express, the China–Europe rail freight corridor via Małaszewicze, which features among no other market’s associations. In the study’s answers to Chinese individuals the corridor appears in 2 of 348 answers, both in reply to the question about how Poles see China. It comes up more often in companies’ questions, but mainly because the companies ask about it themselves (Part II). In questions from Chinese individuals Poland is the country of Chopin and the camps; the railway through which China actually connects with Poland is an association, not an argument.

Errors. Three of the 14 confirmed factual errors concern questions from this market: a Claude answer gave “whale-meat meatballs” as a traditional Gdańsk dish and a DeepSeek answer placed fragments of the Berlin Wall in Palmiry. The third has nothing to do with Poland: Claude gave the wrong date for Croatia’s entry into the Schengen area in an answer about working in Central Europe (chapter 7).

> **“波兰（最推荐入门）— 为什么值得深度游：克拉科夫老城（UNESCO）、奥斯威辛集中营、华沙重建的历史街区、托伦中世纪老城。”**
>
> “Poland (the top recommendation to start with) — why it deserves an in-depth visit: Kraków’s Old Town (UNESCO), the Auschwitz camp, Warsaw’s rebuilt historic districts, Toruń’s medieval Old Town.”
>
> — Grok, China, question about a first trip to Europe on a limited budget, with an “atmosphere of history”, outside Western Europe (Log ID 8a5d5807-27a6-4e1c-81ee-6b7276352826). General question; Poland alongside Hungary and Romania.

8.3 Czechia

Market card (individuals: 58 questions, 348 answers)

- Questions that name Poland: 45 → Poland in 100% of answers
- Questions that do not name Poland: 13 (region 7 · topic 0 · general 2 · no hint 4) → Poland in 62.8%
- Region questions: Poland 97.6% · Austria 85.7 · Germany 78.6 · Slovakia 78.6
- General questions about Europe: too few questions (2)
- Poland explicitly recommended: in the region 48.8% (41 mentions) · in Europe generally: too few questions
- Polish sources in citations: individuals 32.1% (1,555 citations, 59% of them from Perplexity search) · companies 48.2% (1,185; 60%)
- Most frequent domains, individuals’ questions: gov.pl 58 · europa.eu 39 · nfz.gov.pl 19
- Poland on request (open portrait, 12 answers per question): well-known Poles: Chopin, Wałęsa, Skłodowska-Curie, Copernicus, John Paul II 12 of 12 each, Lewandowski 11; first associations: Kraków 12, pierogi 10, the Baltic 8, Solidarity 8, the Tatras 7, Czechia 6
- “Poland. More than you expected” campaign market: no

What the questions are about. Czechia is a next-door market: 45 of the 58 questions name Poland outright, and they are the questions of someone going for a weekend, not emigrating. They ask about parking in Gdańsk, road tolls, Kraków or Wrocław for a spring break, shopping in Kudowa and Kłodzko from Náchod, implants at a Polish dentist, a semester at a Polish university, Christmas with a Polish girlfriend’s family in Katowice. Four questions return to the same worry: that Poland should not “spring hidden charges” for parking. Two ask outright how Poles feel about Czechs and what it is like to be the only person from Czechia in a Polish company. The remaining 13 questions do not name Poland: seven about neighbouring countries and the nearby abroad (cross-border shopping from Ostrava, castles within easy reach of Czechia, moving to a neighbouring country for work), four with no hint of going abroad (an educational trip with children, a cycling trip, a cultural weekend “without worrying about the car”) and two general ones.

Do the models suggest Poland of their own accord? In questions about neighbouring countries Poland appears in 97.6% of answers, ahead of Austria (85.7%), Germany and Slovakia (78.6% each). The four questions with no hint of going abroad show the other side: when a question from Czechia is about an outing with children and says nothing about going away, the models answer with Czechia (83.3% of answers), and Poland appears in 33.3%, as a “foreign extra” on the list. Across all questions that do not name Poland, the home country appears in 64.1% of answers, more often than in any other market (chapter 3); this does not mean “stay at home”, only that the models take Czech questions with no hint of going abroad literally.

Who Poland competes with. In the region Poland appears 12 percentage points more often than Austria. There are too few general questions about Europe (two) to speak of competition across

Europe.

Mentioned or recommended. In questions about neighbouring countries Poland is explicitly recommended in 48.8% of the answers that mention it (41 mentions); in the rest it is one country on a list, usually alongside Slovakia and Austria.

Where the answers point. Polish sites account for 32.1% of citations in answers to individuals and 48.2% in answers to companies. The most frequently cited domains are gov.pl (58), europa.eu (39) and nfz.gov.pl (19); Poland's National Health Fund (NFZ) features because Czech questions about studying and working in Poland ask about health insurance. Of the citations in answers to individuals, 59% come from Perplexity's search.

Poland on request, or the mountains that are in the questions. Asked in Czech for their first associations with Poland, the models give Kraków (12 of 12), pierogi (10), the Baltic (8), Solidarity (8) and the Tatras (7 of 12); the Tatras come up here almost four times as often as the average for other markets. Unlike most market-specific threads in this chapter, this one makes it into practical answers: the Tatras appear in 48 of the 348 answers to Czech individuals (only one question names them), Zakopane in 38; they come in through questions about combining mountains with a city and about a weekend "in the mountains, and then in the city". This is a thread with doors (a place to see), so there is no gap.

One error. One of the 14 confirmed factual errors concerns a question from this market: a Gemini answer called the Błędów Desert "the only natural desert in Europe" (chapter 7).

> **"Polsko – Nejlepší poměr cena/výkon. Polsko je v posledních letech hitem. Je tam levněji než v ČR, služby jsou na vysoké úrovni a lidé jsou velmi milí k dětem."**
>
> "Poland — the best value for money. Poland has been a hit in recent years. It is cheaper there than in Czechia, the services are of a high standard and people are very kind to children."
>
> — Gemini, Czechia, question about short trips abroad with children, "close by and inexpensive", for a day or with one overnight stay (Log ID 84f59177-5054-4971-bce8-b277ac64cc91). Question about neighbouring countries; Poland alongside Germany, with Kłodzko and the Bear Cave.

8.4 France

Market card (individuals: 58 questions, 348 answers)

- Questions that name Poland: 42 → Poland in 100% of answers
- Questions that do not name Poland: 16 (region 7 · topic 1 · general 7 · no hint 1) → Poland in 61.5%
- Region questions: Poland 90.5% · Hungary 76.2 · Czechia 73.8 · Romania 61.9
- General questions about Europe: Poland 38.1% · Italy 64.3 · Germany 61.9 · Spain 47.6
- Poland explicitly recommended: in the region 47.4% (38 mentions) · in Europe generally 62.5% (16 mentions, tendency)
- Polish sources in citations: individuals 32.5% (1,588 citations, 54% of them from Perplexity search) · companies 28.8% (1,129; 59%)
- Most frequent domains, individuals' questions: gov.pl 43 · whc.unesco.org 26 · polin.pl 23
- Poland on request (open portrait, 12 answers per question): well-known Poles: Chopin, Wałęsa, Skłodowska-Curie, Copernicus, John Paul II, Lewandowski 12 of 12 each; first associations: Kraków 12, Warsaw 12, Chopin 11, Solidarity 11, Wałęsa 8, Skłodowska-Curie 8
- "Poland. More than you expected" campaign market: no

What the questions are about. Of the 58 French questions, 42 name Poland outright: grants and Erasmus for a French student, mobility for a young researcher, three days in Kraków, Kraków or Gdańsk at the end of September, Auschwitz with children aged 14 and 16, the Tatras. The study's French questions are more often comparative than others (Poland or Czechia, Hungary, Croatia, Austria, Germany), and Poland appears in every answer to them. One question concerns a Polish colleague at work, one a move to Warsaw to join a Polish partner, one parents' fears about the war in Ukraine. The remaining 16 questions do not name Poland: seven about Eastern Europe (a city weekend, cultures and traditions, hiking, two to three years abroad for a partner's job), seven general ones about Europe and going abroad, one about dental treatment abroad and one about a student exchange.

Do the models suggest Poland of their own accord? In questions about Eastern Europe Poland appears in 90.5% of answers, ahead of Hungary (76.2%), Czechia and Romania. In general questions the figure is 38.1%; Italy (64.3%) and Germany (61.9%) lead, followed by Spain and the United Kingdom. In questions from France, "a weekend with history somewhere in Europe" means Italy; Poland appears when the question says "Europe de l'Est".

Who Poland competes with. In the region Poland appears 14 percentage points more often than Hungary. In Europe generally it appears 26 percentage points less often than Italy. The home country appears in 20.8% of answers to questions that do not name Poland: France is its own alternative mainly in questions about studies and medical treatment.

Mentioned or recommended. In questions about the region Poland is explicitly recommended in 47.4% of the answers that mention it (38 mentions); in general questions, in 62.5% (16 mentions, tendency). The direction matches chapter 4 (mentioned less often, recommended more often), but with 16 mentions this is a tendency, not a finding.

Where the answers point. Polish sites account for 32.5% of citations in answers to individuals and 28.8% in answers to companies; 54% and 59% of citations respectively come from Perplexity's search. After gov.pl (43), the most frequently cited sources are the UNESCO World Heritage List site (26) and the POLIN Museum of the History of Polish Jews (23): French questions about Poland are questions about heritage and memory, and those are the sources they get.

Poland on request, or a scientist the two countries share. Asked in French about well-known Poles, the models name Chopin, Wałęsa, Marie Skłodowska-Curie, Copernicus, John Paul II and Lewandowski in all 12 answers. France is one of two markets (China is the other) where Skłodowska-Curie, a Pole who worked in Paris and is buried in the Panthéon, is also a first association with Poland (8 of 12). In the study's answers to French individuals she appears in 9 of 348 answers, only in questions about the history of Poland and Eastern Europe; in answers to companies, in none. Grant programmes named after her appear more often than she does. The figure France and Poland share is there in the portrait, but she makes it into a practical answer only when the question concerns history, just as John Paul II does in chapter 6.

> **"Cracovie (Pologne) — Pourquoi y aller : la place du marché médiévale (Rynek Główny) parmi les plus belles d'Europe ; château du Wawel, quartier juif de Kazimierz très vivant ; ville universitaire, énergie jeune et tarifs très accessibles."**
>
> "Kraków (Poland) — why go: the medieval market square (Rynek Główny), among the most beautiful in Europe; Wawel Castle, the very lively Jewish quarter of Kazimierz; a university city, youthful energy and very affordable prices."
>
> — Claude, France, question about a short city break in Eastern Europe offering the best mix of culture, history and atmosphere (Log ID 403883f8-91b1-4c4c-acf0-035b48af4093). Region question; the answer marks Kraków as the first choice.

8.5 Germany

Market card (individuals: 60 questions, 360 answers)

- Questions that name Poland: 47 → Poland in 100% of answers
- Questions that do not name Poland: 13 (region 12 · topic 0 · general 1 · no hint 0) → Poland in 82.1%
- Region questions: Poland 87.5% · Czechia 86.1 · Germany 50.0 · Austria 45.8
- General questions about Europe: too few questions (1)
- Poland explicitly recommended: in the region 52.4% (63 mentions) · in Europe generally: too few questions
- Polish sources in citations: individuals 25.2% (1,800 citations, 47% of them from Perplexity search) · companies 35.0% (1,329; 63%)
- Most frequent domains, individuals' questions: europa.eu 46 · gov.pl 44 · auswaertiges-amt.de 29
- Poland on request (open portrait, 12 answers per question): well-known Poles: Chopin, Wałęsa, Skłodowska-Curie, Copernicus, John Paul II, Lewandowski 12 of 12 each; first associations: Warsaw 12, Kraków 9, the Baltic 8, pierogi 8, Solidarity 7, Germany 5
- "Poland. More than you expected" campaign market: yes

What the questions are about. Germany is a next-door market that makes comparisons: 47 of the 60 questions name Poland outright, and four of them set the Polish Baltic coast against the German one. People ask about work in Warsaw and Wrocław, an Erasmus semester, implants in Szczecin or Budapest, a lakeside house in Masuria, retiring to Poland, shopping and filling up from Frankfurt (Oder), buying a holiday home on the Polish Baltic rather than in Mecklenburg. Three questions return to a single theme: how Poles talk about German–Polish history today and what they think of Germans; three others ask about the rule of law and minority rights. One concerns Polish colleagues in a German company. The remaining 13 questions do not name Poland: twelve about the "eastern neighbour" or Eastern Europe (a city weekend, a semester, dentures, national parks, a move for two to three years) and one general question about a region of Europe with lakes and mountains.

Do the models suggest Poland of their own accord? In questions about neighbouring countries Poland appears in 87.5% of answers, level with Czechia (86.1%) and ahead of Austria (45.8%). The second country in these answers is Germany itself (50.0%): in every second answer to questions that do not name Poland, the model suggests the persona's home country, usually as a benchmark for prices and quality. One question shows the limit: a family looking for a "region with lakes and mountains, not too far away" gets Austria, Germany and Italy, and Poland in one answer out of six. There are too few general questions about Europe (one) to say anything about Poland's presence beyond the region.

Who Poland competes with. In the region, a dead heat with Czechia; Austria appears 42 percentage points less often.

Mentioned or recommended. In questions about neighbouring countries Poland is explicitly recommended in 52.4% of the answers that mention it (63 mentions); in the rest it is one of several countries, alongside Czechia and Austria.

Where the answers point. Polish sites account for 25.2% of citations in answers to individuals and 35.0% in answers to companies; 47% and 63% of citations respectively come from Perplexity's search. The most frequently cited domains are europa.eu (46), gov.pl (44) and the Auswärtiges Amt, the German Federal Foreign Office (29): like British and Spanish questions, German questions about safety and formalities get their own ministry as a source.

Poland on request, or the lakes that have to be named. Asked in German about well-known Poles, the models name Chopin, Wałęsa, Marie Skłodowska-Curie, Copernicus, John Paul II and Lewandowski in all 12 answers. Alongside Warsaw, Kraków, the Baltic and pierogi, the first associations with Poland include two names that occur in no other market: Masuria (4 of 12) and the Oder (4 of 12). In the study's answers to German individuals Masuria appears in 34 of 360 answers, but 24 of these respond to the four questions that themselves name Masuria; beyond those it appears in 10 of 336, and in questions that do not name Poland, in one of 78. The Oder appears in 7 of 360, six of them answers to the question from Frankfurt (Oder). Anyone who writes "Masuren" gets a lakeside holiday plan; anyone who asks about lakes not far away does not get Masuria. The doors are there, but only the question itself opens them.

The campaign. Germany is one of the five markets of the "Poland. More than you expected" campaign. The faces from the canon (Chopin, Wałęsa, Lewandowski) appear in every German answer about well-known Poles, and Wałęsa also among first associations (4 of 12); the campaign's new faces appear in none of the 120 open-portrait answers in this market.

> **"Polen – Der Spitzenreiter bei Infrastruktur und Internationalität. Polen hat sich zu einem der beliebtesten Ziele für deutsche Studierende in Osteuropa entwickelt. Die Universitäten sind modern und stark internationalisiert."**
>
> "Poland — the front-runner for infrastructure and internationalisation. Poland has become one of the most popular destinations for German students in Eastern Europe. The universities are modern and highly internationalised."
>
> — Gemini, Germany, a parent's question about a daughter's semester abroad in Eastern Europe (the quality of study and recognition in Germany) (Log ID 2fb8f899-6671-4c28-a901-9c820057845e). Region question; Poland named.

8.6 Italy

Market card (individuals: 58 questions, 348 answers)

- Questions that name Poland: 50 → Poland in 100% of answers
- Questions that do not name Poland: 8 (region 6 · topic 1 · general 1 · no hint 0) → Poland in 89.6%
- Region questions: Poland 97.2% · Czechia 86.1 · Hungary 86.1 · Romania 55.6
- General questions about Europe: too few questions (1)
- Poland explicitly recommended: in the region 57.1% (35 mentions) · in Europe generally: too few questions
- Polish sources in citations: individuals 36.1% (1,518 citations, 59% of them from Perplexity search) · companies 33.3% (1,184; 61%)
- Most frequent domains, individuals' questions: numbeo.com 32 · gov.pl 22 · ec.europa.eu 19
- Poland on request (open portrait, 12 answers per question): well-known Poles: Chopin, Wałęsa, Skłodowska-Curie, Copernicus, John Paul II, Solidarity 12 of 12 each, Lewandowski 10; first associations: Kraków 12, Solidarity 11, pierogi 11, Auschwitz-Birkenau 10, John Paul II 10, Warsaw 10
- "Poland. More than you expected" campaign market: yes

What the questions are about. Of the 58 Italian questions, 50 name Poland outright, more than in any other market in the study. From Italy come questions about the Erasmus of a son, daughter, sister, brother or nephew (ten questions), Kraków, Gdańsk and Warsaw for a city break (eight), what a week or a weekend really costs (six), IT and engineering jobs in Warsaw and Wrocław (six), national parks

and mountains (five), communism and the transition, including for a term paper (three), Auschwitz with teenage children, parents' fears about the war, and what it is like to live in a "very Catholic and conservative country". One question concerns a new Polish colleague on the team, one a cousin who works in Poland. Eight questions do not name Poland: six about Eastern or Central Europe (post-war history, a semester, a programmer's job, a weekend "beyond Prague and Budapest"), one about dental treatment abroad and one general question about a country for a few years with low costs and an Italian community.

Do the models suggest Poland of their own accord? In questions about Central and Eastern Europe Poland appears in 97.2% of answers, 11 percentage points more often than Czechia and Hungary (86.1% each), and ahead of Romania. In the question about dental treatment abroad it appears in four of six answers. There are too few general questions about Europe (one) to say anything about Poland's presence beyond the region. The home country appears in 8.3% of answers.

Who Poland competes with. In the region Poland appears 11 percentage points more often than Czechia and Hungary. This is one of the larger regional leads in the study, though it rests on six questions.

Mentioned or recommended. In questions about the region Poland is explicitly recommended in 57.1% of the answers that mention it (35 mentions), usually as Kraków for a weekend, cheaper than Prague.

Where the answers point. Polish sites account for 36.1% of citations in answers to individuals and 33.3% in answers to companies; 59% and 61% of citations respectively come from Perplexity's search. The most frequently cited domain is not a Polish site but the cost-of-living comparison site numbeo.com (32), ahead of gov.pl (22) and ec.europa.eu (19): Italian questions ask about prices and get a price list.

Poland on request, or the canon with John Paul II. Asked in Italian about well-known Poles, the models name Chopin, Wałęsa, Marie Skłodowska-Curie, Copernicus, John Paul II and Solidarity in all 12 answers, and Lewandowski in 10. The first associations with Poland are Kraków (12 of 12), Solidarity and pierogi (11 each), and Auschwitz-Birkenau, John Paul II and Warsaw (10 each). John Paul II is so obvious an association here that one of the study's Italian questions rules him out itself: "apart from vodka, pierogi and John Paul II, what is Poland known for today?" The two countries' shared history barely registers among the associations: Queen Bona Sforza appears in 2 of 12 answers. The Italian portrait of Poland has no thread of its own: it is the canon from chapter 6.

The campaign. Italy is one of the five markets of the "Poland. More than you expected" campaign. The faces from the canon appear in the answers about well-known Poles (Chopin and Wałęsa in all 12, Lewandowski in 10), and Chopin (7 of 12) and Wałęsa (5 of 12) also among first associations; the campaign's new faces appear in none of the 120 open-portrait answers in this market.

Errors. Two of the 14 confirmed factual errors concern questions from this market: a Claude answer placed the Museum of the Second World War in Warsaw instead of Gdańsk, and a DeepSeek answer to a company's question listed Trieste among Poland's logistics hubs (chapter 7).

> **"Cracovia, Polonia. Perché: Città medievale magnifica, storia profonda (incluso il quartiere ebraico Kazimierz e Auschwitz a poca distanza), e costi tra i più bassi d'Europa."**
>
> "Kraków, Poland. Why: a magnificent medieval city, a deep history (including the Jewish quarter of Kazimierz, with Auschwitz a short distance away) and some of the lowest costs in Europe."
>
> — DeepSeek, Italy, question about an inexpensive long weekend in Central Europe beyond Prague and Budapest (Log ID 0dcdc15d-af31-44aa-984c-dfbdbec60e92). Region question; Kraków among the alternatives.

8.7 Norway

Market card (individuals: 58 questions, 348 answers)

- Questions that name Poland: 44 → Poland in 100% of answers
- Questions that do not name Poland: 14 (region 6 · topic 1 · general 7 · no hint 0) → Poland in 70.2%
- Region questions: Poland 94.4% · Czechia 86.1 · Hungary 86.1 · Slovakia 66.7
- General questions about Europe: Poland 45.2% · Germany 42.9 · Spain 42.9 · Czechia 38.1
- Poland explicitly recommended: in the region 55.9% (34 mentions) · in Europe generally 84.2% (19 mentions, tendency)
- Polish sources in citations: individuals 36.2% (1,596 citations, 52% of them from Perplexity search) · companies 38.4% (1,253; 61%)
- Most frequent domains, individuals' questions: gov.pl 31 · poland.travel 24 · regjeringen.no 23
- Poland on request (open portrait, 12 answers per question): well-known Poles: Chopin, Wałęsa, Skłodowska-Curie, Copernicus, John Paul II, Lewandowski 12 of 12 each; first associations: Kraków 9, Solidarity 9, Warsaw 9, Auschwitz-Birkenau 8, Gdańsk 6, Wałęsa 6
- "Poland. More than you expected" campaign market: no

What the questions are about. Questions from Norway about Poland are most often about choosing a city: eight of the 44 questions that name Poland are "Gdańsk, Kraków or Warsaw" for three to four days, and five more ask about regions beyond the big cities (mountains, parks, the coast) for a week. Then come studies (recognition, insurance for a Norwegian student), a job offer and taxes, implants in Szczecin or Budapest, a flat on the Polish Baltic rather than in Spain, and parents' fears about the war. Four questions concern Poles in the persona's life: a daughter's Polish boyfriend and Christmas with his family, Polish tradespeople on a building site, a mother from Poland, relatives in Poland. The remaining 14 questions do not name Poland: seven general ones (moving to another European country, a cheap city weekend, culture without the tourist traps, a daughter's studies), six about Eastern or Central Europe and one about dental treatment abroad.

Do the models suggest Poland of their own accord? In questions about Eastern Europe Poland appears in 94.4% of answers, ahead of Czechia and Hungary (86.1% each). In general questions it appears in 45.2%, level with Germany and Spain (42.9% each) and ahead of Czechia. One general answer in five (21.4%) names no country at all: these are the answers to two questions about a residence permit "in another European country", which get a description of procedures instead of a list of countries; Poland does not appear in them once. The home country rarely appears (2.4%): anyone in Norway asking about going abroad is not offered Norway.

Who Poland competes with. In the region Poland appears 8 percentage points more often than Czechia and Hungary. In Europe generally, a dead heat with Germany and Spain.

Mentioned or recommended. In questions about the region Poland is explicitly recommended in 55.9% of the answers that mention it (34 mentions); in general questions, in 84.2% (19 mentions, tendency). The direction matches chapter 4: Poland is mentioned less often and recommended more often, usually in the form of Kraków for a weekend, with a direct flight from Oslo.

Where the answers point. Polish sites account for 36.2% of citations in answers to individuals and 38.4% in answers to companies; 52% and 61% of citations respectively come from Perplexity's search. After gov.pl (31), the most frequently cited are poland.travel (24) and regjeringen.no (23). Norway is the only one of the twelve markets where Poland's official tourism portal is among the three most

frequently cited domains, and it is also a market where the answers point to the questioner's own government as a source.

Poland on request, or the canon without a thread of its own. Asked in Norwegian about well-known Poles, the models name Chopin, Wałęsa, Marie Skłodowska-Curie, Copernicus, John Paul II and Lewandowski in all 12 answers. The first associations with Poland are Kraków, Solidarity and Warsaw (9 of 12 each), Auschwitz-Birkenau (8), and Gdańsk and Wałęsa (6 each). No name comes up here clearly more often than in other markets: the Norwegian portrait of Poland is the canon from chapter 6, with no thread of its own. Gdańsk, one of the three cities to choose between in the study's Norwegian questions, appears in half of the association answers.

Errors. Two of the 14 confirmed factual errors concern questions from this market: a DeepSeek answer to a company's question stated that Poland is not in the Schengen area, and a Perplexity answer put the name of a country that does not exist on a list of embassies (chapter 7).

> **"Krakow, Polen. Hvorfor: Direktefly fra Oslo (Wizz Air, Ryanair) tilhører de rimeligste rutene fra Norge – ca. 2 t 20 min. Krakow er vakker, billig og har et rikt kulturliv."**
>
> "Kraków, Poland. Why: direct flights from Oslo (Wizz Air, Ryanair) are among the cheapest routes from Norway — approx. 2 h 20 min. Kraków is beautiful, inexpensive and has a rich cultural life."
>
> — DeepSeek, Norway, question about European cities for a short, inexpensive weekend away from Norway (Log ID efdcd363-f838-4355-9e1a-5998e9aa33e2). General question; Kraków alongside Copenhagen and Riga.

8.8 Slovakia

Market card (individuals: 58 questions, 348 answers)

- Questions that name Poland: 44 → Poland in 100% of answers
- Questions that do not name Poland: 14 (region 14 · topic 0 · general 0 · no hint 0) → Poland in 88.1%
- Region questions: Poland 88.1% · Hungary 83.3 · Czechia 81.0 · Austria 69.0
- General questions about Europe: no questions
- Poland explicitly recommended: in the region 50.0% (74 mentions) · in Europe generally: no questions
- Polish sources in citations: individuals 40.7% (1,715 citations, 46% of them from Perplexity search) · companies 50.3% (1,681; 59%)
- Most frequent domains, individuals' questions: gov.pl 27 · study.gov.pl 22 · muzeumkrakowa.pl 17
- Poland on request (open portrait, 12 answers per question): well-known Poles: Chopin, Wałęsa, Copernicus, John Paul II 12 of 12 each, Lewandowski 11, Solidarity 10; first associations: Kraków 12, the Tatras 12, pierogi 9, Warsaw 8, Zakopane 7, the Baltic 5
- "Poland. More than you expected" campaign market: no

What the questions are about. The study's Slovak questions are not about Europe, only about the nearby abroad. All 14 questions that do not name Poland are about neighbouring countries, the V4 countries or "blízkeho zahraničia", the nearby abroad: cross-border shopping from Žilina and eastern Slovakia (four questions), studying as an alternative to Czechia and Austria, working in a neighbouring country close to family, the culture and history of the neighbours, a mountain holiday with some culture "not too far away". The other 44 questions name Poland outright: shopping trips from Orava to Nowy Targ and Żywiec, a weekend in Kraków (four questions), the Tatras (five questions: from the Tatras to Silesia, the Tatras and the north, the Tatras and the sea), IT and medical degrees compared with Czechia and Austria (six), IT work in Gdańsk, implants, an aquapark and thermal baths near the

border, living in Zakopane rather than in the High Tatras. Two questions, as in Czechia, worry about "hidden parking charges" in Kraków; one asks outright how Poles feel about Slovaks and about the Tatras; one, what it is like to be the only person from Slovakia in a Polish team.

Do the models suggest Poland of their own accord? In questions about neighbouring countries Poland appears in 88.1% of answers, 5 percentage points more often than Hungary (83.3%) and 7 more than Czechia (81.0%), and ahead of Austria (69.0%). The home country appears in 16.7% of answers. This market has no topical, general or no-hint questions, so the portrait of Slovakia stops at the region.

Who Poland competes with. With Hungary and Czechia: all three countries appear in more than four answers out of five, Poland most often; Austria appears clearly less often.

Mentioned or recommended. In questions about neighbouring countries Poland is explicitly recommended in 50.0% of the answers that mention it (74 mentions); in the rest it is one country on a list, alongside Czechia and Hungary.

Where the answers point. Polish sites account for 40.7% of citations in answers to individuals and 50.3% in answers to companies; 46% and 59% of citations respectively come from Perplexity's search. The most frequently cited domains are gov.pl (27), study.gov.pl (22) and muzeumkrakowa.pl (17): the portal on studying in Poland and the Kraków city museum answer the two commonest Slovak topics, studying and a weekend in Kraków.

Poland on request, or mountains in, beer out. Asked in Slovak for their first associations with Poland, the models give Kraków and the Tatras in all 12 answers, Zakopane in 7 of 12 (both more often than in any other market) and, in 4 of 12 answers, two beer brands, Żywiec and Tyskie, which feature among no other market's associations. The mountains make it into practical answers widely: the Tatras appear in 86 of the 348 answers to Slovak individuals (nine questions name them; outside those, 32 of 294), and Zakopane in 84 of 348 (only one question names it; outside that one, 78 of 342). The beer does not make it at all: Żywiec appears in 9 of 348 answers, each time as a town people drive to from Orava for shopping and never as a beer; Tyskie appears in one, as the brewery museum in Tychy. As in Czechia, the shared mountains have doors, because they are a place to see; the shared beer has none.

> **"Najvýhodnejšia destinácia z východného Slovenska je momentálne Poľsko. Prečo práve Poľsko – Ceny potravín: Poľsko má dlhodobo nižšie ceny základných potravín (mäso, mliečne výrobky, vajcia, oleje, alkohol) ako Slovensko."**
>
> "The best-value destination from eastern Slovakia right now is Poland. Why Poland in particular — food prices: Poland has long had lower prices for basic foods (meat, dairy products, eggs, oils, alcohol) than Slovakia."
>
> — Grok, Slovakia, question about the best-value destinations for one-day shopping trips from eastern Slovakia (Log ID 57e93302-5ecc-464e-bdb2-0ca69c0099c7). Question about neighbouring countries; Poland as the first and main choice.

8.9 Spain

Market card (individuals: 58 questions, 348 answers)

- Questions that name Poland: 40 → Poland in 100% of answers
- Questions that do not name Poland: 18 (region 11 · topic 1 · general 6 · no hint 0) → Poland in 67.6%
- Region questions: Poland 81.8% · Czechia 92.4 · Hungary 75.8 · Austria 69.7
- General questions about Europe: Poland 36.1% · the Netherlands 58.3 · Portugal 55.6 · Spain 50.0
- Poland explicitly recommended: in the region 50.0% (54 mentions) · in Europe generally 76.9% (13 mentions, tendency)
- Polish sources in citations: individuals 25.1% (1,390 citations, 53% of them from Perplexity search) · companies 37.7% (1,100; 60%)
- Most frequent domains, individuals' questions: numbeo.com 21 · exteriores.gob.es 18 · europa.eu 18
- Poland on request (open portrait, 12 answers per question): well-known Poles: Chopin, Wałęsa, Skłodowska-Curie, Copernicus, John Paul II, Solidarity 12 of 12 each, Lewandowski 11; first associations: Warsaw 12, Auschwitz-Birkenau 10, Kraków 10, Wałęsa 10, John Paul II 10, Lewandowski 10
- "Poland. More than you expected" campaign market: yes

What the questions are about. Of the 58 Spanish questions, 40 name Poland outright: a honeymoon in Poland (two questions), a daughter's or son's Erasmus, a job offer in Wrocław and the climate there, the cost of a weekend in Kraków with a flight from Madrid, Auschwitz with teenage children, parents' fears about the war in Ukraine. The Spanish questions are often written in "act as an adviser" mode and compare Poland with Italy, Croatia, Czechia, Germany or Hungary. One question concerns a Polish colleague at work. The remaining 18 questions do not name Poland: eleven about Central and Eastern Europe (a city weekend in winter, where to live cheaply in the East, historical sensitivities), six general ones (two years in a cheaper country with a Spanish community, studying in English, moving to another EU country) and one about dental treatment.

Do the models suggest Poland of their own accord? In questions about Central Europe Poland appears in 81.8% of answers; in general questions, in 36.1%, with the Netherlands (58.3%), Portugal (55.6%) and Spain itself (50.0%) in the lead. Someone in Spain asking "where to live more cheaply than in Spain" also gets their own country in every second answer.

Who Poland competes with, or the regional exception. Spain is the only market where Poland does not come first in questions about the region: Czechia appears in 92.4% of answers, 11 percentage points more often than Poland (81.8%). For Spaniards, "Central Europe for a weekend" means above all Prague, with Kraków and Wrocław behind it. In Europe generally Poland appears 22 percentage points less often than the Netherlands.

Mentioned or recommended. In questions about the region Poland is explicitly recommended in 50.0% of the answers that mention it (54 mentions). In general questions the figure is 76.9% (13 mentions, tendency): when a model picks Poland itself for someone in Spain looking for a cheaper country to live in, it usually recommends it, often as Central Europe's "surprise".

Where the answers point. Polish sites account for 25.1% of citations in answers to individuals and 37.7% in answers to companies; 53% and 60% of citations respectively come from Perplexity's search. The most frequently cited domain in individuals' questions is not a Polish site but the cost-of-living comparison site numbeo.com (21), followed by the Spanish Ministry of Foreign Affairs at

exteriores.gob.es (18): safety questions ask outright “what Exteriores says”, and the answer points to a home-state source, as in the United Kingdom.

Poland on request, or the Barcelona footballer. Asked in Spanish about well-known Poles, the models name Chopin, Wałęsa, Marie Skłodowska-Curie, Copernicus, John Paul II and Solidarity in all 12 answers, and Lewandowski in 11. Among Spanish first associations with Poland, people stand alongside Warsaw, Auschwitz and Kraków: Wałęsa, John Paul II and Lewandowski, each in 10 of 12 answers; Lewandowski appears among first associations here 3.6 times as often as in other markets, and FC Barcelona comes up in 5 of 12 answers even to the question about well-known Poles. In the study’s answers to Spanish individuals Lewandowski appears in 1 of 348 and FC Barcelona in 3 of 348; there is no question about sport among the Spanish questions, so this is a lack of opportunity, not a gap. The most Spanish of Poland’s threads is in the portrait, and it has no door into a question about a weekend in Kraków.

The campaign. Spain is one of the five markets of the “Poland. More than you expected” campaign. The faces from the canon are present in Spanish answers: Chopin and Wałęsa in every answer about well-known Poles, Wałęsa and Lewandowski among first associations (10 of 12 each); Maryla Rodowicz appears in two answers about culture (ChatGPT), while Eisenberg, Rubik, Rozbicki and JIMEK, the new faces, appear in none of the 120 Spanish open-portrait answers. Spain is the only one of the five campaign markets where Lewandowski is among the first associations with Poland.

Errors. Two of the 14 confirmed factual errors concern companies’ questions from this market: a Claude answer warned of “customs checks within the EU” between Spain and Poland, and a Gemini answer located a plant in the “Grand Duchy of Poland” (chapter 7).

> **“los países que más sentido tienen como punto de partida son Polonia, Rumanía, Bulgaria, Hungría y, fuera de la UE, Montenegro o Serbia. Los más equilibrados para tu caso suelen ser Polonia y Rumanía, porque combinan precios relativamente bajos, ciudades con vida cultural y opciones donde es más fácil encontrar otros expatriados, incluida comunidad española.”**
>
> “the countries that make most sense as a starting point are Poland, Romania, Bulgaria, Hungary and, outside the EU, Montenegro or Serbia. The most balanced for your situation are usually Poland and Romania, because they combine relatively low prices, cities with a cultural life and options where it is easier to find other expatriates, including a Spanish community.”
>
> — Perplexity, Spain, question about a European country for two years, cheaper than Spain, open to foreigners and with a Spanish community (Log ID 115a457c-8d8a-4c7f-9d6c-c551ced22fb0). General question; Poland named, although the question does not name it.

8.10 Sweden

Market card (individuals: 58 questions, 348 answers)

- Questions that name Poland: 40 → Poland in 100% of answers
- Questions that do not name Poland: 18 (region 2 · topic 3 · general 5 · no hint 8) → Poland in 48.1%
- Region questions: Poland 100% · Hungary 100 · Czechia 100 · Slovakia 75.0 (2 questions)
- General questions about Europe: Poland 60.0% · Germany 56.7 · Denmark 50.0 · the Netherlands 50.0
- Poland explicitly recommended: in the region 75.0% (12 mentions, tendency) · in Europe generally 55.6% (18 mentions, tendency)
- Polish sources in citations: individuals 36.4% (1,703 citations, 48% of them from Perplexity search) · companies 36.0% (977; 65%)
- Most frequent domains, individuals' questions: numbeo.com 28 · gov.pl 27 · auschwitz.org 20
- Poland on request (open portrait, 12 answers per question): well-known Poles: Chopin, Wałęsa, Skłodowska-Curie, Copernicus, John Paul II, Lewandowski 12 of 12 each; first associations: Solidarity 12, Kraków 10, pierogi 10, Warsaw 9, the Baltic 7, Gdańsk 7
- "Poland. More than you expected" campaign market: no

What the questions are about. Of the 58 Swedish questions, 40 name Poland outright: a semester in Kraków and a student's costs (six questions), "Kraków, Gdańsk or Warsaw" for a weekend (five), work in Gdańsk and the language, Second World War museums and Jewish heritage (five), national parks, the coast or Masuria for a week (six), a flat on the Polish Baltic rather than in Spain, parents' fears about the war. Many Swedish questions concern Poland in family life: four are about family in Poland (family ties and visits to relatives), two about tracing an ancestral Polish village whose name has changed, and one about a team half made up of Poles. The remaining 18 questions do not name Poland, and they include eight with no hint of going abroad, more than in any other market: "a longer trip with nature", "the region that is the heart of the country", "what makes this country's culture special", "the ancestral village that is not on the maps". Five are general, three topical (dental treatment, Second World War museums) and two about Central Europe.

Do the models suggest Poland of their own accord? When a question from Sweden says nothing about going away, the model answers with Sweden (41.7%) and Norway (39.6%); Poland appears in 8.3% of such answers (4 of 48). When the question is about Europe in general, Poland appears in 60.0%, level with Germany (56.7%) and ahead of Denmark and the Netherlands. When it is about Central Europe, Poland always appears, as do Hungary and Czechia (two questions). In the topical questions, on Second World War museums and dental treatment abroad, Poland appears in every answer.

Who Poland competes with. In Europe generally, a dead heat with Germany; in questions with no hint of going abroad, the competitor is not another country but Sweden itself.

Mentioned or recommended. In general questions Poland is explicitly recommended in 55.6% of the answers that mention it (18 mentions, tendency), and in questions about the region in 75.0% (12 mentions, tendency). Both figures sit at the lower limit at which we report the recommendation rate, so we do not interpret the direction.

Where the answers point. Polish sites account for 36.4% of citations in answers to individuals and 36.0% in answers to companies; 48% and 65% of citations respectively come from Perplexity's search. The most frequently cited domain is not a Polish site but the cost-of-living comparison site

numbeo.com (28), ahead of gov.pl (27) and auschwitz.org (20): Swedish questions about the cost of studying and about places of remembrance get the sources they ask for.

Poland on request, or the ferry no one takes. Asked in Swedish about well-known Poles, the models name Chopin, Wałęsa, Marie Skłodowska-Curie, Copernicus, John Paul II and Lewandowski in all 12 answers; Solidarity appears in all 12 answers on first associations. Two names come up here more often than elsewhere: Gdańsk (7 of 12) and Ystad (4 of 12). Ystad, the port the ferry to Świnoujście sails from, features in no other market's associations. Gdańsk makes it into practical answers widely: 129 of the 348 answers to Swedish individuals, only 30 of them in reply to the five questions that name it. The ferry does not: Ystad appears in 3 of 348 answers, in one as a ferry port and in two as a Swedish town; the ferry in general appears in 15 of 348, usually as an alternative to flying to Gdańsk, for those taking their own car. The route by which people from Sweden actually reach Poland is in the models' portrait, but rarely in their answers about a trip.

> **"Polen är förmodligen det bästa valet för en skandinav – relativt nära, EU-land, säkert, rikt på upplevelser och betydligt billigare än Sverige/Norge/Danmark."**
>
> "Poland is probably the best choice for a Scandinavian — relatively close, an EU country, safe, full of experiences and considerably cheaper than Sweden/Norway/Denmark."
>
> — Claude, Sweden, question about weekend cities "in a country to the east that gives a lot for the money" (Log ID 573a7000-102c-4e38-989d-4dcf96a9293f). General question; Poland the only country in the answer, tone 9. A counterpoint to chapter 7.

8.11 Ukraine

Market card (individuals: 56 questions, 336 answers)

- Questions that name Poland: 40 → Poland in 100% of answers
- Questions that do not name Poland: 16 (region 6 · topic 0 · general 9 · no hint 1) → Poland in 83.3%
- Region questions: Poland 88.9% · Romania 86.1 · Slovakia 80.6 · Hungary 72.2
- General questions about Europe: Poland 77.8% · Germany 79.6 · Czechia 66.7 · the Netherlands 44.4
- Poland explicitly recommended: in the region 37.5% (32 mentions) · in Europe generally 57.1% (42 mentions)
- Polish sources in citations: individuals 39.4% (1,481 citations, 54% of them from Perplexity search) · companies 46.4% (1,411; 56%)
- Most frequent domains, individuals' questions: gov.pl 66 · dpsu.gov.ua 27 · home-affairs.ec.europa.eu 22
- Poland on request (open portrait, 12 answers per question): well-known Poles: Chopin, Wałęsa, Skłodowska-Curie, Copernicus, John Paul II, Lewandowski 12 of 12 each; first associations: Warsaw 11, Kraków 8, the European Union 3, Krakowiec 3, Rzeszów 3, Szehyni 3
- "Poland. More than you expected" campaign market: no

What the questions are about. In questions from Ukraine Poland is not a foreign country to be discovered but a country where the persona already lives or has family. Of the 56 questions, 40 name Poland outright, and almost none is about a holiday: legalising residence, PESEL UKR (the Polish identification number for Ukrainians under temporary protection) and what happens after 4 March 2027 (six questions), work compared with Czechia and Germany (six), schooling for children and the recognition of school certificates and diplomas (eight), renting or buying a flat in Wrocław or Kraków (four), border crossings, queues and documents (six), how Poles feel today about Ukrainians and

about Volhynia, a sister in Poland, relatives working there, running one's own business. The remaining 16 questions do not name Poland: nine general ones (studying, working and moving with family "to an EU country where Ukrainians are treated kindly"), six about neighbouring countries (documents for a short trip, Eastern European cities to live in) and one, with no hint of going abroad, about attractive cities to move to.

Do the models suggest Poland of their own accord? In questions about neighbouring countries Poland appears in 88.9% of answers, level with Romania (86.1%) and ahead of Slovakia (80.6%) and Hungary. In general questions about Europe it appears in 77.8%, level with Germany (79.6%) and ahead of Czechia (66.7%). This is the exception described in chapter 2: in no other market does a general question give Poland such a presence, because in questions from Ukraine Poland is the default destination even when the question says only "an EU country". The home country rarely appears (7.3%).

Who Poland competes with. In the region with Romania and Slovakia, in Europe generally with Germany; in both cases the differences are within measurement variability.

Mentioned or recommended. In questions about neighbouring countries Poland is explicitly recommended in 37.5% of the answers that mention it (32 mentions), and in general questions in 57.1% (42 mentions). The direction matches chapter 4, but the level is lower than in most markets, for the reason chapter 4 states outright: answers to questions from Ukraine are more likely to describe the conditions for living and working than to encourage.

Where the answers point. Polish sites account for 39.4% of citations in answers to individuals and 46.4% in answers to companies; 54% and 56% of citations respectively come from Perplexity's search. The most frequently cited domains are gov.pl (66), the Ukrainian border guard's site dpsu.gov.ua (27) and the European Commission's page on temporary protection (22): questions about the border and residence status get home-state sources and sources from the EU.

Poland on request, or the canon cut down to two cities. Asked in Ukrainian about well-known Poles, the models name Chopin, Wałęsa, Marie Skłodowska-Curie, Copernicus, John Paul II and Lewandowski in all 12 answers. First associations with Poland are shorter than elsewhere: Warsaw (11 of 12) and Kraków (8 of 12); no other association exceeds 3 of 12, the threshold from which we count a market-specific thread. The Ukrainian portrait of Poland has no thread of its own: it is the canon from chapter 6, cut down to two cities.

> **"Найкращий баланс для сім'ї зазвичай дають Польща, Чехія та Литва […]. Польща — одна з найзручніших для старту: тут велика українська громада, відносна близькість до України, схожа мова й простіша адаптація […]."**
>
> "For a family, the best balance is usually offered by Poland, Czechia and Lithuania […]. Poland is one of the easiest places to start: there is a large Ukrainian community, it is relatively close to Ukraine, the language is similar and settling in is easier […]."
>
> — Perplexity, Ukraine, question about moving with family to an EU country where Ukrainians are made welcome, work is easy to find and living costs are low (Log ID 9c678cee-e1ae-4fac-9f5b-85b3a18d8426). General question; Poland named.

8.12 United Kingdom

Market card (individuals: 60 questions, 360 answers)

- Questions that name Poland: 43 → Poland in 100% of answers
- Questions that do not name Poland: 17 (region 7 · topic 1 · general 8 · no hint 1) → Poland in 59.8%
- Region questions: Poland 97.6% · Czechia 78.6 · Hungary 73.8 · Romania 73.8
- General questions about Europe: Poland 29.2% · Portugal 70.8 · Spain 45.8 · Germany 43.8
- Poland explicitly recommended: in the region 29.3% (41 mentions) · in Europe generally 85.7% (14 mentions, tendency)
- Polish sources in citations: individuals 26.5% (1,374 citations, 54% of them from Perplexity search) · companies 29.0% (1,260; 64%)
- Most frequent domains, individuals' questions: gov.uk 74 · gov.pl 46 · numbeo.com 33
- Poland on request (open portrait, 12 answers per question): well-known Poles: Chopin, Wałęsa, Skłodowska-Curie, Copernicus, John Paul II, Lewandowski 12 of 12 each; first associations: Kraków 11, Warsaw 11, pierogi 10, Solidarity 8, vodka 6, the Battle of Britain 5
- "Poland. More than you expected" campaign market: yes

What the questions are about. Of the 60 British questions, 43 name Poland outright, and Poland appears in every answer to them. They ask about a master's degree, a weekend in Gdańsk or Wrocław, moving for a few years, dental treatment in Kraków. Several are about Poland within the United Kingdom itself: Polish co-workers in a warehouse, Christmas Eve with a Polish family in Reading, a grandfather's Polish roots. The remaining 17 are the questions of someone still looking for a destination. Eight of them ask about Europe in general ("a cheap autumn getaway", "moving for a few years with children"), seven about Central or Eastern Europe.

Two leagues in their purest form. Someone in the United Kingdom asking about Central Europe almost always gets Poland (97.6%), 19 percentage points more often than Czechia, ahead of Hungary and Romania. Someone asking about cheap Europe without naming a region gets Portugal (70.8%), Spain and Germany; Poland appears in fewer than one answer in three (29.2%). In the British market "cheap Europe" means the south, and Poland becomes the answer only when the question says "Central Europe".

Mentioned or recommended. The same Poland that almost always appears in questions about the region is explicitly recommended in only 29.3% of those answers, usually as one country on a list. In general questions it is the other way round: Poland appears rarely, but when the model picks it of its own accord, it recommends it in 85.7% of answers (14 mentions, tendency). This is the strongest example of the pattern from chapter 4: in the region Poland is mentioned, in Europe it is recommended.

Where the answers point. In answers to individuals the most frequently cited domain is not a Polish site but gov.uk (74 citations), ahead of gov.pl (46) and the cost-of-living comparison site numbeo.com (33). Polish sources make up 26.5% of citations, and more than half of all citations in this market come from Perplexity's search. Answers to questions from the United Kingdom point above all to the questioner's home-state sources.

Poland on request, or the shared war missing from the answers. Asked directly about well-known Poles, all six models name Chopin, Wałęsa, Marie Skłodowska-Curie, Copernicus, John Paul II and Lewandowski in every British answer. Among first associations with Poland a thread appears that features in no other market: the Battle of Britain and the Polish pilots (5 of 12 answers). In the study's

practical answers this thread does not appear at all. It is in none of the 714 answers to questions from British individuals and companies, although 17 British questions are about Polish history and culture. This is the activation gap from chapter 6 at the level of a single market: the shared history is in the models' portrait, but there are no doors into a practical answer.

The campaign. The United Kingdom is one of the five markets of the "Poland. More than you expected" campaign. Lewandowski, Wałęsa and Chopin, the faces from the canon, appear in every British answer about well-known Poles; the campaign's new faces appear in none. The Battle of Britain, which the campaign does not use, comes up in British associations with Poland more often than any of the new faces.

One error. One of the 14 confirmed factual errors concerns a question from this market: a DeepSeek answer stated that Poland does not allow naturalised citizens to hold dual citizenship (chapter 7).

> **"Portugal consistently ranks at the top of Expat Insider surveys for 'ease of settling in' and friendliness."**
>
> — Gemini, United Kingdom, question about moving with young children "somewhere in Europe" for a few years (Log ID 1fab2b6e-52f5-40b8-b8a1-7cd58cbec02b). General question; Poland does not appear in the answer.

8.13 United States

Market card (individuals: 60 questions, 360 answers)

- Questions that name Poland: 48 → Poland in 100% of answers
- Questions that do not name Poland: 12 (region 10 · topic 1 · general 1 · no hint 0) → Poland in 95.8%
- Region questions: Poland 98.3% · Czechia 98.3 · Hungary 96.7 · Austria 90.0
- General questions about Europe: too few questions (1)
- Poland explicitly recommended: in the region 32.2% (59 mentions) · in Europe generally: too few questions
- Polish sources in citations: individuals 31.3% (1,442 citations, 53% of them from Perplexity search) · companies 26.5% (1,338; 66%)
- Most frequent domains, individuals' questions: gov.pl 46 · numbeo.com 28 · travel.state.gov 26
- Poland on request (open portrait, 12 answers per question): well-known Poles: Chopin, Wałęsa, Skłodowska-Curie, Copernicus, John Paul II, Lewandowski 12 of 12 each; first associations: pierogi 11, Auschwitz-Birkenau 10, John Paul II 10, Solidarity 10, kielbasa 9, Warsaw 9
- "Poland. More than you expected" campaign market: yes

What the questions are about. Of the 60 American questions, 48 name Poland outright, and eleven of them are about roots: great-grandparents' village, archives, a genealogy trip, citizenship through a grandmother so as to hold an EU passport. Then come Poland as a cheaper alternative to Western Europe, Italy or Germany (eight questions), the cost of living, bureaucracy and taxes when moving (nine), a semester, a visa and the 90-day Schengen limit (five), Fulbright and a research placement (four), Auschwitz with adult children, the border with Ukraine and Belarus and State Department advisories (two), Kraków or Lisbon for a year of remote work, and a hip replacement in Poland, Mexico or Turkey. One question concerns a Polish team leader who has moved to Chicago. The remaining 12 questions do not name Poland: ten about Central or Eastern Europe (work, safety, value for money, history, studies, a grandmother's "old country"), one about Second World War memorial sites and one general question about a year of remote work in Europe.

Do the models suggest Poland of their own accord? In questions about Central Europe Poland appears in 98.3% of answers, exactly as often as Czechia, and ahead of Hungary (96.7%) and Austria (90.0%). Someone in the United States asking about the region gets all of it: on average, almost seven other countries per answer. In the single topical question, on Second World War memorial sites, Poland appears in every answer. There are too few general questions about Europe (one) to say anything about Poland's presence beyond the region. The home country hardly appears at all (2.8%).

Who Poland competes with. In the region, with no one and with everyone: a dead heat with Czechia, with Hungary and Austria close behind.

Mentioned or recommended. The same Poland that almost always appears in questions about the region is explicitly recommended in 32.2% of the answers that mention it (59 mentions), usually as one country on a list, alongside Prague and Budapest. This is the same picture as in the United Kingdom (chapter 4): in the region Poland is mentioned, not recommended.

Where the answers point. Polish sites account for 31.3% of citations in answers to individuals and 26.5% in answers to companies; 53% and 66% of citations respectively come from Perplexity's search. The most frequently cited domains are gov.pl (46), the cost-of-living comparison site numbeo.com (28) and travel.state.gov (26): American questions about safety ask for State Department advisories, and get them.

Poland on request, or kielbasa and Chicago. Asked in English, in an American context, about well-known Poles, the models name Chopin, Wałęsa, Marie Skłodowska-Curie, Copernicus, John Paul II and Lewandowski in all 12 answers. Here first associations with Poland start with food: pierogi (11 of 12), ahead of Auschwitz, John Paul II and Solidarity (10 each), followed by kielbasa (9 of 12) and Chicago (5 of 12); both come up among associations clearly more often than in other markets, and Chicago in no other market at all. In the study's answers to American individuals kielbasa appears in 2 of 360 answers, in lists of dishes. Chicago appears in 50 of 360, but as a departure airport and as the seat of the consulate where visa applications are filed; as a Polish city in only one: the answer about the Polish team leader. The American doors to Poland do exist: the eleven questions about roots. What comes in through them is archives and consulates, not Chicago.

The campaign. The United States is one of the five markets of the "Poland. More than you expected" campaign. The faces from the canon (Chopin, Wałęsa, Lewandowski) appear in every American answer about well-known Poles, and Wałęsa (7 of 12) and Chopin (6 of 12) also among first associations; the campaign's new faces appear in none of the 120 open-portrait answers in this market. Kielbasa and Chicago, which the campaign does not use, come up in American associations more often than any of the new faces.

Errors. Three of the 14 confirmed factual errors concern questions from this market: a DeepSeek answer presented a grant programme that has closed as current, another DeepSeek answer described a scholarship for developing countries as open to Americans, and a Grok answer called Kraków "one of the cheapest capitals in Europe" (chapter 7).

> **"Some Central European countries—especially Poland, Czechia outside Prague's core, Hungary, Slovakia, Slovenia, and parts of Croatia or Romania—can offer better value for a long trip because the main savings occur after you land."**
>
> — ChatGPT, United States, question about what makes Central European countries a better deal for a long holiday when flying from the United States (Log ID 6e3e89f0-3ac8-4dd2-854b-a883fa66658d). Region question; Poland named.

HOW WE COUNT. Market card: questions from individuals in the study and the answers to them (August–September 2026); question type for questions that do not name Poland follows the classification from chapter 2 (region / topic / general / no hint); “Poland in x%” = the share of answers of a given type that mention Poland; countries alongside Poland = the share of answers of a given type in which the country appears; explicitly recommended = a “high” recommendation rating among the answers that mention Poland (chapter 4); Polish sources = the share of Polish domains among citations in the answers for a given market, separately for individuals and companies, combining both modes (sources from Perplexity’s search and URLs given in the answer by the other five models; chapter 5), that is, a profile of the market, not a ranking; Poland on request = the open portrait of 14–15 October 2026, the number of answers containing the name out of 12 answers in the market’s language (chapter 6); market-specific thread = a name that appears among first associations in a given market at least three times as often as the average and in at least 4 of 12 answers, excluding the name of the home country; presence of the thread in the study = a search of the full text of that market’s answers using a list of the name’s forms. The measurement is of the model’s answer returned via the API.

CHAPTER 9

What this means for promoting Poland

IN BRIEF

- Language models (LLMs) do not hold a single opinion of Poland that could be “improved”. They hold answers to questions, and it is the question that decides whether Poland appears at all: it is named in 91.7% of answers to questions about the region, in 45.1% of answers to questions about Europe in general and in 20.0% of answers to questions that give no hint of travel abroad. Promotion that wants to be in the answers has to know which questions it wants to be in.
- Where Poland comes up rarely, it is recommended more often (62.3% in general questions against 47.0% in questions about the region). The model brings Poland in when it has an argument for doing so. Promotion should therefore supply models with arguments, not with presence.
- Poland’s canon (Chopin, Wałęsa, Nicolaus Copernicus, John Paul II, Skłodowska-Curie, Lewandowski) appears in the answers in each of the 12 markets, but it enters practical answers only through “doors”: a place to see or a situation the question is about. The Tatras and Gdańsk have such doors; Skłodowska-Curie, the Battle of Britain and Lewandowski do not. In the language of brand semantics, an entity enters an answer when it is linked to the entity the question is about. A campaign built on association does not build links.
- In all 12 markets, answers point to the authorities of the persona’s home country: from 6% of answers in Slovakia to 21% in the United Kingdom, where gov.uk appears more often than gov.pl. A Polish source that is to reach the answers has to address the situation of a person from the given market, in a language that person reads: gov.pl appears in 11 of the 12 markets chiefly thanks to its English-language pages for foreigners.
- Two gaps in the sources are plain to see. In questions about the cost of living in Poland, answers point to numbeo.com (one of the three most frequent sources in five markets); no Polish institution is named here. The tourism portal poland.travel appears in 4.6% of answers, almost exclusively in English. The market-language versions, hosted on separate domains, barely appear in the answers.
- The image of Poland in the answers changes faster than campaigns do: in seven weeks, newer model versions shifted the share of explicit recommendations by a dozen or so points. This semantic drift has to be measured at regular intervals, with the same set of questions.

9.1 What this chapter can and cannot say

The report measures the answers of six models, collected via their APIs between 25 August and 28 September 2026, to 1,400 questions in 11 languages. It is a map of the representation in a dated study,

and its subject is the brand semantics of Poland: the entities that models link with the name “Poland”, and the links between those entities and the situations people ask about. Chapters 1–8 draw this semantic map separately for each market and question type. The report does not measure what a consumer app shows, website traffic or the effectiveness of any campaign. The conclusions below concern one thing: how Poland looks on this map, and what follows for the institutions that promote it. These are the Polish Tourism Organisation (POT), ministries, diplomatic missions, local governments and brands.

9.2 Poland is in the answer when it is in the question or in the region

The report’s simplest finding is also the most important one for promotion. When a question names Poland, the model mentions it in 99.9% of answers (chapter 1). When it does not, Poland appears in 72.4% of answers, and that share depends on how the question is phrased (chapter 2): 91.7% in questions about a region that includes Poland, 45.1% in questions about Europe in general (38.5% excluding Ukraine) and 20.0% in questions that mention neither a region nor a trip.

For promotion, this changes the objective. The question “do models recommend Poland?” has no single answer; the useful question is “which questions should Poland be in, and which is it missing from?”. In the region, Poland is already there: first in 11 of the 12 markets, ahead of Czechia (80.6%), Hungary (73.2%) and Austria (52.2%). The exception is Spain, where “Central Europe for a weekend” means Prague (Czechia 92.4%, Poland 81.8%; chapter 8). In “default Europe”, Poland barely features: in general questions Germany heads the list (58.0%), while Poland (45.1%) is in the same group as Czechia, Portugal and Spain (chapter 3). For Ukraine, Poland is the default destination even when the question says only “an EU country”. Excluding Ukraine, the figure falls to 38.5%, into the group of countries named in fewer than four answers in ten.

This is where promotion has work to do, and where being present is not enough. Chapter 4 shows a paradox: in questions about the region, Poland is named almost every time but explicitly recommended in only 47.0% of answers; in questions about Europe in general it is named less often, yet when it is named, the model recommends it in 62.3%. The model brings Poland into “default Europe” when it has a reason to, such as a price, a particular city or a situation that Poland suits. When Poland is in the answer only because it lies in the region the question asked about, the model names it and moves on. The lesson for promotion: what counts is the argument, not presence. Content meant to reach the answers must give the model a reason to name Poland as a choice, not as a neighbour.

9.3 Playing the cards the models know

Asked directly about Poland, the models answer with the canon: Chopin, Wałęsa, Copernicus and John Paul II appear in every one of the 144 answers to the question about the best-known Poles, Skłodowska-Curie in 140 and Lewandowski in 138 (chapter 6). In answers to practical questions about a trip, history or culture, the canon barely appears: Chopin in 12.0% of answers, Copernicus in 4.5%, John Paul II in 2.6%, Skłodowska-Curie in 0.8%. This is the activation gap: the models hold these entities in the brand semantics of Poland, but do not bring them up of their own accord.

The Polish Tourism Organisation’s campaign “Poland. More than you expected” (running since August 2026) plays two hands. One is the cards the models know: Lewandowski, Wałęsa and Chopin appear among the associations in the campaign’s five markets (Germany, Italy, Spain, the United Kingdom and the United States). The other is faces that are missing from the answers: Jesse Eisenberg, Anja Rubik, Nikodem Rozbicki and JIMEK appear in none of the portrait’s 1,440 answers, Maryla Rodowicz in four. The measurement was taken eight weeks after the campaign’s launch; it is not an assessment of its effectiveness.

The lesson reaches beyond one campaign. Models do not learn from billboards; what appears in their answers is what is in the texts they were trained on and, for a model with search, also what it finds on the web. Eight weeks in, the new faces are not yet in the answers. The cards the models know are, today, the only capital that promotion can activate straight away. The question is how, and the next section answers it.

9.4 Doors: an association is not an argument

Chapter 6 showed the doors through which the canon enters practical answers: a place to see (Chopin appears alongside Warsaw or Żelazowa Wola in 97% of his mentions, Copernicus alongside Toruń or Frombork in 54%) or a chapter of history the question asks about. Chapter 8 showed the same thing market by market, using associations specific to a single market. Three cases emerge.

Doors wide open. The Tatras are the first thing Czechs and Slovaks associate with Poland, and they enter practical answers in both markets: 48 of 348 Czech answers and 86 of 348 Slovak ones; Zakopane features in 84 of 348 Slovak answers. Gdańsk is what the Swedes associate with Poland, and it appears in 129 of 348 Swedish answers. These associations are places one can travel to, and the models suggest them of their own accord in questions about trips, with no hint in the question.

Doors that only the question opens. The Germans associate Masuria with Poland, but in practical answers it appears in 34 of 360, 24 of them in four questions that themselves mention Masuria; outside those, in 10 of 336. The model knows where Masuria is, but does not suggest it of its own accord.

No doors. The Battle of Britain is the first British association with Poland, yet it appears in none of the 714 British practical answers. Skłodowska-Curie, the French association, appears in 9 of 348 answers, and only in questions about history. Lewandowski, the Spanish association, in 1 of 348. The Ystad–Świnoujście ferry, a Swedish association, in 3 of 348. Kielbasa, the American association, in 2 of 360; Chicago in 50 of 360, but as a departure airport and a consulate, not as a Polish neighbourhood.

What we have called doors since chapter 6 is entity linking: the model brings an entity into a practical answer only when it has linked it to the entity the question asks about (a place, a route, a situation). The Tatras are linked to a trip to the mountains, Gdańsk to a weekend by the sea; Skłodowska-Curie and the Battle of Britain are linked to Poland, but to nothing individuals ask about.

For promotion, this yields a rule: an association that is neither a place nor a decision situation does not make it into the answer to “where should I go?”. A shared scientist, a shared battle and a shared footballer build goodwill, but give the model nothing it could slot into a weekend plan. Mountains and a city by the sea do. Promotion that wants to activate the canon has to add a door to each card, which means links between entities: a place, a route, an event, a reason to come, described in texts available on the web in the markets’ languages. This is not a call to abandon the canon, but to translate it into the language of the questions people ask.

9.5 Answers point to home-country authorities

The sources in an answer show where the model sends people for details. Chapter 5 showed that half of the URLs models give in the answer without search lead to Polish pages (49.5%), against 22.4% of the sources from search. Only Perplexity had live search in the study. Reading the citations market by market adds three observations.

First: in all 12 markets, answers point to home-country authorities. In the British market gov.uk appears (in 74 of 360 answers, more often than gov.pl, at 46); in the German market the Auswärtiges Amt, the German Foreign Office (29 of 360); in the Ukrainian market dpsu.gov.ua (27 of 336); in the

American market travel.state.gov (26 of 360); in the Norwegian market regjeringen.no (23 of 348); in the Spanish market Exteriores, Spain's foreign ministry (18 of 348); and likewise in the other six markets, from foreign ministries to tax offices. In total, between 6% of answers in Slovakia and 21% in the United Kingdom cite an authority of the home country. What these sources have in common is not Poland as a subject but the persona's language and situation: documents, taxes, travel warnings.

Official Polish sites win where they do the same. The gov.pl portal is among the top three cited domains in 11 of the 12 markets, and first in seven. But the most-cited pages are English-language services for foreigners: those of the Office for Foreigners (152 citations of the English version, 128 of the Polish), the Ministry of the Interior and Administration (MSWiA) and the missions in Washington and London. A Polish source appears in an answer when it addresses the situation of a person from the given market, in a language that person reads. Polish institutions should apply this test to their own content: not "is it about Poland?" but "does it address the situation of a person from this market, and is it in their language or in English?".

Second: the cost of living in Poland has no Polish source. The site numbeo.com, which crowdsources prices from the public, appears in 225 of 4,200 answers and is one of the three most frequent sources in five markets (Italy, Spain, Sweden, the United Kingdom and the United States). Statistics Poland (GUS) appears in 76 answers across 12 markets, but makes the top three in none of them. Ask how much it costs to live in Kraków, and the answer points to a site that nobody in Poland edits. This gap can be closed with a single well-built page in the markets' languages.

Third: the tourism portal is cited in English, and rarely. The Polish Tourism Organisation's portal, with all its language versions counted together, appears in 195 of 4,200 answers (4.6%), ranging from none in Ukraine to 9.2% in Norway. It is mainly the English version, poland.travel, that is cited (134 answers). The market-language versions, hosted on separate domains, barely appear in the answers: polen.travel (27 answers: Germany 10, Sweden 8, Norway 6), polska.travel (15), polonia.travel (9), pologne.travel (7), polsko.travel (4). In the French, German and Spanish markets, answers send people to the portal in English even though a version in the market's language exists. Norway, the only market where poland.travel makes the top three domains, owes this not to the Norwegian version (6 answers) but to its questions: Norwegian questions about regions beyond the big cities, nature and driving routes land on the English portal's pages about regions and national parks. The commercial site getyourguide.com appears in 83 answers, fewer than the portal overall but more often than the portal in China, Czechia and Slovakia.

The three observations lead to a single recommendation, the citability rule from chapter 5. Content meant to be cited must be in the market's language or in English, at an address the models link with Poland, and must address a decision situation (documents, costs, a route, safety) rather than describe Poland. How to do this technically (one domain or many, how to structure the language versions, data in a form the models pick up) is a question for the teams that run these sites; the report shows only that, today, the portal's language versions do not reach the answers.

9.6 The war: arguments the models already have

Box 4.7 showed that, when asked directly about safety in Poland during the war in Ukraine, in none of 792 answers do the models advise against Poland as a destination, but they almost always add a "but"; and in questions about safe countries across Europe for a longer stay, between half and three quarters of answers place Poland in the group "with caution, closer to the conflict". What matters for promotion is what the models set against that "but": NATO appears in 78.1% of answers, the Schengen area in 5.3%. The models have the argument of the alliance and repeat it; the argument of a shared legal and economic area they barely use. Communication about Poland's safety does not need to convince the models that Poland is safe. They say as much. It needs to give them arguments that work

at European scale, because at that scale the models rank countries by distance from the front line. In that ranking Poland lands behind Portugal and Spain.

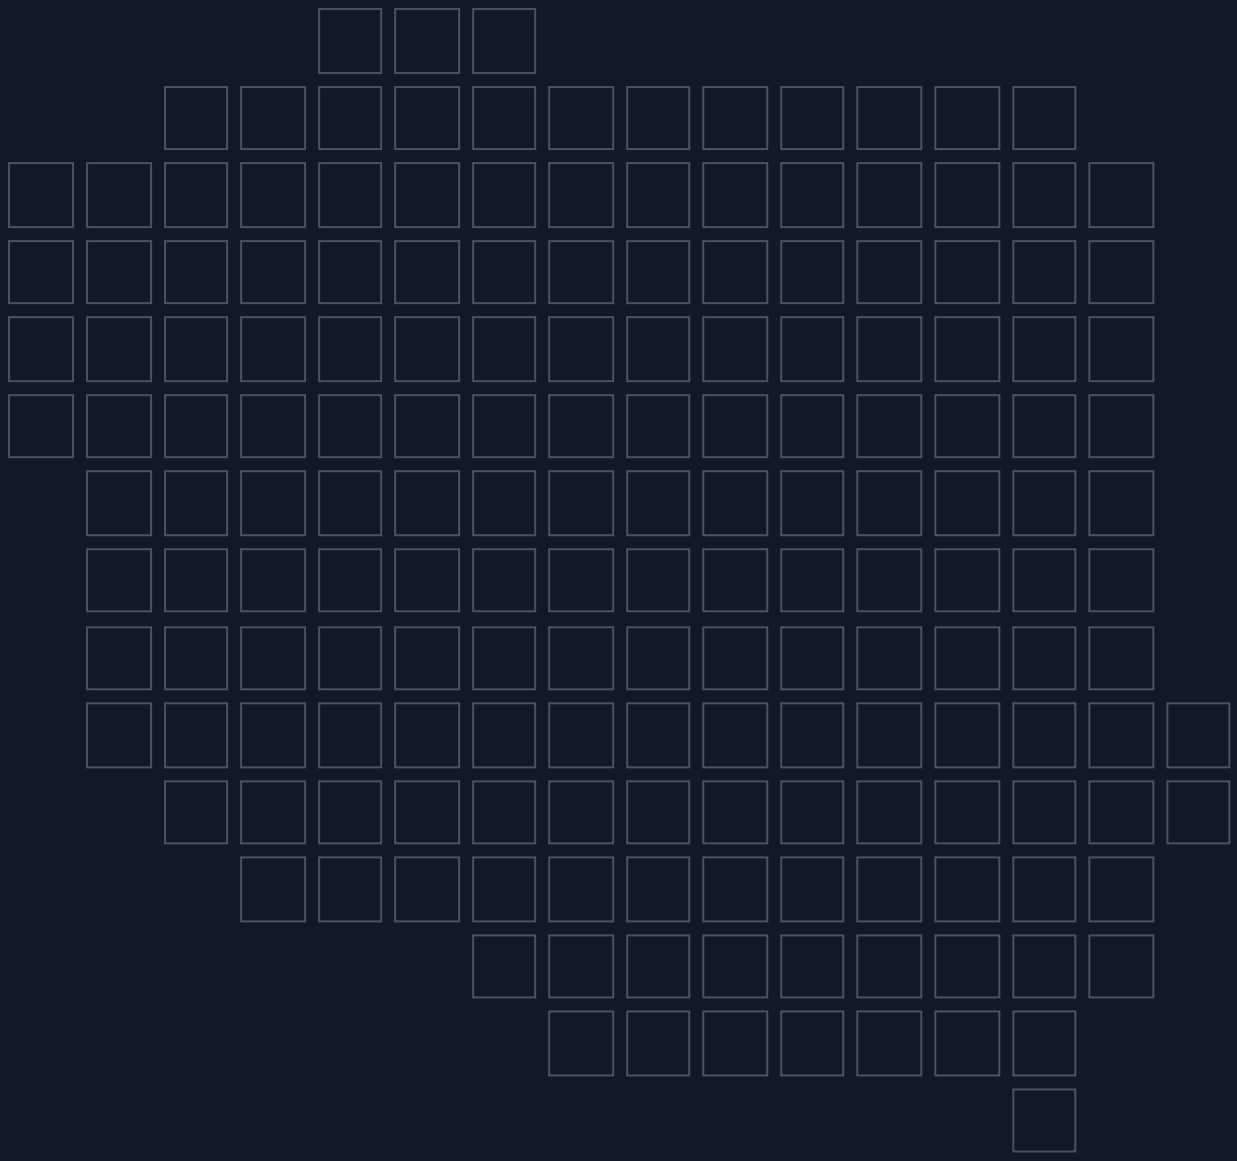
9.7 Measure, don’t audit

The last conclusion concerns method. Chapter 7 showed that after the study ended, four providers replaced their production models, and the newer versions shifted the share of explicit recommendations of Poland by a dozen or so points within seven weeks, in both directions. This is the semantic drift described in chapter 7: it is larger than measurement noise and has a direction. That is why the image of Poland in the answers is not a state that can be studied once and corrected, but a stream that changes with every model version and every change in the sources. A one-off audit shows how things stood. Promotion that wants to know whether it is working needs a measurement repeated with the same set of questions, markets and models.

The same measurement also has a limit that must be understood before models are compared. In the study, four of the six models had an upper length cap on answers, and the truncation simulation on full answers shows that cutting answers lowers the share of explicit recommendations (chapters 4 and 7, annex). The conclusions of this chapter therefore rest on whether and how Poland appears in the answers, not on comparing recommendation rates across models. A measurement intended to compare recommendation rates across models has to set the length cap so that answers are not cut.

A study like this (1,400 questions, 12 markets, six models, a dated window, every answer text recorded) is a dated snapshot, not the truth about Poland in the models. Only a measurement repeated with the same set of questions, in a dated window, can answer whether the August 2026 campaign has reached the answers, whether the cost of living has gained a Polish source and whether Masuria has begun to enter answers without a hint in the question. Each of these questions matters more than whether the models “like Poland”.

HOW WE COUNT (CHAPTER 9). Figures on questions and answers come from chapters 1–8 and from the calculations described in the annex. In this chapter we count source figures as answers in the individuals segment in which a given domain is cited at least once (4,200 answers, 2,331 with at least one citation), unlike chapter 5, which counts citations. We reduce domains to the name without “www” and “en”; we count the language versions of the Polish Tourism Organisation’s portal together as “the portal” and also separately by domain. Most citations come from one model (Perplexity, 56% of citations in the individuals segment), the only one that searched the web live in the study; the other models supplied addresses themselves, three of them when instructed to give sources (annex). The figures on sources therefore describe where an answer points, not where the model got its knowledge.



Part II

Poland for Business

CHAPTER 1

Poland in companies' questions, or tone depends on what the company asks

IN BRIEF

- Companies from 12 markets put 700 questions to six language models (LLMs) about Poland as a market, a location and a partner. Poland appears in 97.4% of the 4,200 answers, and at the consideration and decision stages in almost every one (99.9% and 99.7%).
- When a company does not name Poland in its question, the model names it of its own accord in 75.2% of answers and explicitly recommends it in 65.1% of the answers that mention it. In such questions, the models explicitly recommend Poland to individuals in 52.2% of answers (Part I, chapter 4); these are different question sets, so the figures are simply set side by side.
- The tone of answers for companies is cooler than for individuals (6.78 against 7.12 on a 1–10 scale) and falls from the awareness stage (7.02) to the consideration and decision stages (6.69).
- The question decides the tone. Of the 700 questions, 281 (40.1%) themselves ask about risk, barriers or pitfalls, or express a company's concern. The answers to them have a tone of 6.16, one in eight is negative (12.4%), and Poland is explicitly recommended in one in two (49.7%). In the remaining questions the tone is 7.20, negative answers account for 1.6% and the recommendation rate is 79.2%. The rate is almost the same at every decision stage.
- In answers for companies we confirmed 4 factual errors (a lower bound; Part I, 7.6).

A company asks about risk in 281 of 700 questions: about pitfalls, barriers, what could go wrong. In these questions the tone of answers about Poland is 6.16, and one answer in eight is negative. In the remaining questions the tone is 7.20, negative answers account for 1.6%, and Poland is explicitly recommended in 79.2% of answers, more often than to individuals. The models do not rate Poland lower when a company is asking. They answer what the company asks, and this chapter shows how much that changes the picture.

1.1 What this part measures

The study covers six generative AI models queried through their official APIs: ChatGPT, Gemini, Claude, Perplexity, Grok and DeepSeek. In the business part, 700 questions were asked by company personas from 12 markets, in the languages of those markets, across three decision stages: awareness (212 questions: where in Europe, with whom, at what cost), consideration (274: how to enter, what costs, what regulations) and decision (214: contract, partner, location). Every question

went to every model, giving 4,200 answers from the measurement window of 25 August – 28 September 2026. A "persona" here means a company scenario created for the study, not a real company (annex).

Of the 700 questions, 632 name Poland directly. The remaining 68 questions that do not name Poland are described in chapter 3, and a separate measurement of 96 such questions in chapters 2 and 4.

1.2 Poland is almost always in the answer

decision stage	questions	answers	Poland in the answer	tone (1–10)	negative answers	explicitly recommended
awareness	212	1,272	91.9%	7.02	3.9%	71.6%
consideration	274	1,644	99.9%	6.69	6.8%	63.7%
decision	214	1,284	99.7%	6.69	6.6%	68.0%
total	700	4,200	97.4%	6.78	5.9%	67.3%

Tone, negative answers and recommendation are calculated for answers that mention Poland; "negative" means a tone of 4 or lower. At the awareness stage Poland appears less often, because this is where all the questions that do not name it sit. At the two later stages the company is already asking about Poland, so the question is not "will the model name it?" but "how does it speak of it?". In the control variant without repeated scenarios (annex), the figures in this table change by less than the thresholds of measurement variability.

Companies get Poland in the answer more often than individuals do (97.4% against 92.9% in Part I), but described more coolly. Negative answers account for 5.9% in companies' questions and 3.5% in individuals' questions. The high recommendation rate is 67.3%, against 72.0%.

1.3 When the question does not name Poland

In the 68 questions that do not name Poland, the model mentions it in 75.2% of the 408 answers, and in the answers that mention it, explicitly recommends it in 65.1% (307 answers), with a tone of 7.22. This is the reverse of the picture in Part I: individuals get Poland without asking for it almost as often (71.3%), but recommended less often (52.2%). A company that does not name Poland is therefore more likely than a person planning a trip to find it on the list with an argument in its favour. Chapter 3 describes which countries share that list with Poland.

1.4 Tone depends on what the company asks

Why does tone fall between the stages? The simple answer would be that the closer the decision, the more the model knows about the problems. The data say otherwise. Of the 700 questions, 281 (40.1%) themselves ask about risk: the political and regulatory risk of a 15-year investment, pitfalls in a contract with a distributor, barriers in food labelling, the protection of intellectual property, trade restrictions, or whether a supplier is reliable. Such questions make up 27.8% at the awareness stage, 47.4% at the consideration stage and 43.0% at the decision stage. Risk questions were identified by two models acting as coders, independently of each other and without seeing the answers; they agreed on 98.6% of questions (annex). The share at the decision stage depends on repeated scenarios in the set of company questions: in the control variant with one question from each repetition group (annex), it is 35.6%, while at the first two stages it barely changes (27.4% and 47.6%). In this variant, tone and recommendation at each stage change by less than the variability threshold, so the conclusion of this section does not depend on the repetitions.

decision stage	risk questions: tone	negative answers	explicitly recommended	other questions: tone	negative answers	explicitly recommended
awareness	6.28	10.9%	52.2%	7.30	1.3%	78.9%
consideration	6.13	13.3%	46.4%	7.19	0.9%	79.4%
decision	6.14	12.0%	52.8%	7.11	2.6%	79.5%
total	6.16	12.4%	49.7%	7.20	1.6%	79.2%

At every stage, a risk question lowers the tone by about one point, and negative answers are anywhere from several to more than ten times as frequent. Risk questions account for 205 of the 243 negative answers (84.4%). In the remaining questions, Poland is explicitly recommended in 79% of answers at every stage, and the tone changes little (from 7.30 to 7.11). Roughly three fifths of the fall in tone between awareness and consideration therefore comes from the larger number of risk questions at the consideration stage.

So the model answers the question it is asked. A company that asks what could go wrong gets a list of what could go wrong. This is clearest in a question from Norway about the risk of investing in green energy in Poland compared with Germany: all six models answer with a tone of 2 or 3.

> **"Generelt kan man si at Tyskland lider av et overbyråkratisert, men svært forutsigbart system, mens Polen preges av raske politiske endringer, flaskehalser i infrastrukturen og en mer umoden konsesjonsprosess."**
>
> "Broadly speaking, Germany suffers from an over-bureaucratised but highly predictable system, whereas Poland is marked by rapid political change, infrastructure bottlenecks and a less mature licensing process."
>
> — Gemini, Norway, the management board of a company considering a large green-energy investment in Eastern Europe (Log ID 8f1cf372-de6c-4e9c-bdae-d22f9dde5f20). Consideration stage; tone 3, low recommendation level.

The same contrast recurs in chapter 4, where suppliers are asking: in 9 of 14 answers without an explicit recommendation, Poland is "a base that is growing", while Germany and Italy are the benchmark, described as reliable. In risk questions, such a contrast stops being a description and becomes a judgement.

The conclusion has two sides. First, the lower tone of answers for companies does not mean that the models rate Poland poorly as a place to do business; it is largely a result of what companies ask. Companies that do not ask about risk get Poland explicitly recommended more often than individuals do (79.2% against 72.0% in Part I; these are different question sets). Second, what the models say about risk in Poland is content that a company takes into its decision. If that content calls Poland "unpredictable" and Germany "predictable", the argument against Poland works regardless of the tone of the answer as a whole. The risk of war as a separate topic is covered in chapter 8.

1.5 Errors

The 4,200 answers for companies contain 4 confirmed factual errors (two in answers in the Spanish market, one each in the Italian and Norwegian markets). The verification method is described in Part I, chapter 7, and in the annex.

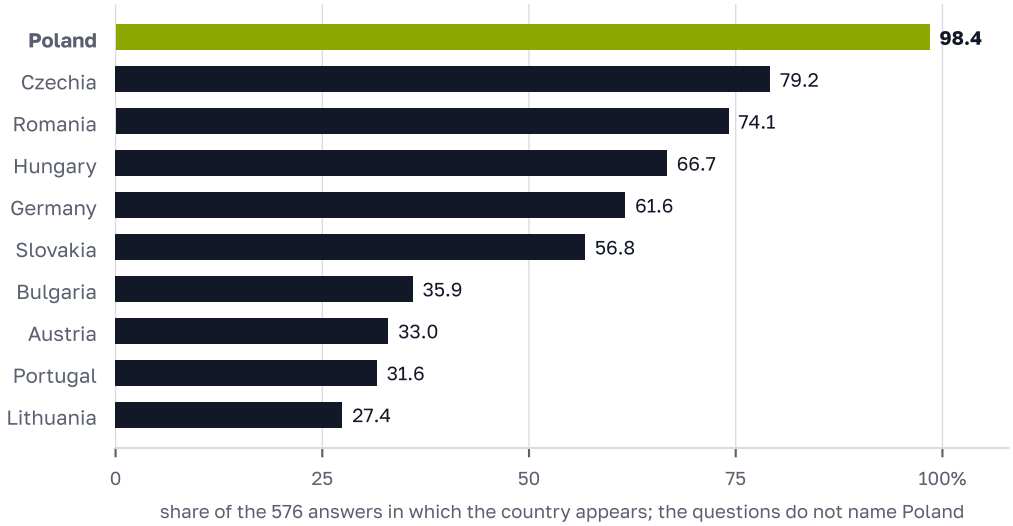
CHAPTER 2

The shortlist, or do the models put Poland forward to companies of their own accord?

IN BRIEF

- When a company asks where in Europe to build a plant, whom to outsource production to, where to find engineers or which market to enter first, and does not name Poland, language models (LLMs) bring Poland up of their own accord in 98.4% of answers (567 of 576), and as a country, not merely through a Polish city, in 97.4%. This is a separate measurement, taken on 30 September 2026: 96 questions that do not name Poland, 8 decision areas, 12 markets, six models.
- Poland is almost always on the shortlist, and it is explicitly recommended there in 82.5% of the answers that mention it. For comparison, when individuals look for a destination without naming Poland, the models explicitly recommend it in 52.2% (Part I, chapter 4). These are two different sets of questions, so the figures stand side by side, not on a single chart.
- Here, the question type does not decide visibility. When the question names a region, Poland appears in 99.6% of answers; when it refers to Europe in general, in 97.7%. For individuals, the same difference was 47 percentage points (Part I, chapter 2).
- What does matter is the decision area. It matters not for whether Poland is named but for whether it is recommended: from 93.1% in questions about sales markets and 93.0% in those about engineers to 69.4% in questions about contract manufacturing and 62.3% in those about a warehouse serving Europe. Why the area changes the recommendation is explained in chapter 4.
- Poland competes against the same list in every market: Czechia (79.2% of answers), Romania (74.1%), Hungary (66.7%), Germany (61.6%), Slovakia (56.8%). Romania, which barely features in questions from individuals, is the second candidate for companies after Czechia (chapter 3).

Chart 2.1 The companies' shortlist



Of the 700 questions from companies in the study, only 68 do not name Poland. Companies that have Poland in mind ask about it directly, so the main study cannot answer the question of whether the models put Poland forward to companies of their own accord. A separate measurement can: 96 questions that do not name Poland, from "where in Central Europe should we build a plant" to "where can we find engineers", put to six models in 12 markets. The result: Poland is on the shortlist in 567 of 576 answers and is explicitly recommended in eight out of ten. Presence, however, is only the start. The rest of this chapter looks at the role in which Poland appears.

2.1 What this chapter measures

In the main study, companies asked 700 questions, but only 68 of them do not name Poland, because companies that already have Poland in mind ask about it directly (chapter 1). To test whether the models place Poland on the shortlist of their own accord, a separate measurement was set up: 96 awareness-stage questions, eight for each of the 12 markets, in the language of the market, with no mention of Poland, a Polish city or the Visegrád Group in the wording. The eight decision areas are: sales markets in Central and Eastern Europe, nearshoring location, contract manufacturing, partners and suppliers, shared-services and IT centres, engineering talent, regional logistics, and grants and economic zones. Each question was put once to six models, giving 576 answers.

The measure is the same as in Part I: the share of answers in which Poland is named even though the question does not name it. The questions did not ask about Poland, but they were questions about Europe or its centre, where Poland is a natural candidate. The measurement therefore shows whether a model places Poland on the shortlist when a company asks an open question about a category, not whether it "remembers" Poland without any cue at all. "Shortlist" is the name of the measurement: we count the presence of Poland in the answer. In 11 logistics answers Poland is only a market served from another country (chapter 4).

2.2 Poland is almost always on the list

	answers	Poland in answer	as a country	explicitly recommended (of mentions)	tone	home country in answer
business questions that do not name Poland (B6)	576	98.4%	97.4%	82.5%	7.85	31.1%

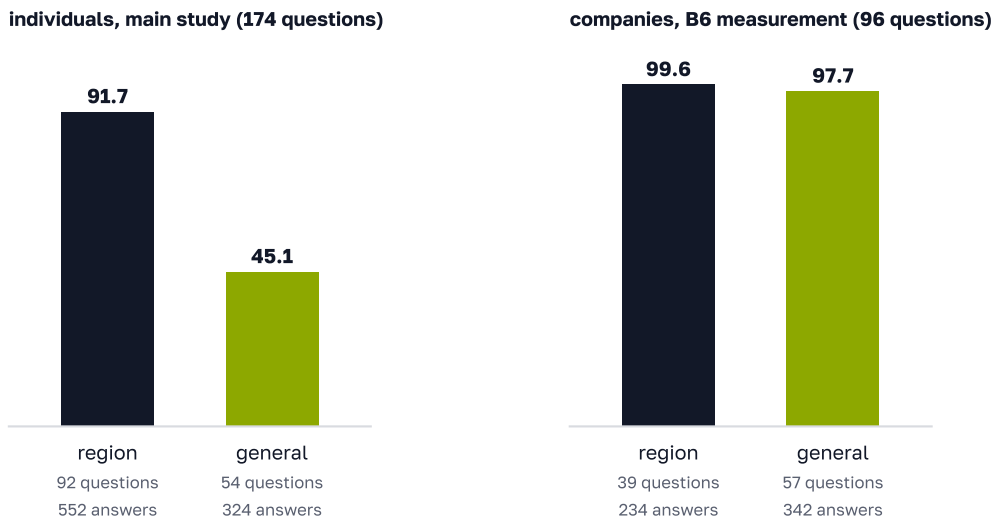
Nine answers in ten mention Poland and eight in ten explicitly recommend it. Only 9 of the 576 answers leave Poland out. Tone is higher than in any breakdown in Part I (7.85 against 6.95 for the study as a whole), but these are different questions with a different denominator, so the difference is a description, not a finding.

The persona's home country appears in 31.1% of answers, again most often in the neighbouring markets: Czechia 75.0%, Germany 72.9%, Slovakia 60.4%. As in Part I (chapter 3), it serves as a reference point for costs and distance, not as advice to "stay at home": in questions from Czechia and Germany, the models compare Polish locations with domestic ones.

2.3 The question type does not change visibility

The questions were classified using the same procedure as in Part I (two independent runs without access to the results; annex): 39 name a region that includes Poland ("in Central Europe", "in a neighbouring country"), and 57 refer to Europe or the Union in general. For individuals, this split decided everything: 91.7% against 45.1%. For companies, it decides nothing as regards visibility and a good deal as regards the recommendation rate.

Chart 2.2 Question type: companies versus individuals



Poland in the answer, % of answers; two measurements with separate denominators: the panels are not connected

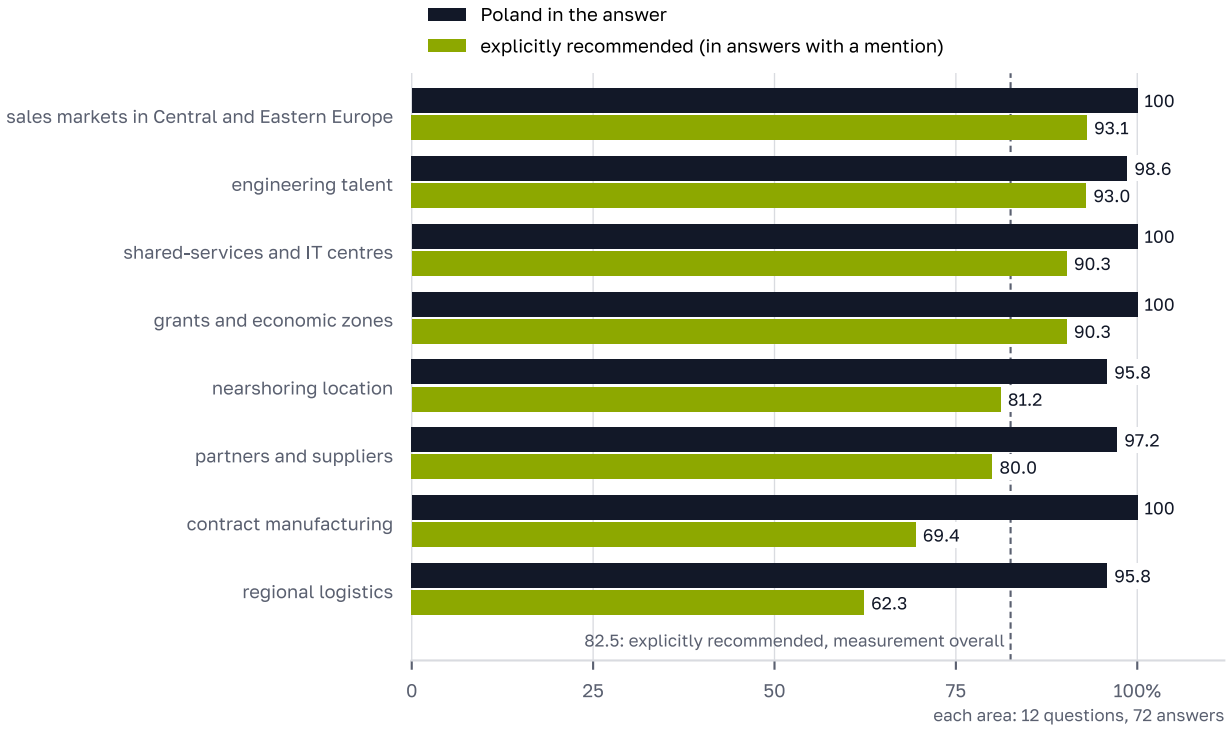
question type (companies, Poland not named in the question)	questions	answers	Poland in answer	explicitly recommended	tone
names a region that includes Poland	39	234	99.6%	89.3%	7.91
general (Europe, EU)	57	342	97.7%	77.8%	7.81

Why does a company get Poland even when it asks about Europe as a whole? Because a company's question usually has a criterion built into it: cost, distance from Germany, access to engineers, an economic zone. The model does not assemble a "default Europe", as it would for someone planning a trip, but a list of the countries that meet the criterion, and Poland is almost always on that list. The difference emerges only in the recommendation: in regional questions Poland is recommended in 89.3% of answers, in general questions in 77.8%, that is, 11.4 percentage points less often. The direction is the reverse of that for individuals, where general questions gave Poland a higher recommendation rate (Part I, chapter 4), and it recurs in five of the six models.

2.4 The decision area decides the recommendation

decision area (12 questions, 72 answers each)	Poland in answer	explicitly recommended	tone	home country
sales markets in Central and Eastern Europe	100%	93.1%	7.97	11.1%
engineering talent	98.6%	93.0%	7.93	41.7%
shared-services and IT centres	100%	90.3%	7.83	20.8%
grants and economic zones	100%	90.3%	8.11	15.3%
nearshoring location	95.8%	81.2%	7.59	25.0%
partners and suppliers	97.2%	80.0%	7.94	43.1%
contract manufacturing	100%	69.4%	8.01	50.0%
regional logistics	95.8%	62.3%	7.41	41.7%

Chart 2.3 The area decides the recommendation



Visibility is close to 100% everywhere, so what the models tell companies about Poland is settled by the “explicitly recommended” column. The spread is 31 percentage points. At the top are sales markets, engineers, service centres and economic zones: there, Poland is not just one of the countries on the list but a country with an argument behind it (market size, supply of engineers, incentives). At the bottom are contract manufacturing and regional logistics. Why this is so, and which countries Poland loses out to in these areas, is described in chapter 4.

2.5 Who Poland competes with for companies’ decisions

Answers to companies, like answers to individuals, take the form of lists. The countries most often named alongside Poland are Czechia (79.2%), Romania (74.1%), Hungary (66.7%), Germany (61.6%)

and Slovakia (56.8%), and Poland comes first in every one of the 12 markets. The make-up of this list, how far it matches the main study and how it differs between markets are described in chapter 3.

2.6 The models

model	answers	Poland in answer	tone
Gemini	96	100%	8.19
Grok	96	100%	8.10
Perplexity	96	94.8%	8.02
DeepSeek	96	100%	7.86
ChatGPT	96	99.0%	7.48
Claude	96	96.9%	7.45

Visibility differs between the models by 5 percentage points (Perplexity 94.8%, i.e. 5 of 96 answers without Poland; four models at 99–100%). Tone differs between the extreme models by 0.74 points (from 7.45 to 8.19), but with 96 answers per model this is a tendency. The table does not give recommendation rates by model. In this measurement, four models had an answer length cap and reached it in anywhere from one answer in five to eight in ten, and the truncation simulation on full answers (Part I, annex) shows that cutting an answer short lowers the share of explicit recommendations. In this measurement, the differences in the recommendation rate between models are smaller than that possible effect, so the order of the models cannot be established from it; the values are given in the annex. Unlike the main study, the business measurement also had no second judge and no repeat runs. A comparison of the models would therefore require control measurements at fixed intervals, for example quarterly, because the models are updated.

2.7 How it sounds in the answers

> **“If your goal is to cut customs friction on EU orders while staying close to UK decision-making, I would shortlist Poland, the Netherlands, Germany, and Ireland.”**
>
> — Perplexity, United Kingdom, a food-packaging manufacturer considering moving part of its production to the EU after Brexit (Log ID d71ad1c0-ebc7-4cef-9aaa-295473ac1b9b). General question about Europe; Poland named, high recommendation level.

> **“Pour des vins et spiritueux haut de gamme, les marchés les plus prometteurs en Europe centrale et orientale sont, en priorité, la Pologne, la République tchèque, la Roumanie et la Hongrie.”**
>
> “For premium wines and spirits, the most promising markets in Central and Eastern Europe are, first and foremost, Poland, the Czech Republic, Romania and Hungary.”
>
> — Perplexity, France, a wine and spirits producer looking for sales markets in Central and Eastern Europe (Log ID afa9c7ee-8c9b-437a-ac8b-1f3232db76e9). Regional question; Poland named, tone 9.

> **“Poland | One of Europe’s deepest electronics-manufacturing bases”**
>
> — ChatGPT, United States, a consumer-electronics start-up looking for a contract manufacturer in Europe so that it can sell “Made in EU” products (Log ID 87d588fe-92c3-4b6f-a880-a23ce0a5797b). General question; Poland in the table of locations, tone 9.

HOW WE COUNT. The measurement of business questions that do not name Poland (B6) consists of 96 awareness-stage questions (8 areas × 12 markets) in the language of the market, put once to each of the six models on 30 September 2026 under separate research identities: 576 answers, all assessed by the same judge as the main study. Poland in answer = a mention of Poland or a Polish city; explicitly recommended = the “high” recommendation level, among answers that mention Poland; home country = the country from which the question was asked (the persona’s home country). Question type follows the classification in the annex (two independent runs). Countries named alongside Poland are counted from the judge’s field after name normalisation, as a share of answers. The business measurement is not a check on the main study and is not combined with its figures; generative visibility in questions from individuals in Part I (71.3%) and the visibility for companies here (98.4%) come from different questions and never appear on the same chart. One run per question (the repeatability test in the main study shows that the result is repeated in about 87% of question × model pairs); there are 8 questions per market, which is why we give market figures with n. Answer length was capped by the research application, as in the main study (annex: “Truncated answers”). The measurement concerns the model’s answer as returned by the API.

CHAPTER 3

Which countries does Poland compete with for companies' decisions?

IN BRIEF

- The list of countries with which Poland competes for companies' decisions does not depend on the set of questions. In the 68 company questions from the main study that do not name Poland (408 answers) and in the 96 questions of the separate measurement (576 answers), the top six places come out in the same order: Czechia, Romania, Hungary, Germany, Slovakia, Bulgaria.
- In company questions, Poland is first on this list even when the question is about Europe in general (69.5% of answers against 55.7% for Germany). For individuals, such questions were headed by Germany (Part I, Chapter 3). A company asks with a criterion in mind (cost, engineers, distance from Germany), and Poland meets it.
- The four questions with no hint of going abroad (about procedures: a representative office, economic zones, choosing a partner, suppliers) produce Poland in 1 answer out of 24, and in eight out of ten produce no country at all. When the question does not say "abroad", the model stays in the country the question comes from.
- For companies from outside Europe, Germany is the point of reference: in questions from the United States and China, Germany is named more often than Poland (78.7% against 69.4%, and 66.7% against 56.7%), and in the separate measurement it is second on the list. The difference comes from procedural questions, such as the sales cycle or registering a company, in which the model describes Europe through Germany. In questions about where to locate a plant, Poland is first or second in questions from the United States and China.
- The models differ more here than in the separate measurement: Poland appears in 83.8% of the answers from Gemini and Grok, and in 57.4% of those from Perplexity. In more than a third of their answers, ChatGPT and Perplexity name no country at all. They answer with a procedure, not a list.

Czechia, Romania, Hungary, Germany, Slovakia, Bulgaria. That is the order in which the models name Poland's competitors in the 68 company questions from the study and in the 96 questions of the separate measurement. These are two sets collected a month apart, with different questions. The list does not depend on the set of questions, and that is the first finding of this chapter. The second: Romania, which barely featured in questions from individuals, is the second candidate for companies, after Czechia.

3.1 Two sets, one list

Chapter 2 showed that when a company asks without naming Poland, language models (LLMs) put it on the shortlist in 98.4% of answers. The question for this chapter is: alongside whom? Answers to companies, like answers to individuals, take the form of lists. The countries on those lists can be counted in two independent sets. The first is the main study: of its 700 company questions, 68 do not name Poland. Of these, 35 point to a region that includes Poland, 29 are about Europe in general and 4 do not point abroad at all; together they produced 408 answers across three decision stages. The second is the Chapter 2 measurement: 96 awareness-stage questions, 576 answers. The questions were written separately, at a different time and for different decision areas. If the list of competitors is a property of the models rather than of the set of questions, both sets should produce the same result.

country alongside Poland	main study, company questions not naming Poland (408 answers)	shortlist measurement, companies, Poland not named in the question (576 answers)
Poland	75.2%	98.4%
Czechia	62.7%	79.2%
Romania	53.9%	74.1%
Hungary	53.7%	66.7%
Germany	50.5%	61.6%
Slovakia	44.1%	56.8%
Bulgaria	32.1%	35.9%
Austria	27.0%	33.0%

Chart 3.1 One list in two sets



share of answers in which the country appears; the same order of countries in both panels

They do. The top six places come out in identical order, and the seventh (Austria) closes both lists. The levels differ because the sets differ: the main study includes procedural questions and the consideration and decision stages, whereas the Chapter 2 measurement contains only questions

about where to go, at the awareness stage. The order is the same. However a company phrases its question about Central Europe, it gets a single map from the models: Poland, Czechia, Romania, Hungary, Germany, Slovakia.

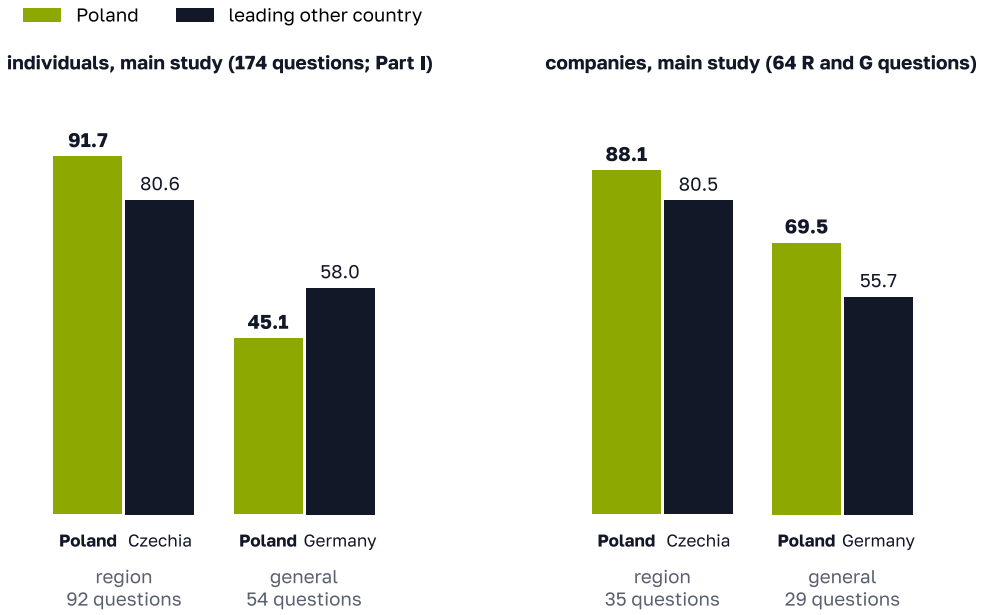
This map differs from the one shown to individuals in two places. Romania, which barely featured in answers to individuals, is the second candidate for companies, after Czechia. Germany, which headed every list in general questions from individuals, is fourth for companies. Companies are given a map of costs and production capacity, not a map of attractions.

3.2 Question type: Poland first in general questions too

In Part I, question type decided everything: when a question from an individual was about Europe in general, Poland appeared in 45.1% of answers and lost out to Germany. For companies, the same three question types (annex: classification by one coder, checked by a second run on a sample) produce a different picture.

question type (companies, Poland not named in the question)	questions	answers	Poland in the answer	explicitly recommended (of mentions)	no other country	home country	three most frequent other countries
points to a region that includes Poland	35	210	88.1%	63.8%	10.5%	18.6%	Czechia 80.5 · Hungary 76.2 · Slovakia 67.1
general (Europe, EU)	29	174	69.5%	67.8%	16.1%	7.5%	Germany 55.7 · Romania 52.9 · Czechia 50.0
no hint of going abroad	4	24	4.2%	–	79.2%	16.7%	France 16.7 · United States 8.3
total	68	408	75.2%	65.1%	16.9%	13.7%	Czechia 62.7 · Romania 53.9 · Hungary 53.7

Chart 3.2 Question type for companies and for individuals



% of answers in which the country appears; region = questions that point to a region, general = general questions; the panels are not connected

In region questions Poland is 7.6 percentage points ahead of Czechia; in general questions it is 13.8 percentage points ahead of Germany. For individuals, the gap in general questions was negative. Why? Because a company's question "about Europe" is seldom a question about Europe. It is a question about a country that meets a condition: labour costs, the supply of engineers, delivery to Germany within 48 hours, an economic zone. The model then assembles not a "default Europe" but a list of countries that meet the condition, and on that list Poland comes first. The make-up of the list changes: in region questions Poland is joined by Czechia, Hungary and Slovakia; in general questions by Germany, Romania and Czechia, followed by Portugal (42.5%) and Spain (39.1%), which are all but absent from region questions (3.8% and 2.9%). Portugal and Spain come in through questions about service centres and software engineers, not through questions about plants.

Here the recommendation rate runs in the opposite direction to visibility: in general questions Poland is explicitly recommended in 67.8% of the answers that mention it, in region questions in 63.8%. The 4-percentage-point difference falls within the threshold of measurement noise and is a description, not a finding; in the Chapter 2 measurement the direction was reversed (89.3 against 77.8). Both sets say the same thing: question type does not decide whether Poland is on a company's list.

> **"Assuming a general industrial operation rather than a highly specialized semiconductor, pharmaceutical, or battery plant, the strongest shortlist is in Central and Eastern Europe. My overall ranking is: […] 1 | Poland | Best overall balance: large labor pool, broad supplier base, competitive costs, resilient economy, good logistics, and substantial EU funding […] 2 | Czechia […] 3 | Slovakia […] 4 | Romania […] 5 | Hungary […]"**
>
> — ChatGPT, United States, a mid-sized manufacturing company planning its first plant in Europe; general question about the countries with the best combination of costs, workforce and stability (Log ID e12878f1-a315-407a-b21a-30a4d6c163a8). High recommendation level; the answer contains a table of seven countries (the quotation gives the first three columns: rank, country, strengths; omitted: the "main drawbacks" column and the last two rows).

3.3 When the question does not say "abroad", the model stays at home

Four of the 68 questions point neither to a region nor abroad: on a representative office without a company, on economic zones for a planned factory, on choosing an engineering partner and on electronics suppliers. These are questions about procedure, not direction, and the models answer with a procedure: 19 of the 24 answers name no country, and one names Poland. One example is instructive. A company from France has been offered "several locations in special economic zones" and asks how to compare them. Four of the six models answer with French zones (ZFU, ZRR, ZAFR), one compares Morocco, India and Turkey, and none names the Polish zones, even though in Poland the "special economic zone" is an established institution. When the question does not say where, the model assumes it means at home. The lesson for economic promotion: Poland appears in the answer once a company already knows it is looking abroad; the models do not bring Poland into the decision of whether to go abroad at all.

3.4 Europe and the rest of the world: for whom Germany is the point of reference

The 68 questions are unevenly distributed across markets: the United States 18 questions (108 answers), France 16 (96), Italy 11 (66), China and Germany 5 each (30 each), and the remaining markets 1–2 questions each. A comparison between markets therefore makes sense only for the first five, and then as a description, not a ranking.

market of the question	questions	answers	Poland in the answer	most frequent other country	home country
China	5	30	56.7%	Germany 66.7%	0%
France	16	96	74.0%	Czechia 60.4%	19.8%
Germany	5	30	70.0%	Romania 66.7%	30.0%
Italy	11	66	68.2%	Czechia 56.1%	12.1%
United States	18	108	69.4%	Germany 78.7%	4.6%

In two markets Germany is named more often than Poland, and both are outside Europe. In questions from France, Germany and Italy, the first country alongside Poland is Czechia or Romania. The separate Chapter 2 measurement shows the same pattern from the other side: there Poland appears in the answers everywhere, but Germany is the second country on the list in questions from China (85.4%) and the United States (85.4%), fourth in those from France (60.4%) and fifth in those from Italy (50.0%).

The table for the Chapter 2 measurement shows this for all 12 markets, with 48 answers per market:

market	Poland	three most frequent other countries
China	100%	Hungary 87.5 · Germany 85.4 · Czechia 85.4
Czechia	97.9%	Slovakia 83.3 · Hungary 77.1 · Romania 75.0
France	100%	Romania 89.6 · Czechia 83.3 · Hungary 66.7
Germany	100%	Czechia 95.8 · Hungary 87.5 · Slovakia 75.0
Italy	95.8%	Romania 83.3 · Czechia 77.1 · Hungary 77.1
Norway	97.9%	Romania 85.4 · Czechia 66.7 · Hungary 56.3
Slovakia	97.9%	Czechia 93.8 · Hungary 79.2 · Romania 70.8
Spain	100%	Czechia 75.0 · Germany 70.8 · Romania 66.7
Sweden	95.8%	Czechia 77.1 · Romania 62.5 · Germany 54.2
Ukraine	100%	Romania 83.3 · Czechia 62.5 · Germany 60.4
United Kingdom	95.8%	Czechia 70.8 · Romania 60.4 · Germany 54.2
United States	100%	Czechia 87.5 · Germany 85.4 · Romania 66.7

The markets differ in the make-up of the list, not in Poland's position. In questions from France, Italy, Norway and Ukraine, the first country alongside Poland is Romania; from Germany, Slovakia, Spain, Sweden, the United Kingdom and the United States, Czechia; from Czechia, Slovakia; from China, Hungary. The differences in Poland's visibility between markets (95.8–100%) fall within measurement variability at 48 answers per market, which is why the table is not a ranking.

Where does the difference come from? From the kind of questions a company from outside Europe asks about Europe. Of the questions from the United States, five concern procedures: the sales cycle, registering a company, the hidden costs of market entry, the timetable for starting production. Poland appears in 7 of those 30 answers; the model describes Europe through Germany, DACH, the GmbH and France. In the other 13 questions (about a plant, a hub, engineers, suppliers), Poland appears in 68 of 78 answers. From China, two questions about CE certification yield 0 of 12, and three location questions 17 of 18. For a company from outside Europe, then, Germany is the name for Europe, and

Poland is the answer to the question of where in Europe to manufacture. This distinction returns in Chapter 7, in the market cards for China and the United States.

> **"Typical B2B sales cycles in Europe are longer and more relationship-driven than in the US [...] Nordics & UK: Closer to 6–12 months [...] DACH region (Germany, Austria, Switzerland): Frequently 12–24 months due to consensus-driven buying [...] France, Benelux, Southern Europe: 9–15 months [...]"**
>
> — Grok, United States, a company asking about the typical B2B sales cycle when entering the European market (Log ID 4483ea75-8660-4672-84d1-39d9695af946). General question; Europe described through four regions, without Poland and without Central Europe.

China has one further distinctive feature. In the question about a first production plant in Europe, ChatGPT recommends Poland, citing "市场规模、劳动力、物流和激励最均衡" ("the best balance of market size, workforce, logistics and incentives"), and DeepSeek rates Hungary higher, justifying this explicitly by the experience of Chinese investors: "奥迪、奔驰、宝马、宁德时代、比亚迪等均在匈牙利布局" ("Audi, Mercedes, BMW, CATL and BYD all have plants in Hungary"), while on Poland it says: "中资制造企业有布局，但不如匈牙利在电池/汽车领域集中" ("Chinese manufacturers are present, but not as concentrated in batteries and automotive as in Hungary"). The model answers a Chinese company with what other Chinese companies have already done. That is an argument Poland does not have in the answers, and Chapter 10 returns to it.

3.5 The models

model	answers	Poland in the answer	no other country	other countries on average	most frequent other country
ChatGPT	68	63.2%	36.8%	5.97	Czechia 55.9%
Gemini	68	83.8%	5.9%	7.01	Romania 64.7%
Claude	68	80.9%	13.2%	5.00	Czechia 64.7%
Perplexity	68	57.4%	35.3%	4.47	Czechia 44.1%
Grok	68	83.8%	4.4%	6.26	Czechia 75.0%
DeepSeek	68	82.4%	5.9%	8.15	Czechia 73.5%

Every model received the same 68 questions, so here the comparison is fair. The spread is 26 percentage points (Gemini and Grok 83.8%, Perplexity 57.4%). That is five times wider than in the Chapter 2 measurement, where it was 5 percentage points. The difference comes from the "no other country" column: in more than a third of their answers ChatGPT and Perplexity name no country, because they answer with a procedure (a list of steps, costs to check, questions to ask), where Gemini, Grok and DeepSeek answer with a list of countries. This is not reluctance towards Poland: when ChatGPT does name countries, Poland comes first (63.2% against 55.9% for Czechia). It is a difference in answering style that a company should be aware of: the same question produces a procedure from Perplexity and a ranking from Grok. Chapter 9 returns to this difference.

> **"L'Europe centrale (Pologne, République tchèque, Slovaquie, Hongrie, Roumanie, Bulgarie) reste attractive pour le nearshoring de la métallurgie (usinage, tôlerie, fonderie, soudure) grâce à l'appartenance à l'UE, des coûts inférieurs à l'Europe de l'Ouest et une logistique raisonnable. [...] Pologne : ~12-14 €/h; République tchèque : ~14-16 €/h; Roumanie : ~8-10 €/h; Allemagne de l'Ouest : ~35-40 €/h"**
>
> "Central Europe (Poland, Czechia, Slovakia, Hungary, Romania, Bulgaria) remains attractive for nearshoring metalworking (machining, sheet-metal work, foundry, welding) thanks to EU membership, costs lower than in Western Europe and reasonable logistics. [...] Poland: approx. €12–14/h; Czechia: approx. €14–16/h; Romania: approx. €8–10/h; western Germany: approx. €35–40/h"
>
> — Grok, France, a company assessing locations for metalworking subcontracting in Central Europe (Log ID f3bc18ea-bfc9-460c-b450-376088c40837). Regional question; Poland on a list of six countries, high recommendation level; the hourly cost of labour as the main criterion.

HOW WE COUNT. Set 1: 68 company questions from the main study that do not name Poland, a Polish city or a Polish region (listed in the open data of the report), with one answer from each of the six models in the measurement window 25.08–28.09.2026, giving 408 answers. Set 2: the Chapter 2 measurement (96 questions, 576 answers, 30.09.2026). Question type follows the classification in the annex (two independent runs without sight of the results; 164 questions = 68 + 96, agreement 97.6%). Countries alongside Poland are counted from the judge's field after name normalisation, as the share of answers in which the country appears; "home country" = the home country of the persona asking the question. The two sets have separate denominators and are not added together; Table 3.1 sets them side by side, not on a single axis. We give per-market figures only for markets with at least five questions (30 answers), and only as a description; seven markets have two questions or fewer in Set 1. We treat differences in recommendation rate below 5 percentage points as noise (annex).

CHAPTER 4

Two maps, or the decision area chooses Poland's competitors

IN BRIEF

- A company that does not name Poland in its question is given Poland by the language models (LLMs) almost every time, in each of the eight decision areas (95.8–100% of answers). What differs is something else: which countries Poland stands alongside, and whether it is recommended. Explicit recommendation ranges from 93.1% in questions about sales markets to 62.3% in questions about logistics (chapter 2).
- The decision area chooses the map. In questions about sales markets, service centres, economic zones and nearshoring, Poland's competitors are its Central and Eastern European neighbours, Romania, Hungary, Czechia and Slovakia, and Poland is recommended in 81–93% of answers. In questions about suppliers, contract manufacturing and logistics, Germany, Italy, Portugal, Austria and the Netherlands join the list, and recommendation is lower, lowest of all in contract manufacturing (69.4%) and logistics (62.3%).
- This is not an effect of question type. Among questions about Europe in general alone, recommendation ranges from over 93% (service centres, engineers) to 62.3% (logistics).
- In the industrial areas the model more often names the company's own country: in contract manufacturing in half of the answers (50.0%), for suppliers in 43.1%, in logistics in 41.7%. In questions about sales markets, in 11.1%.
- To suppliers, the models describe Poland as a base that is growing: that is how 9 of the 14 answers in which Poland is not explicitly recommended read. Germany is the benchmark in them.

A company that does not name Poland is given Poland by the models almost every time: in 95.8–100% of answers in each of the eight decision areas. Yet in questions about sales markets Poland is explicitly recommended in 93.1% of answers, and in questions about a warehouse for Europe in 62.3%. Thirty percentage points apart, with the same presence. This chapter asks what happens between those two figures, and its answer is that the map on which Poland stands changes, and with it the competitors.

4.1 What this chapter measures

Chapter 2 showed that the decision area determines whether the model recommends Poland to a company. This chapter asks why. The material is the same measurement of 30 September 2026: 96 questions that do not name Poland, eight areas of 12 questions each, one question from each of the 12 markets in every area, and the answers of six models. That gives 72 answers per area, but only six per area and market. The chapter therefore compares areas, not markets.

We count competitors as in chapter 3: as the countries the model names in the same answer as Poland. "Home country" means the country of the company on whose behalf the question was asked.

4.2 Question type does not explain the differences

The first suspicion is a simple one. At the bottom of the table in chapter 2 are the areas in which almost all the questions refer to Europe in general (contract manufacturing, suppliers, logistics), and at the top those in which the questions point to a region (sales markets, economic zones). For individuals, question type determined whether Poland came up at all (Part I, chapter 2). Is it the same here? No.

decision area	questions about Europe in general	explicitly recommended	regional questions	explicitly recommended
engineering talent	11	93.8%	1	—
shared-services and IT centres	4	95.8%	8	87.5%
nearshoring location	8	87.0%	4	69.6%
partners and suppliers	10	77.6%	2	—
contract manufacturing	10	66.7%	2	—
regional logistics	12	62.3%	0	—
sales markets in Central and Eastern Europe	0	—	12	93.1%
grants and economic zones	2	—	10	95.0%

Among questions about Europe in general alone, the spread exceeds 30 percentage points: a company looking for engineers gets Poland with a recommendation in 93.8% of answers (a service centre in 95.8%, on four questions), a company looking for a warehouse in 62.3%. The question type is the same; what differs is the subject of the decision. The table omits cells with one or two questions, because with six or twelve answers a percentage says nothing.

4.3 Region and industry draw two different maps for Poland

decision area	explicitly recommended	three most frequent other countries (share of answers)
sales markets in Central and Eastern Europe	93.1%	Romania 100 · Hungary 97.2 · Czechia 88.9
engineering talent	93.0%	Romania 84.7 · Germany 75.0 · Czechia 66.7
shared-services and IT centres	90.3%	Romania 94.4 · Hungary 84.7 · Czechia 79.2
grants and economic zones	90.3%	Hungary 97.2 · Czechia 86.1 · Slovakia 84.7
nearshoring location	81.2%	Czechia 79.2 · Romania 75.0 · Hungary 70.8
partners and suppliers	80.0%	Germany 98.6 · Italy 91.7 · Czechia 87.5
contract manufacturing	69.4%	Germany 76.4 · Portugal 75.0 · Romania 69.4
regional logistics	62.3%	Germany 94.4 · Czechia 88.9 · Austria 80.6

The table falls into two maps. The first is the map of the region. When a company asks where to sell, where to open a service centre, where it will get incentives and where to move production closer to Europe, the model compares Poland with its Central and Eastern European neighbours. On this map Poland has an argument in the answers: the size of its market, its supply of engineers, its incentives (chapter 2). Recommendation stands at 81–93%.

> **"Polsko je největší příležitost, ale vyžaduje vlastní obchodní zastoupení nebo silného lokálního partnera/distributora."**
>
> "Poland is the biggest opportunity, but it requires a sales office of your own or a strong local partner or distributor."
>
> — Grok, Czechia, construction-chemicals manufacturer looking for export markets in Central and Eastern Europe (Log ID 2a5c83d2-513b-4c08-ac73-39ec0ac5ca3e). Regional question; high recommendation level.

The second is the map of industry. When a company asks whom to buy components from, whom to contract production to and where to put a warehouse, Germany, Italy, Portugal, Austria and the Netherlands join the list. The models justify them by their industry (Germany, Italy) or by their location closer to western markets (the Netherlands, Austria; 4.6). Poland is almost always on this map (97–100% of answers), but it is recommended less often: in 80.0% of answers for suppliers, in 69.4% in contract manufacturing and in 62.3% in logistics.

When Poland is not explicitly recommended, it is clear to whom it gives way. In questions about suppliers, Germany is named in all 14 such answers, and Italy in 13. In contract manufacturing (22 answers) the countries named most often are Germany (72.7%), Portugal (68.2%) and Italy (59.1%). In logistics (26 answers), Germany and Czechia (84.6% each), Austria (73.1%) and the Netherlands (65.4%); why the Netherlands and Germany beat Poland as a warehouse for the whole of Europe is described in 4.6.

The maps also differ in their sources. In questions about grants and economic zones, the answers point to the Polish investment agency: paih.gov.pl, the site of PAIH, the Polish Investment and Trade Agency, appears in 29 of the 39 answers with citations, alongside the Hungarian, Czech, Romanian and Slovak agencies. This is the only area in which a Polish source appears in the majority of answers with citations. In the others, the answers point mainly to EU sources, company directories and industry websites (chapter 6).

4.4 In questions about production, the model names the company's own country

decision area	home country in the answer
contract manufacturing	50.0%
partners and suppliers	43.1%
engineering talent	41.7%
regional logistics	41.7%
nearshoring location	25.0%
shared-services and IT centres	20.8%
grants and economic zones	15.3%
sales markets in Central and Eastern Europe	11.1%

In questions about contract manufacturing, every second answer names the company's own country. This happens most often in questions from companies in Czechia, Germany, Italy, Norway, Slovakia and Spain; to companies from China, Ukraine and the United States asking about production in Europe, the model does not suggest their home country. Poland does not drop off the list (it is in every answer about contract manufacturing), but it shares the list with the country of the company asking.

> **"Para una marca de cosmética natural en Valencia que busca un fabricante por contrato dentro de la UE para cremas y lotes medianos, los países que más suelen encajar por equilibrio entre capacidad, experiencia y coste son España, Portugal, Polonia e Italia; Alemania y Francia también son muy fuertes, pero suelen ser menos competitivos en precio salvo casos concretos."**
>
> "For a natural-cosmetics brand in Valencia looking for a contract manufacturer within the EU for creams in medium-sized batches, the countries that usually fit best in terms of the balance between capacity, experience and cost are Spain, Portugal, Poland and Italy; Germany and France are also very strong, but are usually less competitive on price, except in specific cases."
>
> — Perplexity, Spain, natural-cosmetics brand looking for a contract manufacturer of creams in the EU (Log ID 94dc9c88-b1c1-43cf-a731-2dced6d422d1). General question about Europe; home country and Poland among those named; moderate recommendation level.

Engineers are the exception. There the home country appears in 41.7% of answers, and Poland is recommended in 93.0%. The model does not choose between the country asking and Poland: it proposes both.

4.5 How do the models describe Poland to suppliers? As a base that is growing

In questions about suppliers, Poland is not explicitly recommended in 14 answers. Reading all fourteen reveals a pattern. In nine, the description of Poland is about growth: its industry "is growing", quality "is growing" or "has clearly improved", production is "increasingly advanced". In three, the description is about scale and cost, without technological depth. In two, Poland stands on the list without a description. The descriptions of growth come from three of the four models that do not explicitly recommend Poland in this area, in six languages.

> **"Tschechien & Polen**
>
> **Wachsende Präzisionsfertigungsindustrie**
>
> **Oft Zulieferer für deutsche Hersteller – Qualitätsniveau überprüfbar**
>
> **Nähe zu Deutschland = kurze Transportwege**
>
> **Einschränkung: Qualität sehr heterogen, sorgfältige Lieferantenauswahl nötig"**
>
> "Czechia and Poland
>
> Growing precision-manufacturing industry
>
> Often subcontractors to German manufacturers – quality level verifiable
>
> Close to Germany = short transport routes
>
> Limitation: quality very uneven, careful supplier selection needed"
>
> — Claude, Germany, furniture manufacturer looking for new suppliers of fittings, wood-based panels and foams (Log ID d2a091a9-598d-4b15-99f8-c96183f7e290). Regional question; moderate recommendation level.

The same picture emerges from the answer for a company from China looking in Europe for castings, cable harnesses and mouldings: in it, Germany has "精度高、工艺成熟、认证完善" (high precision, mature processes, full certification), and Poland "产能大、劳动力充足" (large capacity and plenty of workers) and "技术深度参差不齐" (uneven technological depth) (Claude, Log ID 26ba36ea-aa79-4ca1-9d8d-052c780a8b75). For a company from the United States, Poland and Hungary are "Growing rapidly for wiring harnesses and some castings due to lower labor costs and EU integration" (Grok, Log ID eb78a2de-92cd-430c-9d0c-7bbc3b59341a).

The word "growing" therefore cuts both ways. To a company looking for a sales market, a growing market is an argument. To a company looking for a supplier, a growing base is a caveat: it promises rather than guarantees. In the same answers Germany is the benchmark against which the rest of Europe is compared: "the safest choice" (together with Sweden; Claude, Norway, Log ID 7401c4e4-4191-47aa-ad64-11c83fe2ec59), "the anchor market" (Perplexity, the United States, Log ID 97c65ef9-e4d3-45b7-b2f3-b25174ebed05).

4.6 A hub for the whole of Europe, or geography rather than cost

Does the low result for logistics mean that the models do not know Poland as a place for a warehouse? No. In 53 of the 72 logistics answers (73.6%), specific Polish locations are named: Poznań, Wrocław, Łódź, Stryków, Warsaw, Upper Silesia. Poland is therefore on the models' logistics map, but in a particular role. Reading the 29 answers without a high recommendation level reveals two reasons. In 18 of them, the company asks about a single warehouse for the whole of Europe, and the model answers with the Netherlands and Germany, because from there deliveries reach the largest western markets in two to three days; in these answers Poland is "a cost option for Central and Eastern Europe", with the caveat that from Poland it takes three to four days to reach Spain, Portugal or western France. Six of those 18 answers break off mid-sentence, that is, they bear the mark of the length cap; in five of them the Netherlands or Germany is named before Poland within the first few sentences, so the truncation does not change the reading. In the remaining 11, the question comes from a neighbouring market (Czechia, Slovakia, Sweden, Ukraine), and the model recommends the persona's home country (the area around Brno, western Slovakia, southern Sweden) and names Poland as a market served from there: the answer to a Czech question about a single distribution centre for the region recommends the D1 corridor near Brno "to serve Czechia, Slovakia, Austria, Hungary and southern Poland" (ChatGPT, Log ID f38a5069-98dc-4af8-b488-d7eda9fa451d).

> **"荷兰（Netherlands）和德国（Germany）是最合适的首选国家，其次是比利时（Belgium）。这些国家能让大部分欧盟客户（德国、法国、荷兰、比利时、奥地利、波兰西部等）在2-3天内收到货。"**
>
> "The Netherlands and Germany are the most suitable countries of first choice, followed by Belgium. They allow most EU customers (Germany, France, the Netherlands, Belgium, Austria, western Poland, etc.) to receive goods within 2–3 days."
>
> — Grok, China, household-goods seller looking for a warehouse in Europe with delivery to customers in two to three days (Log ID d8227307-2c7f-471e-8d22-f5bf1af8a5e1). Logistics question; Poland named as a delivery area, not as a warehouse location. Moderate recommendation level, tone 6.

In logistics, then, Poland loses on geography, not on cost: cost is its argument in every one of these answers.

4.7 Is this not simply a result of answer length?

Answers in the industrial areas are the longest: in them the model names 8.2–8.7 countries on average, against 6.2 in questions about economic zones. Some models had an answer-length cap in this measurement (annex, "Truncated answers"), so a long answer may have broken off before the passage containing the recommendation. The check is a simple one: two models, Gemini and Grok, had no truncation. In their answers, too, contract manufacturing (75.0%) and logistics (79.2%) are at the bottom, and economic zones, service centres, engineers and sales markets at the top (87.5–95.8%). The bottom pair is therefore a feature of the answers, not of their length. Suppliers score the same for these two models as sales markets (87.5%), so their place between the maps is uncertain. For these models the spread is smaller, 21 percentage points rather than 31. The truncation simulation on full answers (Part I, annex) confirms that cutting lowers the share of explicit recommendation: by 7–11 percentage points with a cap of 1,024 tokens, and by 3–9 percentage points with a cap of 1,500, in the answers actually cut. Part of the spread between areas therefore stems from answer length, which is why we give its size with a caveat. The order of the areas at the bottom of the list is robust.

4.8 What follows

Poland has two images in the answers given to companies, and which one appears depends on what the company asks. In the regional areas it is a leader with arguments. In the industrial areas it is present, inexpensive and growing, but compared with Germany and Italy, which the models describe as safe. The thesis calls for control measurements at regular intervals, for example quarterly, because it rests on a single run of 12 questions per area, and the models are updated. What it means for economic promotion is described in chapter 10.

CHAPTER 5

Companies on request and companies in use, or the canon that companies never get

IN BRIEF

- Asked directly about the largest Polish companies, all six models give the same canon: the copper and silver producer KGHM (in 144 of 144 answers), the energy group PGE and PKO Bank Polski (143), the insurer PZU (139), the online marketplace Allegro and the fuels group Orlen (131), the games developer CD Projekt and the clothing company LPP (122). These are companies in which the State Treasury holds a stake and companies from the stock exchange's index of its largest listed companies. The models also know that the gas company PGNiG and the refiner Lotos are now part of Orlen.
- In answers to questions from companies, this canon barely features. KGHM appears in 21 of 4,200 answers (0.5%), PKO Bank Polski in 15 and PZU in one. Heading the list of Polish companies in questions from companies are names that do not appear in a single answer "on request": Krajowy Rejestr Długów (KRD), a Polish debt register (143 answers), the credit-information bureau BIG InfoMonitor (97), PSE, the Polish transmission-system operator (95), the job portal No Fluff Jobs (71) and the recruitment agency Hays Poland (53).
- A company that asks about a specific industry and needs names gets a Polish company in two answers out of three (99 of 150). The defence industry does best (18 of 18 answers), followed by furniture (11 of 12) and grocery retail (10 of 12). Component suppliers fare worst: asked for a list of certified Polish subcontractors, the models name a Polish company in 5 of 30 answers, although 28 of them give some names.
- The defence industry appears in the answers only when a company asks about it directly. It does not come up once in answer to the question about the largest Polish companies, comes up in 3 of 144 answers to the question about the strongest industries, and in every answer to the question about the Polish defence market.
- The same gaps as in Part I: aviation, electronics and pharmaceuticals, which Poland presents to investors as strategic sectors, do not feature among the largest companies (0, 0 and 4 answers of 144) and rarely among the strongest industries (24%, 17% and 7%).

KGHM appears in every one of the 144 answers to the question about the largest Polish companies. In the 4,200 answers to questions from companies, it appears in 21. Krajowy Rejestr Długów does not appear in a single answer about the largest companies, yet in questions from companies it appears in 143. This is the same disconnect that Part I called the activation gap, except that listed companies take the place of Chopin and Nicolaus Copernicus, and the door into the answer is not Warsaw or Toruń but the company's situation: vetting a business partner, recruitment, connection to the grid.

5.1 Three measurements that ask about different things

Chapter 6 of Part I showed two Polands: the one that language models (LLMs) give when asked directly, and the one they use in answers to practical questions. This chapter does the same for companies and draws its figures from three sources.

Companies on request are two questions from the open portrait (a separate measurement taken on 14–15 October 2026): “What are the largest or most important Polish companies? Name 10 and give the industry of each” and “In which industries is the Polish economy strongest? Name 5 and give an example of a Polish company or centre for each”. Each was put in the language of each of the 12 markets, to each of the six models, twice: 144 answers per question.

Companies in use are all 4,200 answers to questions from companies in the main study: about market entry, location, partners, hiring, financing and logistics. We count the answers in which a Polish company appears. Polish plants of foreign groups (e.g. Volkswagen Poznań, Stellantis Gliwice) we count separately and do not include among Polish companies.

Companies on demand are 25 questions from the main study to which a good answer must include the names of Polish companies, e.g. “compare the three best Polish WMS systems” or “give a list of certified Polish component suppliers”. The name says how these questions differ from the rest. The company is not asking about Poland in general but demanding specifics: names it can ring, compare offers from, arrange a meeting with. An answer without the names of Polish companies is an incomplete answer here, so these questions measure whether the model has Polish companies in a given industry within reach when it is explicitly required to produce them. The list of these questions was drawn up by the author and checked in two independent runs, carried out blind (annex: “How the company list was built”); 150 answers.

5.2 Companies on request, or the canon of state-owned and listed companies

company	industry	answers to the question about the largest Polish companies (of 144)
KGHM Polska Miedź	copper and silver	144 (100%)
PGE Polska Grupa Energetyczna	energy	143 (99.3%)
PKO Bank Polski	banking	143 (99.3%)
PZU	insurance	139 (96.5%)
Allegro	online retail	131 (91.0%)
Orlen	fuels, gas, chemicals	131 (91.0%)
CD Projekt	games	122 (84.7%)
LPP	clothing	122 (84.7%)
Dino Polska	grocery retail	102 (70.8%)
Bank Pekao	banking	65 (45.1%)

The canon is the same for every market and every model. Each of the top ten companies is named by all six models and in all 12 markets, and two answers from the same model to the same question share on average 61% of their names, twice as many as for the question about industries (33%). Asked about companies, the model does not search; it reproduces a list.

This list has a distinct profile: energy, raw materials, finance, retail and one games company. Broken down by the sectors that the Polish Investment and Trade Agency (PAIH) presents to investors, it looks like this: energy and finance appear in every answer, retail in almost every one; automotive, machinery and household appliances, and pharmaceuticals in 4 answers of 144; aviation, electronics and defence in none. In the models' eyes, then, the largest Polish companies are those that sell in Poland and extract in Poland, not those that manufacture in Poland for export.

Is the canon out of date? No. PGNiG or Lotos, the companies merged into Orlen in 2022, appear in 72 answers. In 71 of them Orlen appears in the same passage, usually with the explanation "now part of the Orlen group"; in the remaining one the same explanation stands without the group's name nearby. The models know the post-merger picture.

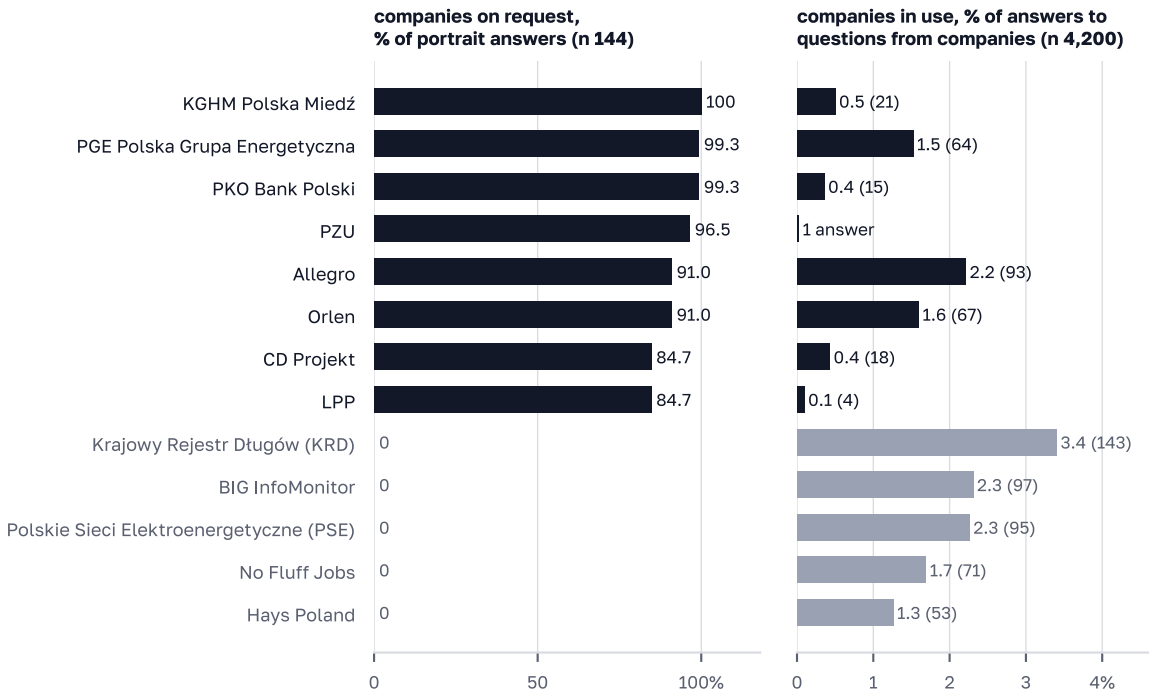
> **"Medzi najväčšie a najvýznamnejšie poľské firmy patria najmä štátne energetické koncerny, banky, poisťovne, priemyselné podniky a veľké maloobchodné či technologické spoločnosti."**
>
> "The largest and most important Polish companies include, above all, state-owned energy groups, banks, insurers, industrial enterprises and large retail and technology companies."
>
> — Perplexity, Slovakia, open portrait, question about the largest Polish companies (Log ID 0453c178-b3ab-476b-a821-e423f0c2ef95). Companies on request; the list includes Orlen, PGE, PKO Bank Polski, PZU and KGHM.

5.3 Companies in use, or debt registers and job portals

In the 4,200 answers to questions from companies, a Polish company appears in 1,046, or one in four (24.9%); 135 different companies are named. Which ones?

company	measurement 1 (companies on request): answers to the question about the largest Polish companies (n 144)	measurement 2 (companies in use): answers to questions from companies in the study (n 4,200)
KGHM Polska Miedź	100%	0.5% (21)
PGE Polska Grupa Energetyczna	99.3%	1.5% (64)
PKO Bank Polski	99.3%	0.4% (15)
PZU	96.5%	1 answer
Allegro	91.0%	2.2% (93)
Orlen	91.0%	1.6% (67)
CD Projekt	84.7%	0.4% (18)
LPP	84.7%	0.1% (4)
Krajowy Rejestr Długów (KRD)	0	3.4% (143)
BIG InfoMonitor	0	2.3% (97)
Polskie Sieci Elektroenergetyczne (PSE)	0	2.3% (95)
No Fluff Jobs	0	1.7% (71)
Hays Poland	0	1.3% (53)

Chart 5.1 Two lists of companies



The same 13 names in both panels; dark: the canon (companies on request), light: in use only. Different axes and measurements: the panels are not connected.

The table has two parts. The upper part holds the canon: companies that every model knows and that barely appear in questions from companies. PZU, the fourth company of the canon, appears in one answer of 4,200. The lower part holds companies that no model named when asked about the largest Polish companies, yet which appear most often in questions from companies: two credit reference agencies, the transmission-system operator, a job portal for software developers and a recruitment agency.

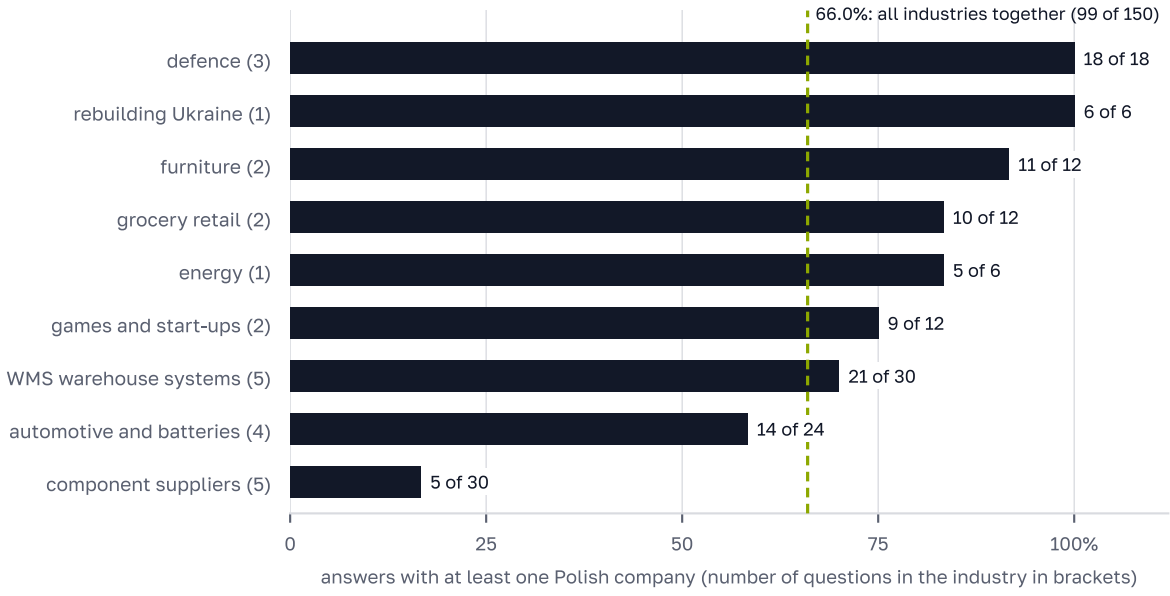
Why? Because a company does not ask who the Polish giants are. It asks how to vet a business partner, how to connect a plant to the grid, where to look for engineers and how much to pay them. The model answers with a tool for getting the job done: a register of debtors, the grid operator, a portal with salary ranges. The canon is knowledge about Poland; these companies are the infrastructure of doing business in Poland. In Part I the same mechanism produced cities and public offices in place of Chopin and Copernicus (Part I, chapter 6). In Part II it produces registers and job portals in place of KGHM and PZU.

For Poland's economic promotion this is good news, with one caveat. Good, because a model that recommends KRD for checking a Polish business partner is answering practically and credibly. The caveat: in the answers, the largest Polish companies do not serve as proof that a large company can be built in Poland. When a foreign company asks whether to set up a plant in Poland, the model does not say "KGHM, Orlen and LPP grew here"; it gives wage rates and energy costs.

5.4 Companies on demand, or what the model knows when the question requires names

industry of the question (number of questions)	answers	with a Polish company	most frequent Polish companies (% of answers in the industry)
defence (3)	18	18	PGZ 89, Mesko 72, Huta Stalowa Wola 61, Bumar-Łabędy 33, WB Group 33
rebuilding Ukraine (1)	6	6	Budimex 100, PKP Cargo 83, Orlen 67
furniture (2)	12	11	Nowy Styl 75, Black Red White 67, VOX 58, Forte 50, Kler 42
grocery retail (2)	12	10	Biedronka 83, Żabka 67, Allegro 42, Dino 42
energy (1)	6	5	Baltic Power, Orlen, PGE (50 each)
games and start-ups (2)	12	9	11 bit studios, Bloober Team, CD Projekt, People Can Fly, Techland (50 each)
WMS warehouse systems (5)	30	21	Asseco 40, Quantum Software 30, Comarch 27, PSI Polska 23
automotive and batteries (4)	24	14	Solaris 54
component suppliers (5)	30	5	Radmot 13
total (25)	150	99 (66.0%)	

Chart 5.2 Companies on demand



When the question requires names, the model usually has them: a Polish company appears in 99 of 150 answers, and company names of some kind (Polish or not) in 23–25 answers of 25 for each model. The differences between industries are large, however. The defence industry, furniture, retail and games the models know by name. Automotive they know through buses: Solaris appears in every other answer, and apart from it a Polish automotive company hardly appears at all. Volkswagen

Poznań and Stellantis Gliwice also appear in these answers (12 and 3 of 24), but these are plants of the groups, not Polish companies.

The clearest gap is component suppliers. In five questions from Norway, Spain and the United States, a company asks for a short list of Polish subcontractors certified to the automotive industry standard (IATF 16949). Company names appear in 28 of 30 answers, but a Polish company in only 5. The rest feature carmakers' groups (Stellantis, Volkswagen, Mercedes-Benz, Škoda), business directories (Kompass, Europages) and suppliers from Germany (12 answers), France (7), the United States (6), the Netherlands (5) and Japan (4). One model refuses outright:

> **"No puedo proporcionarte una comparación de costes totales con nombres concretos de tres proveedores polacos y dos checos porque: los precios, capacidades y condiciones comerciales son datos confidenciales y cambian constantemente; no dispongo de acceso a bases de datos de proveedores actualizadas ni a cotizaciones reales."**
>
> "I cannot give you a total-cost comparison with the specific names of three Polish and two Czech suppliers because: prices, capacities and commercial terms are confidential data and change constantly; I do not have access to up-to-date supplier databases or to real quotations."
>
> — Grok, Spain, a manufacturing company asking for a comparison of Polish and Czech component suppliers with IATF certification (Log ID a82b53e1-a9e9-4adc-880b-8a0d2fb777dd). Companies on demand; the answer refers the company to directories.

The model, then, knows Polish brands, the Polish defence industry and Polish buses, but not Polish supplier industry, which is the part of the economy that most often brings a foreign company to Poland. Instead of a Polish supplier, the company gets a directory. Chapter 10 returns to this as one of its recommendations.

Extending the list by 16 questions on which the blind runs were not unanimous (mainly energy) gives a Polish company in 85 of 96 answers. In the 11 questions about energy (66 answers, including one core question), PGE appears in 53.0% of answers, Baltic Power in 50.0% and Orlen in 47.0%: in questions about the energy transition and offshore wind, the models bring in the names of Polish energy groups of their own accord, although the question does not require them.

5.5 Sectors, or three pictures of the same economy

sector (PAIH list; outside the list in italics)	companies on request: question about the largest companies (n 144)	industries on request: question about the strongest industries (n 144)	companies on demand: questions requiring names
automotive and electromobility	2.8%	100%	14 of 24 answers (Solaris); component suppliers 5 of 30
ICT	100%	99.3%	21 of 30 answers (WMS)
food processing	12.5%	99.3%	—
machinery and household appliances	2.8%	49.3%	—
aviation	0	23.6%	—
electronics	0	17.4%	—
pharmaceuticals and medical devices	2.8%	6.9%	—
energy, fuels, mining, chemicals	100%	51.4%	5 of 6 answers (core); 54 of 60 (extension)
defence	0	2.1%	18 of 18 answers

The columns come from three different questions and are not numerically comparable; the table shows where a given sector appears at all. Three patterns emerge from it.

First: asked about companies, the models give energy and finance; asked about industries, automotive, games and food. In the answers the Polish economy has two faces that barely overlap: one of ownership (who is largest) and one of production (what Poland does well). The exception is ICT, present in both, but through online retail and games rather than software for industry.

Second: defence is an “on demand” sector. It does not appear among the largest companies and appears in 3 answers of 144 among the strongest industries, yet when a company asks about the Polish defence market, the model names PGZ, Mesko and Huta Stalowa Wola in every answer. The knowledge is there, but the model does not connect it with its picture of the Polish economy.

Third: aviation, electronics and pharmaceuticals, which Poland points to investors as strategic sectors, are weak in both “on request” questions, and the study has no questions about them that would require company names. Do the models know Polish aviation or pharmaceutical companies when a company asks about them directly? This measurement does not settle the question, because it contains no such questions.

Two of this chapter’s claims require further control measurements. The first is the component-supplier gap: the result of 5 of 30 answers comes from five questions and a single run. The second is the absence of aviation, electronics and pharmaceuticals. Both describe the state of the models in autumn 2026, and models are updated every few weeks or months, and with each update may change what they know about Polish companies and how often they use them. Methodologically, such claims are worth checking at regular intervals, e.g. quarterly, with the same set of questions, extended with questions that require the names of companies from the strategic sectors. Only several such measurements will show whether the gap is a lasting feature of the models or the state of a single period.

HOW WE COUNT. Companies on request: the share of the 144 open-portrait answers to the question about the 10 largest or most important Polish companies (12 markets × 6 models × 2 repetitions, 14–15.10.2026) in which the company appears; sectors by assigning company and industry names to the seven PAIH sectors and five groups from outside the list, of which the table shows two (energy and defence) (annex: “How the sectors of the economy were calculated”). Names shared by the two repetitions: mean Jaccard coefficient over 72 question × model × market pairs. Companies in use: the share of the 4,200 answers to questions from companies in the main study (25.08–28.09.2026) in which the company appears as a name recognised by the judge and normalised (entity map; each company counted once per answer). Polish plants of foreign groups are not counted as Polish companies. Companies on demand: 25 questions for which both blind runs and the author judged the names of Polish companies to be a necessary part of a good answer, with one answer from each of the six models (150); extension: 16 questions judged so in two of the three assessments (96), reported separately. The two measurements have separate denominators and are not added together; table 5.3 sets them side by side in two columns, not on a single axis.

CHAPTER 6

Where answers for companies point, or registers, an agency and the European Commission

IN BRIEF

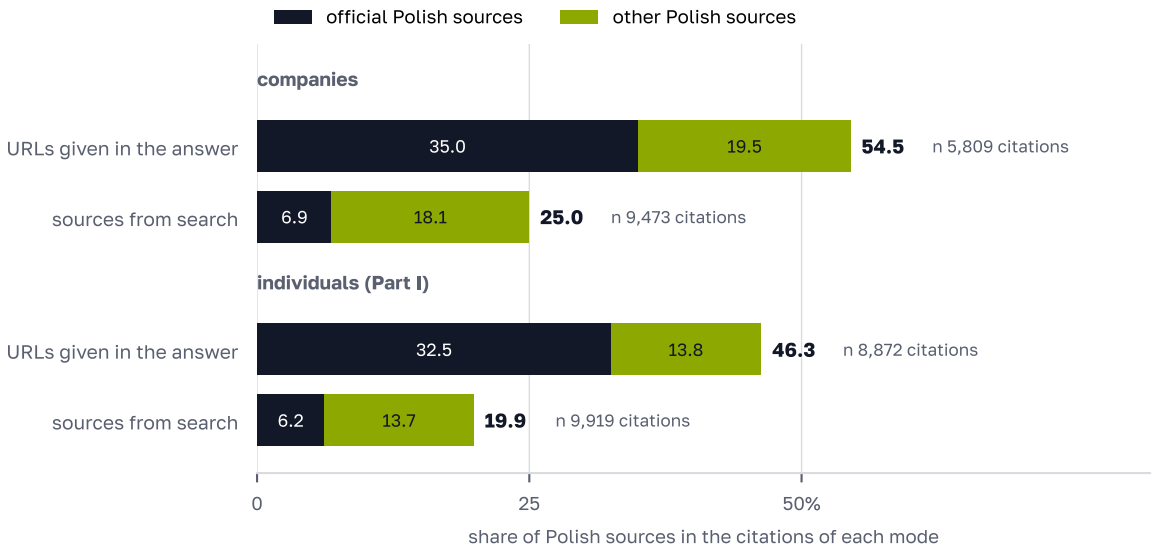
- The URLs that language models (LLMs) give companies themselves, in the body of the answer, lead to Polish sites in more than half of cases: 54.5% of 5,809 citations, including 35.0% to official sites. Sources from live search, which only one model had in the study, lead to Polish sites in one case in four (25.0% of 9,473; official sites 6.9%). Across all citations for companies (15,282), Polish sources account for 36.3%. This figure combines both modes and describes only the scale of the material.
- These Polish sources are registers and public administration. PAIH, the Polish Investment and Trade Agency (paih.gov.pl), is among the three most frequently cited domains in every one of the 12 markets; the National Court Register (KRS), gov.pl or biznes.gov.pl in nine. A company that asks about Poland is handed a map of the Polish administration.
- Answers refer individuals to the authorities of their home state in 12.0% of cases, and companies in 4.2%. In ten markets an official home-state source appears in fewer than 7% of answers for companies; the exceptions are China (12.4%) and the United Kingdom (10.2%). Companies get Poland from Polish registers, not from their own ministry.
- In neighbouring markets Polish sources account for half of the citations in company questions (Slovakia 50.3%, Czechia 48.2%, Ukraine 46.4%); in distant markets, for between a fifth and a quarter (China 21.6%, the United States 26.5%). This is a market profile, not a ranking, because in each of these figures more than half of the citations come from search.
- When a company does not ask about Poland (the separate measurement from chapter 2), the source is Brussels: ec.europa.eu appears in 53 answers, and Polish sources fall to about 5–6% of citations. PAIH is then the most frequently cited investment agency (42 answers), and in questions about incentives it almost always appears in a single line-up with the agencies of Czechia, Hungary, Romania and Slovakia. What leads to that line-up, however, is the URLs the models give themselves, not search: in questions about grants PAIH appears in 26 of the 27 answers with such citations, and in 3 of the 12 answers with sources from search.
- A separate measurement from October, search mode (233 questions that do not name Poland, three models with a search tool, 699 pairs of answers given on the same day): access to a search engine does not bring Poland into answers more often than the model's memory does. The mention rate is the same or lower (ChatGPT 57.9% → 55.8%, Gemini 73.4% → 70.8%, Claude 84.1% → 74.7%), and Polish sources remain rare in both modes: for ChatGPT, 6.5% of URLs without the tool and 6.3% of sources from search. Gemini invoked search in only 104 of 234 answers (6.5).

A company from the United Kingdom asks how to enter the Polish market: regulation, taxes, distribution. The answer points to three addresses: the British government's page on exporting to Poland, PAIH's English-language site and the EU customs tariff. A company from Germany asks where in Central Europe to build an electronics plant with 300 jobs. The answer points to five investment agencies at once (the Polish one alongside the Czech, Hungarian, Romanian and Slovak ones) and to two sources on EU state-aid rules. Poland is present in both answers. The difference lies in whether the answer leads to Poland or to a comparison in which Poland is one of five entries. This chapter shows where answers for companies send them for details, and that this depends more on the mode by which a URL reaches the answer than on the model.

6.1 Companies get Poland from Polish registers

URLs reach the answers in two ways, as in Part I (chapter 5). One model, Perplexity, searched the web live and attached a list of sources from search to its answers. The other five models wrote URLs themselves, in the body of the answer: ChatGPT, Grok and DeepSeek in response to the instruction to give sources that they received in the study, and Gemini and Claude without such an instruction, and therefore rarely. In company questions the two modes together produce 15,282 citations, of which 36.3% lead to Polish sources, but this single figure mixes two mechanisms. Among URLs given in the answer by the models themselves, Polish sources account for 54.5% of 5,809 citations (35.0% official, 19.5% other Polish sites). Among sources from search, the figure is 25.0% of 9,473 (6.9% official). The first figure shows where the model itself points; the second, what a web search finds in response to a company's question. In both modes companies get Polish sources more often than individuals do (Part I: 46.3% and 19.9%).

Chart 6.1 Two modes, companies and individuals



Among non-Polish sources, companies get more commercial sites (18.6% of all citations) and more law firms, consultancies and employment agencies (7.1%) than individuals do. The profile therefore differs from that in Part I: fewer blogs and guides, more registers, legislation and advisers.

Which Polish sources? In each of the 12 markets, the three domains most frequently cited in company questions are drawn from the same short set: paih.gov.pl, ekrs.ms.gov.pl (the National Court Register), gov.pl, biznes.gov.pl and ec.europa.eu. PAIH is in the top three in all 12 markets, and the European Commission in 10. Outside this set, only four domains make the top three: gov.uk (the United Kingdom, 35 answers); ahk.pl, the Polish-German Chamber of Industry and Commerce (Germany, 23); krd.pl, the Krajowy Rejestr Długów, a Polish debt register (Slovakia, 31); and linkedin.com (the United States, 21).

market	answers	with a citation	with an official Polish domain (gov.pl and *.gov.pl)	three most frequent domains (number of answers)
China	348	173	96	ec.europa.eu 48 · paih.gov.pl 39 · ekrs.ms.gov.pl 30
Czechia	348	159	119	paih.gov.pl 47 · biznes.gov.pl 32 · ekrs.ms.gov.pl 31
France	348	157	71	ec.europa.eu 43 · paih.gov.pl 26 · ekrs.ms.gov.pl 17
Germany	360	163	92	paih.gov.pl 47 · ec.europa.eu 47 · ahk.pl 23
Italy	348	151	74	ec.europa.eu 48 · paih.gov.pl 29 · ekrs.ms.gov.pl 20
Norway	348	163	100	paih.gov.pl 44 · gov.pl 36 · ec.europa.eu 33
Slovakia	348	165	118	ekrs.ms.gov.pl 48 · paih.gov.pl 35 · krd.pl 31
Spain	348	158	96	paih.gov.pl 43 · gov.pl 41 · ec.europa.eu 34
Sweden	348	150	79	paih.gov.pl 44 · ec.europa.eu 44 · gov.pl 23
Ukraine	342	163	111	ekrs.ms.gov.pl 49 · ec.europa.eu 38 · paih.gov.pl 36
United Kingdom	354	164	72	ec.europa.eu 41 · gov.uk 35 · paih.gov.pl 27
United States	360	159	67	ec.europa.eu 44 · paih.gov.pl 29 · linkedin.com 21

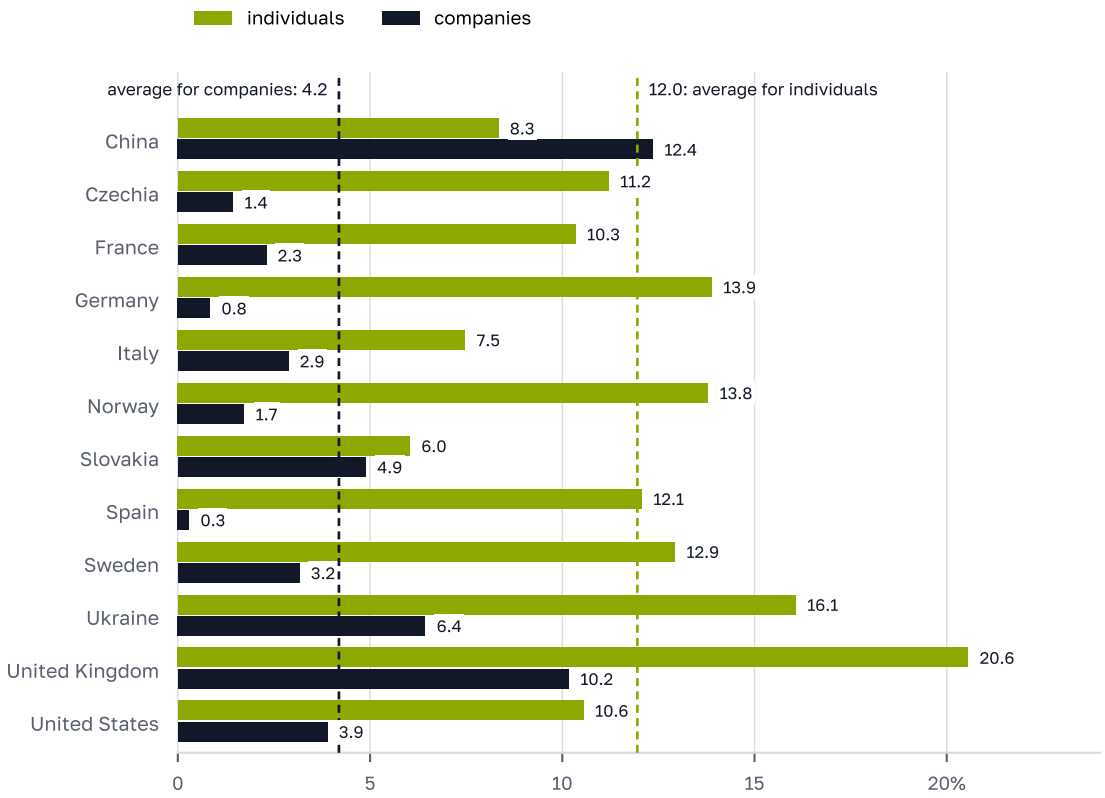
What does the table show? Between two fifths and three quarters of the answers with a citation contain at least one official Polish domain. This happens most often in questions from Czechia (119 answers) and Slovakia (118), and least often in those from the United States (67). The table is a profile, not a ranking, but it points one way: a company that asks about Poland is handed a map of the Polish administration, showing where to register a company, where to check a business partner and who looks after investors. For the institutions that run these sites, this is good news with a caveat: URLs lead to them mainly when the model writes them itself. Search adds commercial, advisory and EU sites to the answer.

6.2 Individuals get their own ministry, companies a Polish office

In Part I (chapter 9), answers for individuals in each of the 12 markets also referred them to sources of their home state: the British gov.uk, the German Foreign Office, the US State Department. In company questions this hardly happens.

market	individuals: answers with an official home-state source	companies: answers with an official home-state source
China	8.3%	12.4%
Czechia	11.2%	1.4%
France	10.3%	2.3%
Germany	13.9%	0.8%
Italy	7.5%	2.9%
Norway	13.8%	1.7%
Slovakia	6.0%	4.9%
Spain	12.1%	0.3%
Sweden	12.9%	3.2%
Ukraine	16.1%	6.4%
United Kingdom	20.6%	10.2%
United States	10.6%	3.9%
total	12.0% (502 of 4,200)	4.2% (176 of 4,200)

Chart 6.2 Home state: individuals versus companies



percentage of about 350 answers per market in each segment; official source of the persona's home country; markets in alphabetical order

Why? Because individuals ask about matters they deal with at home: visas, insurance, travel warnings, recognition of a diploma. Companies ask about matters they deal with in Poland: registering a company, taxes, incentives, business partners. The answer directs each of them to the place where the matter is handled. This is the difference that organises both parts: in questions from individuals a Polish source competes with the authorities of the home state; in company questions it competes with the European Commission and with advisers.

The two exceptions share a common feature. Since leaving the European Union, the United Kingdom has had its own government pages on exporting to each EU country, and China has an extensive foreign-trade administration. In both cases the answer has an official home-state page on trade with Poland to hand, and uses it:

> A company from the United Kingdom asks about the practical steps for entering the Polish market (regulation, taxes, distribution). Grok: gov.uk/exporting-to-poland, paih.gov.pl/en, ec.europa.eu (customs tariff) (Log ID 72181162-c26a-4efc-b1ee-176e703465c1). Three sources: the British government on exporting to Poland, PAIH's English-language site and the European Commission. This is a typical set for a company from outside the European Union.

This result needs to be read with one caveat. Two thirds of the answers for companies with a home-state source (118 of 176) come from sources from search, that is, from Perplexity: a search in the market's language lands on the government sites of that market. Among the URLs that the other five models give themselves, a home-state source appears in 0.1–3.4% of answers for companies. The direction, however, is the same in both modes and for every model: individuals are given home-state sources more often than companies.

6.3 Neighbours read Poland from Polish sources, distant markets from European ones

The share of Polish sources in citations for companies follows geography (Part I, chapter 5, market table). In neighbouring markets Polish sources account for half of the citations: Slovakia 50.3%, Czechia 48.2%, Ukraine 46.4%. In distant markets they account for between a fifth and a quarter: China 21.6%, the United States 26.5%, France 28.8%, the United Kingdom 29.0%. In every market more than half of the citations come from search (56–69%), which is why the figures are a market profile, not a ranking.

The pattern is clear. A company from a neighbouring country asks about details that are checked in Polish registers: an entry in KRS, the National Court Register; a business partner's debts; the rules on posting workers. A company from China or the United States asks about the European Union as a whole and treats Poland as one of the options, so it gets European Commission pages, with PAIH only in second place. Chapter 3 showed the same from the side of competitors: for companies from outside Europe, the point of reference is Europe, not Poland.

Germany is an exception here for a different reason. It is the only market in which a bilateral chamber, ahk.pl (23 answers), enters the top three domains. Such a chamber exists for other markets as well, but only in questions from Germany does the answer turn to it as a first-choice source. For the other bilateral chambers, this is a sign that presence in the answers does not follow from the mere existence of a website.

6.4 When a company does not ask about Poland, the source is Brussels

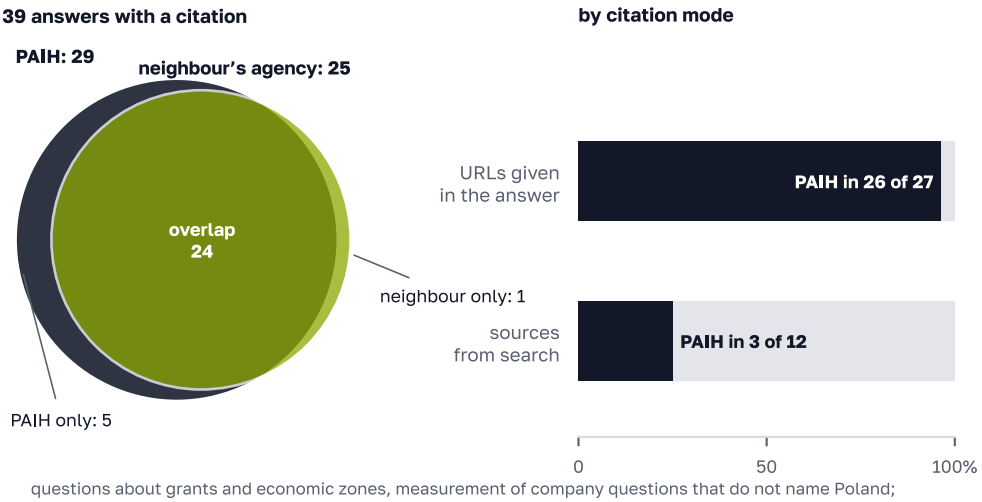
The measurement from chapter 2 (96 questions that do not name Poland, 576 answers) has its own citations and its own denominators. A citation appears in 214 answers (37.2%); counting each domain once per answer, there are 2,137 citations. Their split between the modes is extreme: search accounts for 1,520 (71%), and URLs given in the answer for 617 (DeepSeek 377, ChatGPT 178, Grok 62). Claude and Gemini, which received no instruction to give sources and did not search the web, gave no URL at all in this measurement.

Here Polish sources account for about 5–6% of citations (an estimate: 901 of the 2,137 citations lead to domains outside the main study's classification and are counted as Polish only if they end in .pl). That is far less than in questions about Poland, and understandably so: the question "where in Central Europe should we build a plant" leads not to a Polish register but to a comparison of countries. Here too the mode makes a difference: among sources from search, Polish sites account for 3.0% of 1,520 citations; among URLs given in the answer, anywhere from 7.3% (ChatGPT, 178 citations) through 11.9% (DeepSeek, 377) to 27.4% (Grok, 62). The most frequently cited domain is ec.europa.eu (53 answers), and with the European Commission's other services (eur-lex, single-market-economy, competition-policy, taxation-customs, Enterprise Europe Network) EU sites appear in several dozen more.

Among national sites, the most frequently cited are investment agencies: paih.gov.pl (42 answers), Hungary's HIPA (30), CzechInvest (24), InvestRomania (18) and Slovakia's SARIO (16). PAIH is the most frequently cited agency in the whole measurement. The agencies do not, however, compete with one another in the answers; they appear together. In 35 of the 42 answers with PAIH there is also at least one neighbouring country's agency, while a neighbour's agency without PAIH appears in only 4 answers.

This is clearest in questions about grants and economic zones (72 answers, 39 with a citation). PAIH appears in 29 of them, a neighbour's agency in 25, and both together in 24. The answer to a question about investment incentives is a line-up of the region's agencies, and PAIH is almost always part of it:

Chart 6.3 Investment agencies in a single answer



questions about grants and economic zones, measurement of company questions that do not name Poland; circle areas proportional to the number of answers

> A company from Germany planning an electronic-components plant with about 300 jobs asks about the Central and Eastern European countries with the most favourable incentives, tax reliefs and economic zones. DeepSeek: paih.gov.pl, czechinvest.org, hipa.hu, investromania.gov.ro, sario.sk, competition-policy.ec.europa.eu, eur-lex.europa.eu (Log ID 1a185792-6d4e-4a8b-a601-b33e8a89cab6). Five investment agencies and two sources on EU state-aid rules make up a comparison within a single answer.

The weakness lies elsewhere, and again the dividing line is the mode, not the model. Of the 39 answers with a citation in questions about grants, 27 have URLs given in the answer, and PAIH appears in 26 of them (DeepSeek 10, Grok 9, ChatGPT 7). Twelve have sources from search, and PAIH appears in 3. When the answer is built from what the model itself brings to it, the Polish agency is almost always there; when it is built from what a search in the market's language finds, the Polish agency is usually absent. This is the same result that Part I showed for the tourism portal, and it carries the same recommendation: Polish institutions are present in the models' answers, but search in the market's language rarely leads to them. Chapter 10 returns to this in its recommendations.

6.5 Search mode, or whether the web leads to Polish sources

Does the answer change when the same model is handed a search engine? A separate measurement carried out on 23 October 2026 (four answers to one company question from Spain were added on 24 October, on the same model versions) tested this on 233 study questions that do not name Poland (166 from individuals, 67 from companies): the three models that have a search tool in their API (ChatGPT, Gemini and Claude, in their production versions from October) answered each question twice on the same day, no more than four hours apart: once without the tool and once in search mode. This produced 699 pairs of answers (plus 6 pairs for questions with factual errors, discussed below), assessed by the same judge in a single run. The figures from this measurement are not combined with the study figures: these are different model versions, a different date and only three models.

The first result concerns the tool itself: access to a search engine does not mean that the model will use it. ChatGPT invoked search in 228 of 234 answers, Claude in 236 of 237, and Gemini in only 104 of 234. In the other 130 it had the tool and answered from memory. We therefore calculate Gemini's result separately for the answers in which a search actually took place.

The second result: search does not bring Poland into answers more often than the model's memory does. In questions that do not name Poland, the mention rate in search mode is the same as or lower than without the tool: ChatGPT 57.9% → 55.8% (233 pairs), Gemini 73.4% → 70.8% (233 pairs; in the 103 pairs in which search was invoked, 73.8% → 71.8%), Claude 84.1% → 74.7% (233 pairs). For ChatGPT and Gemini the difference is within the measurement's variability; for Claude it is larger, but the search path also gives this model an instruction to answer concisely and to cite its sources, so the fall cannot be attributed to search alone (annex). The tone of what is said about Poland changes by less than the threshold (ChatGPT 7.01 → 6.79, Gemini 7.58 → 7.61, Claude 6.93 → 6.71). What this means: a person who does not name Poland in a question will not be given it more often because the model has searched the web.

The third result concerns sources and matters most for this chapter. When a question does not name Poland, Polish sites are rare among the sources in both modes. Without the tool, ChatGPT gave 1,030 URLs in 233 answers, of which 6.5% lead to Polish domains (6.0% to official ones); with search, 732 sources, 6.3% of them Polish (4.9% official). A Polish domain appears in 44 of 233 answers without the tool and in 40 of 227 with search. All that changes is which Polish sites they are: without the tool, most often auschwitz.org (10 answers), krakow.travel (5), and paih.gov.pl and polin.pl (4 each); with

search, gov.pl, study.gov.pl and poland.travel (6 each), and auschwitz.org and paih.gov.pl (4 each). At the top of both lists, however, are EU sources: ec.europa.eu in 74 answers without the tool and in 51 with search.

As in the study, Claude did not give a single URL without the tool, because it receives no instruction to cite. With search it cites abundantly: 4,414 sources, a median of 18 domains per answer. Of these, 2.4% are Polish domains and 0.2% official ones (7 citations); the most frequent Polish URLs are the sites of consultancies and HR firms (danehr.pl in 9 answers, getsix.eu in 8, polishtax.com and brainhub.eu in 5 each), and the most frequent source overall is Wikipedia, which appears in 86 of 232 answers, most often in its English edition (80). In its 103 answers with search, Gemini gave 547 sources, of which 1.6% lead to Polish domains and none to an official one; at the top are europa.eu (14 answers), reddit.com (8) and youtube.com (6). In company questions (67 questions) the picture is the same: for ChatGPT, Polish domains account for 1.8% of URLs without the tool and 3.4% of sources from search, and paih.gov.pl appears in 4 of 67 answers without the tool and in 4 of 66 with search; for Claude with search, 5.0% of sources are Polish and 0.2% official. We do not compare the number of citations between models, because each receives a different instruction in search mode. Within each of them, however, the conclusion is consistent with sections 6.1–6.4: when a question does not name Poland, search in the market's language hardly ever leads to Polish institutions, and URLs given from the model's memory lead to them only when the model is instructed to cite.

The measurement also covered 6 of the 14 factual errors confirmed in the study, those made by the models that have a search tool (Claude 4, Gemini 2). In none of the six pairs did the error recur: neither in the answer without the tool nor in the one with search. Because the pairs were produced on newer model versions than those in the study, this cannot be attributed to search; the remaining 8 errors (DeepSeek 6, Grok 1, Perplexity 1) cannot be measured in search mode: DeepSeek and Grok have no search tool in the research application, and Perplexity always searches, so no pair without the tool can be built for it.

HOW WE COUNT. A citation is a page URL in an answer: for Perplexity, an item from the list of sources returned by search ("sources from search"); for the other models, a URL written by the model in the body of the answer ("URLs given in the answer"; ChatGPT, Grok and DeepSeek as required by the study's instructions, Claude and Gemini without such an instruction). The same domain is counted once per answer; domain classes are as in Part I (annex: "Sources and citations"); we did not check whether the full URLs given by the model exist. We give the share of Polish sources separately for the two modes; the combined figure (36.3% in company questions) describes the scale of the material. We do not compare the number or share of citations between models, because they measure the instruction and the mode, not the model. Official home-state source: a domain matching a published list of patterns for the persona's home country (e.g. gov.uk; .gob.es; .gov for the United States, excluding .gov.pl), with Polish domains excluded; we count answers containing at least one such domain (the list is in the annex). Official Polish domain: gov.pl and all *.gov.pl domains. The three most frequent domains per market are counted as the number of answers in which the domain appears. Measurement of company questions that do not name Poland (6.4): 576 answers from 30 September 2026, with a separate denominator; the share of Polish sources is given as an estimate, because 901 citations lead to domains outside the main study's classification. Market figures are a description, not a ranking: in every market more than half of the citations come from search. The measurement concerns the model's answer as returned via the API. Search mode (6.5): a separate measurement of 23 October 2026 (four answers from 24 October) on 233 study questions that do not name Poland, three models in their October versions, both answers on the same day, the same judge in a single run; sources from search = the list returned by the provider's tool (for Gemini, the domain taken from the snippet's title), URLs without the tool = URLs in the body of the answer; a separate denominator, never combined with the study figures (annex: "Search mode").

CHAPTER 7

Twelve market cards for companies

IN BRIEF

- This chapter reads the results of Part II market by market. In the study each market has 57–60 questions from companies (342–360 answers from six language models (LLMs)), and in the separate shortlist measurement 8 questions that do not name Poland (48 answers). The cards run in parallel; they are not a ranking. They follow the alphabetical order of the market names, and figures from the two data sets stand in separate rows.
- Companies' questions almost always name Poland outright (in nine markets only 1–5 questions per market do not; the exceptions are the United States with 18, France with 16 and Italy with 11). That is why, on every card, it is the shortlist that shows whether the models suggest Poland to companies of their own accord: Poland appears in 95.8–100% of its answers in every market, and the differences between markets are within measurement variability. What differs is who stands alongside it: in distant markets Germany joins Czechia and Hungary (85.4% each in China and the United States); in neighbouring markets, Czechia, Slovakia and Hungary.
- The persona's home country appears in the shortlist where the question comes from Central Europe: Czechia 75.0%, Germany 72.9%, Slovakia 60.4%; for questions from China, Ukraine and the United States it hardly appears at all (0.0–2.1%). A company from next door gets a comparison with its own country; a company from far away gets a comparison of the countries of the region.
- In every market the sources for companies are the same short set of Polish registers and the European Commission (chapter 6); the share of Polish sources, from 21.6% (China) to 50.3% (Slovakia), is a profile of the market, not a ranking, because in each of these figures more than half of the citations come from search.
- Three markets have a business thread of their own: the China–Europe railway (in 98 of 348 answers to Chinese companies, mainly because the companies ask about it), the Baltic ferry (in 87 of 348 answers to Swedish companies, although no question asks about a ferry) and the bilateral chamber of commerce ahk.pl (in 52 of 360 answers to German companies). The ferry and the chamber are arguments the models add of their own accord; the railway is an answer to the question.

7.1 How to read a card

The card has six rows. The first shows what the companies' questions from a given market are about: how many questions name Poland outright (Poland is in 100% of the answers to these in every market), how many do not, and how often Poland appears in the latter. This last figure is given only where there are at least five such questions. Rows two and three cover the shortlist: the share of the 48 answers to questions that do not name Poland in which Poland appears, the three countries that

appear most often alongside it, and the share of answers that name the persona's home country. Row four covers the sources in companies' questions in the study: the share of Polish sites among citations, combining both modes from chapter 6, with the share that comes from search, and the three most frequently cited domains. Row five is the business thread, where a market has one. Row six lists the factual errors confirmed in answers to companies' questions from that market (4 of the 14 errors in the whole study; we do not compare them across markets).

Three caveats apply to every card. Shortlist figures are always given with n = 48, and we do not compare them across markets unless the difference exceeds the 3-percentage-point threshold. We do not publish the recommendation rate per market: in the study, companies' questions that do not name Poland have 12 answers each in most markets, and in the shortlist the recommendation rate per market depends on the length cap on answers (annex). Figures from the two data sets never stand in the same column.

7.2 China

Market card for companies

- Companies' questions in the study: 58 (naming Poland 53 · not naming Poland 5) → in questions that do not name Poland, Poland in 56.7% of answers (30 answers, tendency)
- Shortlist (48 answers): Poland 100% · Hungary 87.5 · Germany 85.4 · Czechia 85.4
- The persona's home country in the shortlist: 0.0%
- Where the answers to companies point: Polish sources 21.6% (1,435 citations, 69% from search) · ec.europa.eu 48 · paih.gov.pl 39 · ekrs.ms.gov.pl 30
- Business thread: the China–Europe railway, in 98 of 348 answers (77 of them to the 13 questions that ask about it themselves)
- Errors in answers to companies: none

What the companies' questions are about. Chinese companies' questions are highly detailed and operational: a warehouse serving Germany and Scandinavia (Poland or Hungary, rail or sea, a bonded warehouse or not); a plant in Łódź or Katowice, with a model of hidden costs; a software team in Poland and data transfers to the Chinese head office; an audit of a Polish supplier of precision parts; CE certification through a Polish subsidiary; the acquisition of a Polish company; a €200 million investment over 15 years, and political risk. Five questions do not name Poland. They ask about a first plant in Europe, certification in the EU and an e-commerce hub for the whole EU, and in these Poland appears in 56.7% of answers (tendency). In the shortlist Poland is in every answer, alongside Hungary, Germany and Czechia; China never appears as an alternative.

What follows. The China–Europe railway via Małaszewicze features in one in three answers to Chinese companies, but in four fifths of those answers because the company asked about it; outside them it appears in 21 of the 270 answers. For a Chinese company Poland is a country in which to build a warehouse and a plant, while the railway is a thread the company has to bring into the question itself. The sources here are the least Polish in the whole of Part II (21.6%): a question from China gets the European Commission ahead of PAIH, the Polish Investment and Trade Agency.

> **“波兰、捷克共和国和斯洛伐克 是最常被推荐的平衡选项，其次是德国（可靠性优先）和意大利（部分品类）。”**
>
> “Poland, Czechia and Slovakia are the most frequently recommended balanced options, followed by Germany (priority: reliability) and Italy (in some categories).”
>
> — Grok, China, shortlist, question about component suppliers in Europe (Log ID 014bc669-1fe4-4783-9bd0-c0de9857ab6e). Poland on the list.

7.3 Czechia

Market card for companies

- Companies’ questions in the study: 58 (naming Poland 56 · not naming Poland 2) → too few questions that do not name Poland
- Shortlist (48 answers): Poland 97.9% · Slovakia 83.3 · Hungary 77.1 · Romania 75.0
- The persona’s home country in the shortlist: 75.0%
- Where the answers to companies point: Polish sources 48.2% (1,185 citations, 60% from search) · paih.gov.pl 47 · biznes.gov.pl 32 · ekrs.ms.gov.pl 31
- Business thread: —
- Errors in answers to companies: none

What the companies’ questions are about. A Czech company does not ask whether to enter Poland, but how: its own subsidiary or a distributor for machine tools and machinery; a Polish supplier of a warehouse system, because it is cheaper; two offers from Polish foundries and injection-moulding plants for the automotive industry; a cross-dock serving Central Europe; a production hall in Poland or Slovakia, Poland or Hungary; cooperation with Polish universities on research projects; a service centre for 50 people in Kraków, Wrocław or Katowice. Only two questions do not name Poland. In the shortlist Poland appears in 97.9% of answers, alongside Slovakia, Hungary and Romania, and Czechia itself appears in three quarters of the answers.

What follows. Czechia is a market in which Poland is at once the most frequent option and a point of comparison: 75.0% of the answers to questions that do not name Poland name the persona’s home country, more than in any other market, because a Czech company measures Poland against what it has just across the border. The answers point to Polish registers more often than anywhere except Slovakia (48.2%): PAIH, biznes.gov.pl and KRS, the National Court Register, are the typical trio in a Czech answer.

> **“Polsko je velmocí v oblasti SSC v Evropě.”**
>
> “Poland is a powerhouse in shared services centres in Europe.”
>
> — Gemini, Czechia, shortlist, question about the location of a shared services centre (Log ID 6a68b302-93d6-4be2-84e2-c55793c0414a). Poland, with Kraków, Wrocław and Katowice.

7.4 France

Market card for companies

- Companies' questions in the study: 58 (naming Poland 42 · not naming Poland 16) → in questions that do not name Poland, Poland in 74.0% of answers (96 answers)
- Shortlist (48 answers): Poland 100% · Romania 89.6 · Czechia 83.3 · Hungary 66.7
- The persona's home country in the shortlist: 33.3%
- Where the answers to companies point: Polish sources 28.8% (1,129 citations, 59% from search) · ec.europa.eu 43 · paih.gov.pl 26 · ekrs.ms.gov.pl 17
- Business thread: —
- Errors in answers to companies: none

What the companies' questions are about. More often than others, French companies' questions start from Central Europe rather than from Poland: 16 of the 58 do not name Poland. They concern a comparison of labour costs and state aid across several countries in the region, a logistics hub for Germany, a university research partner in Central Europe and software nearshoring in a nearby time zone. When Poland is in the question, it comes as one of a pair: Poland or Czechia for an automotive factory, Poland or Hungary for .NET developers, Poland or Romania for nearshoring, a branch in Warsaw or a distributor for medical equipment and chemicals. In questions that do not name Poland the models mention it in 74.0% of answers; in the shortlist, in every answer, alongside Romania, Czechia and Hungary.

What follows. France is a market of comparisons, so Poland wins or loses on argument, not on presence. Romania appears in the shortlist more often than in any other Western market (89.6%), because French questions about nearshoring and costs bring it up of their own accord. The sources point to the European Commission ahead of PAIH, the Polish Investment and Trade Agency; there is no French domain in the top three.

> **"Pour un deuxième hub dédié au centre et à l'est de l'Europe, les options les plus solides sont la Pologne occidentale — en particulier la zone Poznań / Wrocław / Łódź — ou, si vous privilégiez un rôle plus pan-européen, la Belgique ou le sud des Pays-Bas."**
>
> "For a second hub dedicated to Central and Eastern Europe, the strongest options are western Poland — particularly the Poznań / Wrocław / Łódź area — or, if a more pan-European role matters more, Belgium or the south of the Netherlands."
>
> — Perplexity, France, shortlist, question about a warehouse for Central and Eastern Europe (Log ID df263b55-9184-4a57-94fc-c26e889f3069). Poland as the regional hub, Benelux as the pan-European one. This is the argument of chapter 4 in a single sentence.

7.5 Germany

Market card for companies

- Companies' questions in the study: 60 (naming Poland 55 · not naming Poland 5) → in questions that do not name Poland, Poland in 70.0% of answers (30 answers, tendency)
- Shortlist (48 answers): Poland 100% · Czechia 95.8 · Hungary 87.5 · Slovakia 75.0
- The persona's home country in the shortlist: 72.9%
- Where the answers to companies point: Polish sources 35.0% (1,329 citations, 63% from search) · ec.europa.eu 47 · paih.gov.pl 47 · ahk.pl 23
- Business thread: the Polish-German Chamber of Industry and Commerce (ahk.pl), in 52 of 360 answers; no question asks about it
- Errors in answers to companies: none

What the companies' questions are about. A German company asks as one that already has Polish suppliers and wants to compare them: two offers for hydraulics a fifth apart in price; three audits of precision-parts suppliers with different results; a software partner for an Industry 4.0 project; a research team in Poland working to the same quality standards as in Germany; a plant in Silesia, with a full list of hidden costs and an energy balance; a warehouse in western Poland instead of Hungary; a distributor or a subsidiary, with a decision due by the end of the month. Five questions do not name Poland (a new production line in Eastern Europe, nearshoring, an engineering service provider); in these Poland appears in 70.0% of answers (tendency). In the shortlist Poland is in every answer, Czechia in almost every one (95.8%), and Germany itself in 72.9%.

What follows. Germany is the only market in which a bilateral chamber of commerce makes it into the three most frequently cited domains. Ahk.pl appears in one in seven answers to German companies, although no question asks for the chamber. The models add it of their own accord. Bilateral chambers exist for each of the twelve markets; only this one appears in the answers. For the other eleven, this is the shortest recommendation in this part (chapter 10).

> **"Polen sollte Ihr erster Markt sein – mit Tschechien als sinnvoller zweiter Schritt."**
>
> "Poland should be your first market — with Czechia as a sensible second step."
>
> — Claude, Germany, shortlist, question about a heat-pump manufacturer's expansion into Central and Eastern Europe (Log ID 962d20f9-74c9-4011-a2e0-e1664604f697). The answer opens with the decision.

7.6 Italy

Market card for companies

- Companies' questions in the study: 58 (naming Poland 47 · not naming Poland 11) → in questions that do not name Poland, Poland in 68.2% of answers (66 answers)
- Shortlist (48 answers): Poland 95.8% · Romania 83.3 · Czechia 77.1 · Hungary 77.1
- The persona's home country in the shortlist: 18.8%
- Where the answers to companies point: Polish sources 33.3% (1,184 citations, 61% from search) · ec.europa.eu 48 · paih.gov.pl 29 · ekrs.ms.gov.pl 20
- Business thread: —
- Errors in answers to companies: 1 of 14, a DeepSeek answer that listed Trieste among Poland's logistics hubs (Part I, chapter 7)

What the companies' questions are about. Italian companies' questions speak the language of cost: a CFO wants a report on the full costs of a plant in Poland, a comparison with Romania and Slovakia, a calculation for moving the production of mechanical components; beyond that come an urgent order for a German client and a Polish supplier, a logistics hub in Poland for the German market, a Polish software house for an IoT platform, a distributor or a branch for food and machinery, premium food on the Polish market, furniture, and the acquisition of a mechanical-engineering company. Eleven questions do not name Poland (new suppliers in Eastern Europe, the location of a plant, a nearshore partner); in these Poland appears in 68.2% of answers. In the shortlist it appears in 95.8%, alongside Romania, Czechia and Hungary.

What follows. Italy is, with France and Norway, a market in which Romania stands close to Poland (83.3% against 95.8% in the shortlist): Italian questions about labour costs bring it up of their own accord, just as French questions do. Poland wins where the question concerns logistics to Germany and industrial maturity; it loses on price. The only error in this market concerns logistics, with Trieste listed among Poland's hubs.

> **"La Polonia è attualmente il 'motore logistico' dell'Europa centro-orientale."**
>
> "Poland is currently the 'logistics engine' of Central and Eastern Europe."
>
> — Gemini, Italy, shortlist, question about a logistics hub for Central and Eastern Europe (Log ID c2d5f0c0-f99a-41ae-8eaf-d944ad70fa5b). Western and central Poland as the "main strategic choice".

7.7 Norway

Market card for companies

- Companies' questions in the study: 58 (naming Poland 56 · not naming Poland 2) → too few questions that do not name Poland
- Shortlist (48 answers): Poland 97.9% · Romania 85.4 · Czechia 66.7 · Hungary 56.3
- The persona's home country in the shortlist: 25.0%
- Where the answers to companies point: Polish sources 38.4% (1,253 citations, 61% from search) · paih.gov.pl 44 · gov.pl 36 · ec.europa.eu 33
- Business thread: —
- Errors in answers to companies: 1 of 14, a DeepSeek answer stating that Poland is not in the Schengen area (Part I, chapter 7)

What the companies' questions are about. Norwegian companies' questions have one thread found nowhere else: the sea. A supply chain for offshore wind in the Baltic and Polish ports; research partners for hydrogen and CO_2 capture; exports of salmon and seafood to the Polish market. Beyond that, the questions resemble those from other markets: industrial sensors and B2B software, with Poland as a test market after Scandinavia; a production plant and regulatory risk; a development centre for 40 people in Kraków, Gdańsk or Warsaw; suppliers of CNC parts. Only two questions do not name Poland. In the shortlist Poland appears in 97.9% of answers, ahead of Romania (85.4%) and Czechia; Estonia and Lithuania (56.3% and 52.1% respectively) make the top five here, the only Baltic states to do so in the whole of Part II.

What follows. For a Norwegian company Poland is a country of ports and shipyards, and the models have that map: Gdańsk, Gdynia or Szczecin appear in 33 of the 48 shortlist answers, although none of its eight questions asks about ports or the coast. The only factual error in this market concerns, of all things, a border: an answer to a company placed Poland outside the Schengen area.

> **"Polen er det beste kompromisset for et kjøle- og distribusjonslager som skal betjene Sentral- og Øst-Europa raskt."**
>
> "Poland is the best compromise for a cold-storage and distribution warehouse that has to serve Central and Eastern Europe quickly."
>
> — Grok, Norway, shortlist, question about a cold-storage warehouse for Central and Eastern Europe (Log ID 646ca236-d1de-4a5c-8244-77d64b889fca). Poland as the regional hub, with Poznań, Łódź and Gdańsk.

7.8 Slovakia

Market card for companies

- Companies' questions in the study: 58 (naming Poland 56 · not naming Poland 2) → too few questions that do not name Poland
- Shortlist (48 answers): Poland 97.9% · Czechia 93.8 · Hungary 79.2 · Romania 70.8
- The persona's home country in the shortlist: 60.4%
- Where the answers to companies point: Polish sources 50.3% (1,681 citations, 59% from search) · ekrs.ms.gov.pl 48 · paih.gov.pl 35 · krd.pl 31
- Business thread: —
- Errors in answers to companies: none

What the companies' questions are about. A Slovak company asks about southern Poland: a pilot entry into the Silesian Voivodeship; a machinery plant in the southern regions; a warehouse near Kraków or a 3PL partner; an industrial zone near Katowice as a family firm's first export market; entry into Polish automotive supply chains; a service centre for 40 people in Kraków, Katowice or Rzeszów. Six questions return to the same point: where can the solvency of a Polish company proposing a distribution deal be checked before the contract is signed? Only two questions do not name Poland. In the shortlist Poland appears in 97.9% of answers, Czechia in 93.8%, and Slovakia itself in 60.4%.

What follows. Slovakia is the market with the highest share of Polish sources in the whole study (50.3%), and the only one in which KRD (Krajowy Rejestr Długów), a Polish debt register, makes it into the three most frequent domains. This answers the Slovak questions directly: whoever asks about a counterparty's solvency gets KRS, the National Court Register, and KRD. Here Poland is a country in which a partner is checked before the contract is signed, and the answers know where to do it.

> **"Program 'Polska Strefa Inwestycji' (PSI) – nahradil staré špeciálne ekonomické zóny."**
>
> "The 'Polska Strefa Inwestycji' (Polish Investment Zone, PSI) programme — replaced the old special economic zones."
>
> — DeepSeek, Slovakia, shortlist, question about investment incentives and economic zones in the region (Log ID 945c5490-2e1d-4388-acd1-386e3a60b402). The Polish incentive programme described by name, alongside the programmes of Czechia, Hungary and Romania.

7.9 Spain

Market card for companies

- Companies' questions in the study: 58 (naming Poland 56 · not naming Poland 2) → too few questions that do not name Poland
- Shortlist (48 answers): Poland 100% · Czechia 75.0 · Germany 70.8 · Romania 66.7
- The persona's home country in the shortlist: 37.5%
- Where the answers to companies point: Polish sources 37.7% (1,100 citations, 60% from search) · paih.gov.pl 43 · gov.pl 41 · ec.europa.eu 34
- Business thread: —
- Errors in answers to companies: 2 of 14, a Claude answer that warned of "customs checks within the EU" between Spain and Poland and a Gemini answer that placed a plant in "the Grand Duchy of Poland" (Part I, chapter 7)

What the companies' questions are about. Spanish companies' questions treat Poland as a substitute for Asia and as a sales market: Polish manufacturers of metal parts in place of an Asian supplier; an automotive-components plant in Poland or Romania, Poland or Hungary; a logistics hub for northern Europe; wind and solar farms; a recycling plant financed with EU funds; a research centre for a software company; olive oil, wine and tinned food for the Polish market; an online footwear shop for Polish customers. Only two questions do not name Poland. In the shortlist Poland appears in every answer, alongside Czechia, Germany and Romania, and Spain itself appears in 37.5% of answers.

What follows. In Part I Spain was the only market in which Poland lost out to Czechia in questions about the region. In companies' questions that exception disappears: Poland is in every shortlist answer, Czechia in three quarters of them. For a Spanish company Poland is a candidate for a plant, not for a weekend. The sources here are more Polish than in any other Western market except Norway (37.7%): PAIH, the Polish Investment and Trade Agency, and gov.pl come ahead of the European Commission.

> **"Si la prioridad es experiencia probada en centros de servicios compartidos, los destinos más consolidados en Europa del Este son Polonia, Rumanía, República Checa, Hungría y Bulgaria."**
>
> "If the priority is proven experience in shared services centres, the most established destinations in Eastern Europe are Poland, Romania, Czechia, Hungary and Bulgaria."
>
> — Perplexity, Spain, shortlist, question about a shared services centre (Log ID e63a25d1-c5a1-4da0-9e23-90253d7003ca). Poland on the list.

7.10 Sweden

Market card for companies

- Companies' questions in the study: 58 (naming Poland 56 · not naming Poland 2) → too few questions that do not name Poland
- Shortlist (48 answers): Poland 95.8% · Czechia 77.1 · Romania 62.5 · Germany 54.2
- The persona's home country in the shortlist: 31.3%
- Where the answers to companies point: Polish sources 36.0% (977 citations, 65% from search) · ec.europa.eu 44 · paih.gov.pl 44 · gov.pl 23
- Business thread: the Baltic ferry, in 87 of 348 answers; Ystad in 23 of 348; no question asks about a ferry
- Errors in answers to companies: none

What the companies' questions are about. Swedish companies' questions come down to three decisions repeated in many variants: moving customer service to Poland, and a list of cities with multilingual staff; a trading company in Poland, in Warsaw or Kraków; energy storage in Poland, and the regions with access to green energy. On top of these: production moving from China to Poland; offers for components from Poland and Latvia, or from Poland and Czechia; a logistics hub for Scandinavia and Central Europe, in Warsaw or Poznań. Only two questions do not name Poland. In the shortlist Poland appears in 95.8% of answers, ahead of Czechia (77.1%) and Romania.

What follows. In Part I the ferry from Ystad was an association the Swedes have with Poland that was almost absent from the answers to individuals (3 of 348). In the answers to companies the reverse holds: the ferry appears in one in four answers, and Ystad in 23, although no question asks about either. The same route across the Baltic that is no argument for individuals is one for logistics. The models add it of their own accord when the question concerns a warehouse, deliveries and transit times. The ferry is an argument for companies, not for individuals.

> **"Polen har den största och mest mogna talangpoolen i regionen."**
>
> "Poland has the largest and most mature talent pool in the region."
>
> — Grok, Sweden, shortlist, question about the location of a shared services centre (Log ID 9ed5c23b-077c-4e3a-b4a6-008573911258). Poland "the strongest overall".

7.11 Ukraine

Market card for companies

- Companies' questions in the study: 57 (naming Poland 56 · not naming Poland 1) → too few questions that do not name Poland
- Shortlist (48 answers): Poland 100% · Romania 83.3 · Czechia 62.5 · Germany 60.4
- The persona's home country in the shortlist: 2.1%
- Where the answers to companies point: Polish sources 46.4% (1,411 citations, 56% from search) · ekrs.ms.gov.pl 49 · ec.europa.eu 38 · paih.gov.pl 36
- Business thread: —
- Errors in answers to companies: none

What the companies' questions are about. A Ukrainian company asks about Poland as a gateway to the EU and as a back-up base: a grain store in Lublin or Rzeszów ahead of delivery to Germany; CE certification of food and electronics through a Polish subsidiary; moving part of the welding and packaging work to Poland from a plant running at 60% capacity; its own sales branch in Warsaw instead of an agent; a Polish partner for tenders for the reconstruction of Ukraine; suppliers of steel sections and fasteners for building sites in Ukraine; an IT office in Wrocław or Kraków for clients in Czechia and Germany. Only one question does not name Poland. In the shortlist Poland appears in every answer, ahead of Romania (83.3%), Czechia and Germany; Ukraine as the home country, in 2.1%.

What follows. For Ukrainian companies Poland is an operational base in the EU, not one option among several: in the shortlist it appears in every answer, and Ukraine as the home country almost never (2.1%), because a question from Ukraine about a plant or a warehouse is by definition a question about going abroad. The answers point most often to the National Court Register (49 answers): Ukrainian questions about a tender partner and a steel supplier ask for verification and get it. The Wrocław and Legnica example below shows how detailed this map can be.

> **"Для заморожених овочів і ягід, орієнтованих на супермаркети Центральної Європи, найкращий базовий варіант — південно-західна Польща, передусім район Вроцлав–Легниця, якщо товар заходить до ЄС автотранспортом через Україну."**
>
> "For frozen vegetables and berries aimed at Central European supermarkets, the best base option is south-western Poland, above all the Wrocław–Legnica area, if the goods enter the EU by road through Ukraine."
>
> — ChatGPT, Ukraine, shortlist, question about a cold-storage warehouse for deliveries to Central Europe (Log ID 153e5884-2043-4013-899f-bdb430c9ef54). The answer points not to a country but to an area.

7.12 United Kingdom

Market card for companies

- Companies' questions in the study: 59 (naming Poland 57 · not naming Poland 2) → too few questions that do not name Poland
- Shortlist (48 answers): Poland 95.8% · Czechia 70.8 · Romania 60.4 · Germany 54.2
- The persona's home country in the shortlist: 14.6%
- Where the answers to companies point: Polish sources 29.0% (1,260 citations, 64% from search) · ec.europa.eu 41 · gov.uk 35 · paih.gov.pl 27
- Business thread: —
- Errors in answers to companies: none

What the companies' questions are about. A British company asks about Poland after Brexit: a plant in the EU to avoid customs friction; Polish suppliers of precision parts after quality problems with an Asian supplier; a shared services centre for finance and HR; a distribution hub (Poland or Hungary, the road network and freight reliability); a team of developers for an urgent project; a food exporter with Poland as its first EU market; a games studio to acquire; a £30 million investment over 10–15 years. Only two questions do not name Poland. In the shortlist Poland appears in 95.8% of answers, ahead of Czechia (70.8%) and Romania.

What follows. The United Kingdom is one of two markets (with China) in which the answers to companies point to the authorities of the home country more often than elsewhere (10.2% of answers;

chapter 6): gov.uk appears in 35 answers, ahead of PAIH, the Polish Investment and Trade Agency, because the British government has pages on exporting to every EU country. In Part I the same market pointed individuals to gov.uk more often than to gov.pl; for companies the pattern is similar, only the stakes differ. This is a page about trade with Poland, not about travel safety.

> **"If your goal is to reach most of Central Europe within two days, I would put the warehouse in western Poland or the Germany–Netherlands corridor depending on whether you prioritize reach/cost or fastest pan-EU access."**
>
> — Perplexity, United Kingdom, shortlist, question about a warehouse for Central Europe (Log ID 58a5c8ee-c593-4d0b-a6ef-c75b1e3b45da). Poland for reach and cost; Germany and the Netherlands for the fastest access to the whole EU.

7.13 United States

Market card for companies

- Companies' questions in the study: 60 (naming Poland 42 · not naming Poland 18) → in questions that do not name Poland, Poland in 69.4% of answers (108 answers)
- Shortlist (48 answers): Poland 100% · Czechia 87.5 · Germany 85.4 · Romania 66.7
- The persona's home country in the shortlist: 2.1%
- Where the answers to companies point: Polish sources 26.5% (1,338 citations, 66% from search) · ec.europa.eu 44 · paih.gov.pl 29 · linkedin.com 21
- Business thread: —
- Errors in answers to companies: none

What the companies' questions are about. American companies' questions are the questions asked before a board decision: "one hard justification for the board", "a final go/no-go", "three reasons to choose Poland over...". In them Poland is paired with Romania (a plant), Czechia (an R&D centre) and Germany (a logistics hub with 48-hour delivery to France and Germany); beyond that come precision-parts suppliers and their certificates, a sales office as a way into the EU market, the acquisition of a Polish software company, defence procurement, a base for contracts on the reconstruction of Ukraine, and a $60 million investment over 15 years. As many as 18 questions do not name Poland, the most in the study. They concern a hub for France and Germany, a first plant in Europe, an engineering team in Europe and the B2B sales cycle in the EU. In these, Poland appears in 69.4% of answers. In the shortlist it appears in every answer, alongside Czechia and Germany; the United States as the home country, in 2.1%.

What follows. A company from the United States asks about Europe and gets Poland as one of the options. It is there in seven answers out of ten when the question does not name it, and always when the question narrows the choice to the region. The sources here are the second least Polish in Part II (26.5%), with the European Commission at the top and the professional network linkedin.com in the top three: American questions about talent get a source about people, not about the country.

> **"For a ~50-person semiconductor R&D center, the strongest European cost-to-talent options are generally Poland, the Czech Republic, Romania, and Portugal."**
>
> — ChatGPT, United States, shortlist, question about an R&D centre in Europe (Log ID 11de16ac-d240-4496-931c-cf754b435511). Poland on the list.

HOW WE COUNT (ALL CARDS). Market card: companies' questions from the study (August–September 2026), counting the questions with and without Poland in the text (the "Poland named in the question" flag); mention of Poland in answers to questions that do not name Poland only for markets with ≥ 5 such questions (with 5 questions: 30 answers, tendency). Shortlist: a separate measurement of 30 September 2026 (chapter 2): 8 questions that do not name Poland per market × 6 models = 48 answers; "Poland in x%" = the share of answers that mention Poland; countries alongside Poland = the share of answers in which a given country appears (the judge's field "other entities mentioned", after normalisation); the persona's home country = the share of answers in which the market's country appears. Sources: the share of Polish domains among citations in answers to a given market's companies' questions, combining both modes (sources from search and URLs given in the answer; chapter 6), with the share of citations that come from search (a profile of the market, not a ranking); three domains = the number of answers in which the domain appears. Business thread: a search of the full text of companies' answers in the study using a list of the name's forms (China–Europe railway, ferry and Ystad, the ahk.pl chamber), split between questions that contain the thread and the rest. We do not publish the recommendation rate per market (study: 12 answers without Poland in most markets; shortlist: dependence on the length cap on answers, see the annex). Figures from the two data sets do not stand in the same column. The measurement is of the model's answer returned via the API.

CHAPTER 8

Poland in the shadow of war, or what companies hear about risk

IN BRIEF

- Asked directly about investing, nearshoring, energy, workers or insurance in Poland, the language models (LLMs) did not advise against Poland as a choice because of the war in any of 864 answers.
- The recommendation almost always comes with a condition: in the four questions on investment risk, answers with no caveat at all account for 3.8%, and “yes, but” answers for 81.9%.
- It is the same Poland in two roles: about two thirds of answers to the questions on investment, nearshoring and insurance present Poland as clearly more exposed than the countries it is compared with, and nine in ten of those same answers recommend it with caveats.
- As a base for Ukraine’s reconstruction, Poland is explicitly recommended in every answer, including when the question does not name it.
- When a company asks about safe countries in the region without naming Poland, Poland always appears, but in 37.5% of answers it is placed among the riskier countries, and in 3 of 144 answers it is advised against.
- Companies hear different arguments from individuals: the European Union (86.9% of answers) and energy (75.0%) rather than Belarus and foreign-ministry warnings.

In October 2026, separately from the study, the six models were given eight questions from companies about Poland during the war in Ukraine (on a €20 million investment, nearshoring, war-risk insurance and a base for serving Ukraine’s reconstruction) in 12 markets, each asked twice. In none of the 864 answers to the questions that name Poland did a model advise against it as a choice because of the war. Yet only 3.8% of answers on investment risk recommend it without any caveat. This chapter measures how large that “but” is, and what the models weigh it against.

8.1 People first, or what individuals hear

Individuals ask about the same things differently, and it is their picture of Poland that employees, students and expatriate families later encounter. Part I describes this measurement in box 4.7; this section gives only what the comparison with companies requires.

Reminder from Part I (section 4.7). In October 2026, separately from the study, we put eight questions from individuals about safety in Poland during the war in Ukraine to the six models, in 11 markets, each asked twice (1,056 answers). In none of the 792 answers to the questions that name Poland did a model advise against Poland as a destination because of the war, but only 15.7% answer without any caveat. The arguments offered to individuals are consistent: Belarus (88.3% of answers), NATO (78.1%) and warnings from the foreign ministry of the persona's home country (70.5%); the Schengen area comes up in 5.3%. When a question about safe countries anywhere in Europe does not name Poland, between half and three quarters of answers put it in the "with caution, closer to the conflict" group.

"Poland is a heavily defended NATO ally. It's essentially one of the most secure places in Europe right now."
— Gemini, United States, question from a person whose parents are worried about them going to Poland (Log ID 7e6fa012-6c92-4071-ba1b-e4c877719720). The model suggests what to tell the parents.

How we count. The thematic measurement of 18–19 October 2026 is a separate dataset: different questions, a different date, and ChatGPT, Gemini, Claude and DeepSeek running in newer versions than in the study. We do not combine its figures with those of the study. We assigned each answer a stance on risk in one of four classes: no caveat, "yes, but", "with caution", advises against. This was done by a single coder (a language model) blind to which model had answered; a second run on a random sample of 576 answers agreed with it on 87.5%. In the question about Europe as a whole, the boundary between "yes, but" and "with caution" depends on the classification instructions, which is why we give a range. The method is described in the annex.

8.2 Questions from companies: eight situations, 1,152 answers

On the same day as the questions from individuals, we put eight questions from companies to the six models in all 12 markets, each asked twice. Six of them name Poland: a €20 million investment, nearshoring to Poland or further west, Poland as a base for Ukraine's reconstruction, energy, workers from Ukraine and war-risk insurance. Two ask without naming Poland: about safe countries in the region for investment and about the best base for Ukraine's reconstruction.

Asked directly: yes, with conditions. In none of the 864 answers to the six questions that name Poland did a model advise against it as a choice because of the war. Some answers on investment and nearshoring, however, point a board that ranks war risk highest to a location further west, most often in Czechia; under other priorities they recommend Poland, so they belong to the "yes, but" or "with caution" classes (we did not count their share). In the risk questions (investment, nearshoring, insurance, safe countries in the region; 576 answers) the typical answer is "yes, but": 81.9%. No caveat at all: 3.8%; flagging significant risk: 13.7%; advises against: 0.5%. The most common condition runs as follows: the risk is higher than in Czechia or Germany, so it has to be priced and insured, not "do not invest".

8.3 The same Poland in two roles

A "yes, but" stance does not yet say how large the "but" is. In the three questions in which a company compares Poland with other countries (a €20 million investment, nearshoring, war-risk insurance; 432 answers), we therefore coded a relative measure separately: whether the answer presents Poland as clearly more exposed to war risk than the countries it is compared with.

The result cuts both ways. About two thirds of answers (67%) present Poland as clearly more exposed. This is most often the case in the investment question (76%) and the insurance question (73%), less

often in the nearshoring question (53%); 2% of answers make no comparison at all. And yet nine in ten of the same answers (91%) recommend Poland with caveats. In them, Poland is at once a country closer to the front and a country worth entering. Both judgements appear in a single answer and do not contradict each other: the model first places Poland on the risk map, then works out whether that risk can be priced and insured, and usually concludes that it can.

This measure has a band of uncertainty. Two halves of the same question, coded in different runs, differ by a dozen or so points (85% and 67% in the investment question), so we read it as ranges: investment, from two thirds to more than four fifths of answers; nearshoring, about half; insurance, about three quarters. The extreme stance classes, by contrast, give firm figures: answers with no caveat never present Poland as more exposed, while answers flagging significant risk almost always do (22 of 25). The share of "clearly more exposed" depends more on the question than on the model: for the same model it ranges from 0% in the nearshoring question to 96% in the insurance question (Gemini); across models the spread runs from 53% (Gemini) to 90% (Perplexity). These are, however, the October 2026 versions, which differ from those in the study, so we do not carry this ordering over to the main study (Part I).

Tone depends on the topic, not the market. The models are coolest about war-risk insurance (tone 5.47; explicitly recommended in 2.8% of answers, low recommendation level in 12.5%) and about the €20 million investment (5.94; 14.6%). They are warmest about Poland as a base for Ukraine's reconstruction (7.50; explicitly recommended in 98.6%) and about workers from Ukraine (7.05; 70.8%). The energy question barely touches on the war in Poland itself: 78.5% of answers do not assess war risk at all, but every one of them discusses energy prices and sources (explicitly recommended in 4.2%, moderate recommendation level in 93.8%). The tone of answers for companies across the 12 markets differs by less than half a point (China 6.46 – Ukraine 6.90). On nearshoring ("Poland or further west"), between 8% and 14% of answers put Poland in the "with caution" group; the rest recommend it with conditions.

8.4 Without Poland in the question: a reconstruction base, yes; a safe investment, conditionally

When a company asks about the best base for serving Ukraine's reconstruction, Poland appears in every answer and is explicitly recommended in every one; in three quarters of these answers the model does not address the war risk to Poland itself at all. When a company asks about safe countries in the region for investment, Poland again always appears, but in 37.5% of answers it is placed in the "riskier" group, and in 3 of 144 (2.1%) the model recommends another country. Proximity to Ukraine is an asset when the subject is Ukraine, and a liability when the subject is the safety of capital.

Here the models differ more than the markets. In the question about safe countries in the region, ChatGPT places Poland among the riskier countries in 87.5% of answers (21 of 24), and Gemini in none (0 of 24); the other models fall between the two. This is a difference between the October 2026 versions: in the August study both models were running older versions, so we neither carry this difference over to the main study (Part I) nor read it as a permanent trait of either model. For a company it means one thing: which model an investor asks changes the answer about Poland more than which country the investor asks from.

Arguments for companies. The models justify their answers to companies with European Union membership and EU funds (86.9% of answers), energy (75.0%), NATO (51.1%) and Belarus (43.2%). Foreign-ministry warnings, present in 70.5% of answers for individuals, are hardly ever invoked for companies (3.4%).

8.5 How it sounds in the answers

> **“For most businesses serving Ukraine’s reconstruction, Poland is the best all-round base.”**
>
> — ChatGPT, United States, question about the best base for Ukraine’s reconstruction, without naming Poland (Log ID b299068d-33b1-48dd-9ffc-5ceb73991740).

> **“Setting up a regional hub in southeastern or central Poland (Rzeszów or Warsaw) allows you to project power into Ukraine while maintaining a safe EU base for your capital, contracts, and talent.”**
>
> — Gemini, United States, question about Poland as a base for Ukraine’s reconstruction (Log ID fb9b4ddc-0cc5-4ec7-9ddc-9b774aa1a6bc).

> **“Z čisto bezpečnostného hľadiska je Poľsko stabilné. Z hľadiska vnímania rizika vedením je však často jednoduchšie a rozumnejšie zvoliť Česko alebo západné Slovensko”**
>
> “From a purely security standpoint, Poland is stable. From the standpoint of how management perceives risk, however, it is often simpler and more sensible to choose Czechia or western Slovakia”
>
> — Grok, Slovakia, question about nearshoring to Poland or further west (Log ID de23699c-b279-48a2-a955-ba66e2bd8c1b). Poland is safe, but for a risk-sensitive board the model suggests Czechia. Classified as “with caution”.

HOW WE COUNT. Thematic measurement of 18–19 October 2026: 16 questions (8 from individuals in 11 markets, 8 from companies in 12 markets), six models queried through their official APIs, each question asked twice, giving 2,208 answers, 1,152 of them from companies. It is a separate dataset (different questions, a different date, four models in newer versions than in the study), and we do not combine its figures with those of the study. We assigned each answer a stance on risk (no caveat / “yes, but” / “with caution” / advises against / not rated) by a single coder (a language model) that did not know which model had answered; an independent second run covered a stratified random sample of 576 answers (agreement 87.5%); in the investment and nearshoring questions the boundary between “yes, but” and “with caution” depends on the instructions, which is why we give ranges or a description. The relative measure (“clearly more exposed than the countries compared”) was coded separately for 432 answers in six independent runs; halves of the same question differ by up to 18 percentage points, so we give ranges. A factual error was confirmed in one business answer (the year the Yamal gas pipeline was shut down: 2002 instead of 2022). Some answers from three models are truncated by the length cap set in the research application (see the annex).

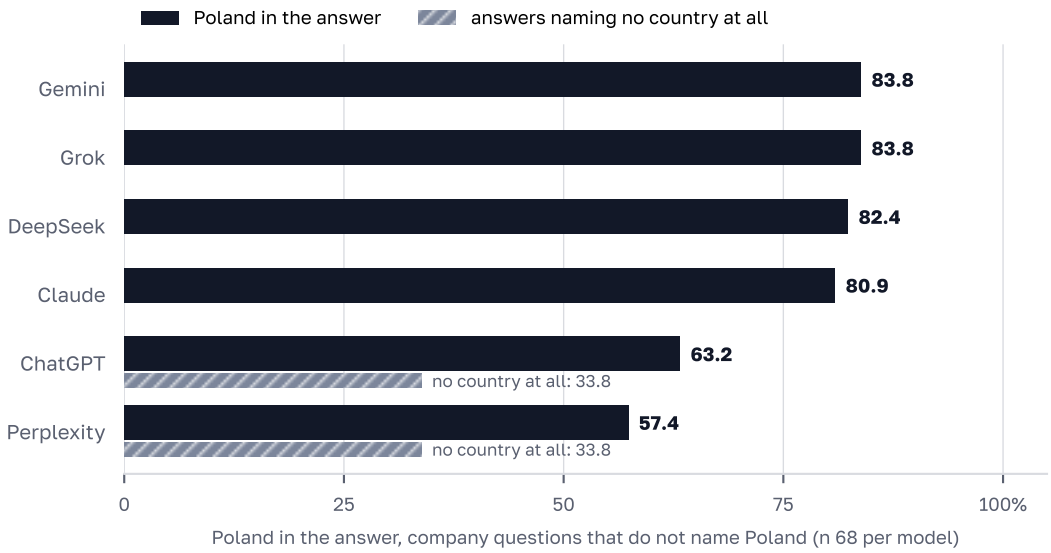
CHAPTER 9

The models, or the differences a company should know about

IN BRIEF

- The six language models (LLMs) agree on whether to tell companies about Poland: in 665 of the 700 company questions, all of them answer the same way (95.0%). The dispute lies where the question does not name Poland: 30 of the 35 split questions are questions that do not name Poland, and the study contains 68 of them.
- In these 68 questions the models differ sharply: Gemini and Grok bring Poland into 83.8% of their answers, DeepSeek into 82.4%, Claude into 80.9%, ChatGPT into 63.2% and Perplexity into 57.4% (68 answers each). What makes the difference is not the assessment of Poland but the way of answering: in more than a third of such answers, ChatGPT and Perplexity name no country at all, because they answer with a procedure rather than a list (chapter 3).
- The tone in which the models speak of Poland in company questions puts them in the same order as in Part I: on the same questions, Gemini speaks of Poland more warmly than the model average (+0.47 points) and Claude the most coolly (−0.39). None of these divergences exceeds the 0.5-point threshold, but on a single question the warmest and the coolest answer differ by an average of 1.8 points, and in one question in five (20.5%) by 3 or more.
- When a company asks about risk, every model lowers its tone by 0.9–1.3 points and every model answers negatively more often. They differ in scale: in risk questions, negative answers account for 6.1% at Gemini and 17.5% at Claude; in the other questions, 0.7–3.4% at each.
- Which models recommend Poland to companies most explicitly is reported only in the annex: four of the six models had a length cap on their answers, and a cut answer more often loses its explicit recommendation (Part I, annex). In the shortlist, the order of the models by recommendation rate cannot be established at all. For this reason there is no “best model for Poland”. There are six different ways of answering, and a company that wants to know how it looks in the answers has to measure across several models at once.

Chart 9.1 Poland in company questions that do not name Poland, six models



Same question, same day, two models. A German company is planning a new production line and asks for a comparison of Poland and Czechia for its chief financial officer. Gemini opens with a verdict: Poland is the “scaling winner”, with the better cost profile over ten years. Claude opens with a caveat about the lack of current data and closes with the statement that Poland’s cost advantage is “melting”. Tone 8 against 4 (both quotations in 9.5). Both models gave the same orders of magnitude for wages and the same sources to check; what separates them is what they made of them. This chapter measures exactly that difference and shows that less follows from it for companies than it might seem.

9.1 What this chapter measures

Part I (chapter 7) compared the six models across all 1,400 questions. This chapter does the same on the 700 company questions of the study and adds the shortlist (96 questions that do not name Poland, a separate measurement). The measures are the same: agreement on mention (whether all the models talk about Poland), mention in questions that do not name Poland, the tone in which Poland is spoken of, and **model divergence**, that is, how far a given model’s tone departs from the average tone of all six in their answers to the same question. Comparing on the same questions removes the effect of question selection, which makes it the only fair measure of which model speaks of Poland more warmly.

One measure is absent here: the recommendation rate per model. This is not because the models do not differ on it (they do), but because in the study four of the six models had an upper length cap on their answers and often reached it, and the truncation simulation on full answers shows that a cut answer more often loses its explicit recommendation (Part I, chapters 4 and 7). The per-model values, with a caveat, are given in the annex. The text keeps the measures that are robust to truncation: mention, tone and divergence.

9.2 The models agree on whether to tell companies about Poland

In 657 of the 700 company questions, all six models name Poland; in 8, none does. Agreement is thus almost complete (95.0%), but for the same reason as in Part I: 632 of the 700 company questions

name Poland outright, and when they do, every model mentions it. The split begins in the 68 questions that do not name Poland: in 30 of them, some models bring Poland in and others do not.

model	Poland in the answer, questions that do not name Poland (n 68)	Poland in the answer, overall (n 700)
Gemini	**83.8%**	98.4%
Grok	**83.8%**	98.3%
DeepSeek	82.4%	98.1%
Claude	80.9%	97.7%
ChatGPT	63.2%	96.4%
Perplexity	57.4%	95.6%

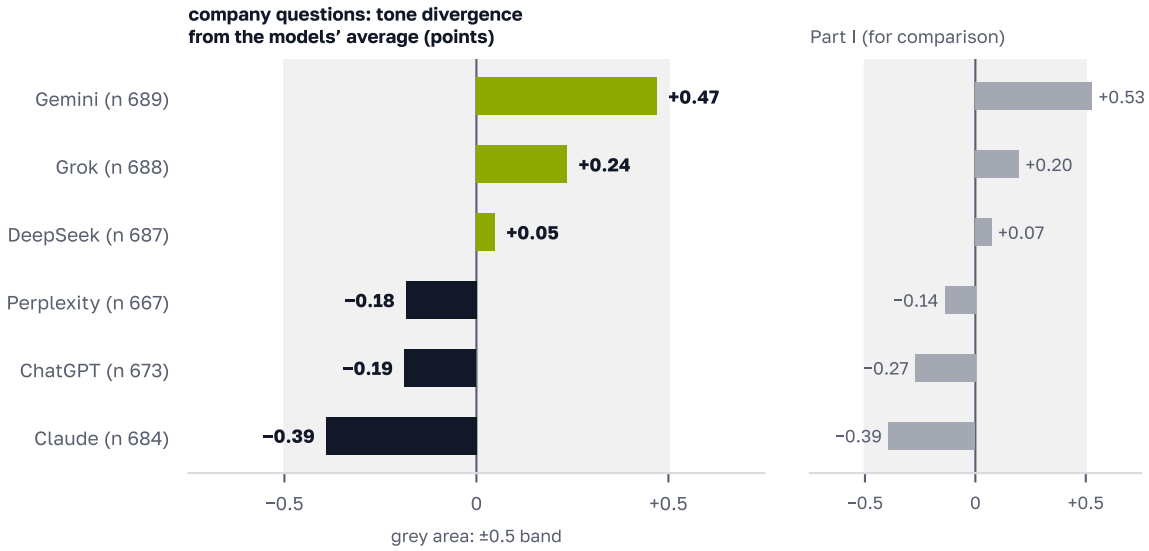
What follows from this table? Overall, the differences fall within 3 percentage points, that is, within measurement variability. In questions that do not name Poland, the spread reaches 26 percentage points and is real. It is not, however, a difference in the assessment of Poland. Chapter 3 showed where it comes from: in more than a third of their answers to company questions that do not name Poland, ChatGPT and Perplexity name no country at all. Asked about the sales cycle in the EU, company registration or certification, they answer with a description of the procedure rather than a list of countries. Gemini, Grok and DeepSeek answer the same question with a list, and Poland is on it. A company that wants a model to suggest Poland of its own accord therefore stands a better chance with the models that answer with a list; a company that wants a procedure will get one without Poland, but also without Czechia and Germany.

In the shortlist (96 questions that do not name Poland, 8 decision areas) there are almost no differences: Gemini, Grok and DeepSeek name Poland in 100% of their answers, ChatGPT in 99.0%, Claude in 96.9% and Perplexity in 94.8% (96 answers each; chapter 2). When the question narrows the choice to a location decision in Europe, Poland is on the list at every model.

9.3 Tone, or the same order as in Part I

model	tone divergence on the same company questions	tone overall	negative answers (≤ 4)	n
Gemini	**+0.47**	7.25	2.9%	689
Grok	+0.24	7.01	5.4%	688
DeepSeek	+0.05	6.83	5.7%	687
Perplexity	−0.18	6.61	7.9%	667
ChatGPT	−0.19	6.60	4.8%	673
Claude	**−0.39**	6.40	9.1%	684

Chart 9.2 Tone divergence in company questions



The order is the same as across all 1,400 questions in Part I (Gemini +0.53, Claude −0.39): Gemini speaks of Poland to companies the most warmly, Claude the most coolly. No divergence in companies' questions exceeds the 0.5-point threshold, so the report calls no model "favourable" or "unfavourable". The difference between the two extreme models (0.86 points) is nevertheless larger than measurement noise, and on a single question it can be much larger: in the 655 company questions to which all six models answer with a mention of Poland and a tone score, the warmest and the coolest answer differ by an average of 1.8 points, and in one question in five (20.5%) by 3 points or more. For a company checking "what the model says about Poland", this means something simple: one question put to one model does not show what the models say. It shows what one of them said on that day.

Negative answers (tone 4 or lower) account for 5.9% of answers in company questions (Part I, chapter 4; 4,088 answers with a mention and a tone score), from 2.9% at Gemini to 9.1% at Claude. Where they come from is the subject of the next section.

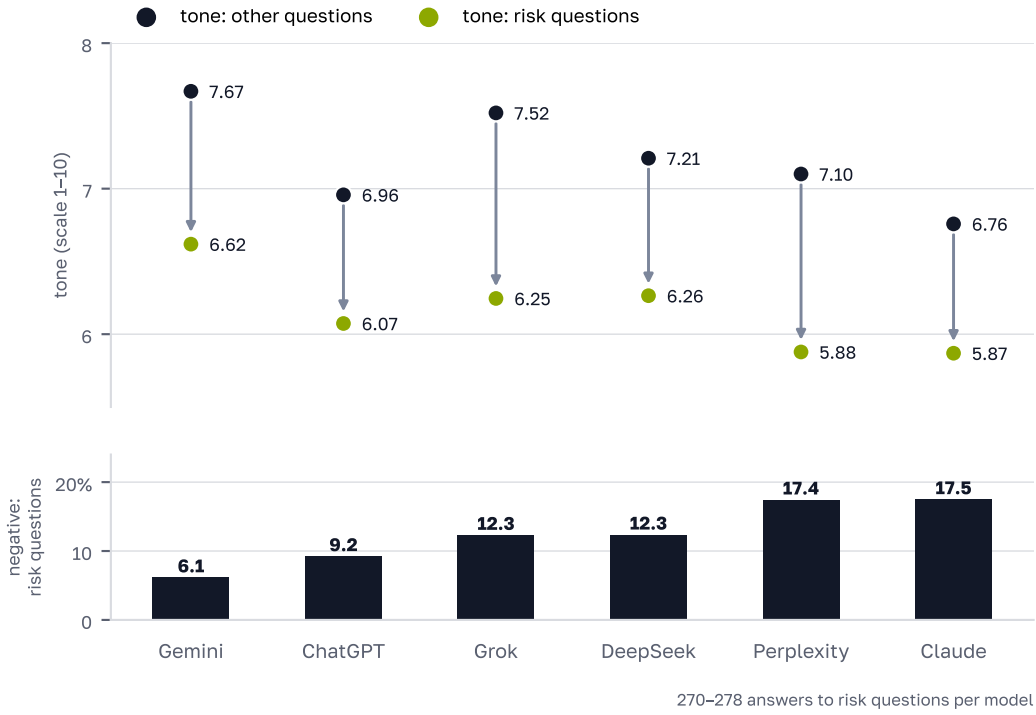
9.4 Every model issues warnings when a company asks about risk, but not equally strongly

Chapter 1 showed that the tone of answers to companies is decided by the question: 281 of the 700 company questions themselves ask about risks, barriers or pitfalls, and it is in these that the answers are cool (6.16) and one in eight is negative (12.4%); in the other questions, tone is 7.20 and 1.6% of answers are negative. This relationship holds for each of the six models, and each lowers its tone by a similar amount.

model	tone: risk questions	tone: other questions	difference	negative: risk	negative: other
Gemini	6.62	7.67	−1.05	6.1%	0.7%
Grok	6.25	7.52	−1.28	12.3%	0.7%
DeepSeek	6.26	7.21	−0.94	12.3%	1.2%
ChatGPT	6.07	6.96	−0.88	9.2%	1.7%
Perplexity	5.88	7.10	−1.22	17.4%	1.5%
Claude	5.87	6.76	−0.89	17.5%	3.4%

270–278 risk answers and 397–411 other answers per model.

Chart 9.3 Every model issues warnings



What follows from this? First, no model "scares companies off Poland" of its own accord: in questions that do not ask about risk, negative answers are rare at every model (0.7–3.4%). Second, when a company asks about risk, every model answers with risk, and the drop in tone of 0.9–1.3 points is common to all six. Third, they differ in how often they turn a list of risks into a negative answer: Claude and Perplexity do so almost three times as often as Gemini (17.5% and 17.4% against 6.1%). For a company asking about the pitfalls of investing in Poland, this means that it will get a list of risks from every model, but how heavy that list is will depend on which model it asks. For the economic promotion of Poland, it means that the task is not a "bad opinion held by the models" but the content the models hold on specific risks: work-permit legalisation, energy costs, regulatory change (chapters 1, 8 and 10).

9.5 How it sounds in the answers

Two answers to the same question from a German company that is planning a new production line and asks for a comparison of Poland and Czechia for its chief financial officer over a ten-year horizon:

> **“Polen (Der Skalierungssieger): Bietet über die nächsten 10 Jahre das bessere TCO-Profil für arbeits- und flächenintensive Produktionen.”**
>
> “Poland (the scaling winner): offers the better total-cost-of-ownership profile over the next 10 years for labour- and space-intensive production.”
>
> — Gemini, Germany, question comparing Poland and Czechia for a new production line (Log ID 6248c287-7479-4b9f-9b7c-9abbf6e14a35). Tone 8, high recommendation level; the answer opens with a summary for the chief financial officer.

> **“Fazit Arbeitskosten: Polen ist heute günstiger, der Kostenvorteil schmilzt über 10 Jahre erheblich.”**
>
> “Conclusion on labour costs: Poland is cheaper today; the cost advantage melts away considerably over 10 years.”
>
> — Claude, the same question (Log ID aa344af7-2c2f-46cb-a7bb-762785c545ed). Tone 4, moderate recommendation level; the answer opens with a caveat that the model has no current market data.

Both answers give similar wage ranges and the same sources to check. They differ in framing: one says “Poland wins, unless...”, the other “Poland is cheaper, but...”. The judge scored this difference at four points of tone and one recommendation level. A company that puts this question to one model will get one of these two framings and will never learn of the other.

9.6 What this measurement cannot say about the models

Three things. First, which model recommends Poland to companies most explicitly. The values are in the annex, with a caveat about the length cap; in the shortlist the differences between the models are smaller than the possible effect of truncation, so the order cannot be established (chapter 2). Second, whether that order is lasting: Part I showed that within seven weeks newer versions of the models shifted the share of explicit recommendations of Poland by a dozen or so points, in both directions; the order of the models is a property of the August 2026 versions, not of the model brand. Third, which model is “best for Poland”: the question has no single answer, because it has as many answers as there are goals. A company that wants to be suggested without being asked looks at mention in questions that do not name Poland; a company that wants to be spoken of more warmly looks at tone divergence; a company asking about risk will get it everywhere, but in different doses. The conclusion for each of them is the same: measure across several models at once, on the same questions, in a dated measurement window, because each of these figures is a property of a version, not a truth about the model.

HOW WE COUNT. Company questions: 700 questions × 6 models = 4,200 answers of the study (August–September 2026); agreement on mention = the share of questions in which either all six models name Poland or none does; mention in questions that do not name Poland: 68 questions × 6 = 408 answers (the “Poland named in the question” flag). Model divergence: the average difference between a given model’s tone and the average tone of all the models in their answers to the same question (only answers with a mention of Poland and a tone score, in questions that at least five models answer in this way); the tone range is calculated on the 655 questions with six such answers. “Risk” questions (281) follow the classification from the blind round described in chapter 1 (10 disputed questions counted as other questions). Shortlist: 96 questions × 6 models (30 September 2026), a separate denominator. Recommendation rate per model in both sets is given in the annex only (length cap on answers at four models; truncation simulation). Differences below 3 percentage points for mention, 0.5 points for tone and 5 percentage points for recommendation rate are not interpreted. The measurement concerns the model’s answer as returned by the API.

CHAPTER 10

What this means for the economic promotion of Poland

IN BRIEF

- Language models (LLMs) suggest Poland to companies almost always: in the shortlist it appears in 98.4% of the 576 answers to questions that do not name it, and it is explicitly recommended in 82.5%. In this measurement, the problem for economic promotion is not presence but role: in questions about sales markets, incentives, service centres and engineers, Poland is recommended in 90–93% of answers; in questions about contract manufacturing, in 69.4%; about a warehouse for Europe, in 62.3%. The goal is not "for the models to recommend Poland" but the decision area in which Poland is to be a candidate rather than a backdrop.
- Poland loses on geography, not on cost. In logistics and industrial answers, cost is its argument in every answer; what it loses on is transit time to the west and its reputation for maturity. Its competitors there are Germany, the Netherlands, Italy and Austria, not Czechia and Hungary. The missing argument is proof of maturity (certificates, references, names), not proof of growth.
- A company that asks for names gets a Polish company in two answers in three (99 of 150), but when it asks for certified subcontractors, in only 5 of 30. The "Polish companies in use" are headed by registers and portals (Krajowy Rejestr Długów, a Polish debt register; BIG InfoMonitor, a credit-information bureau; PSE, the Polish transmission-system operator; No Fluff Jobs), while aviation, electronics and pharmaceuticals, the sectors Poland presents to investors as strategic, appear rarely even on request. The canon on request (KGHM, PGE, PKO) does not appear in practical answers.
- Polish institutions are present in the models' answers, but search in the market's language rarely leads to them. PAIH, the Polish Investment and Trade Agency, is among the three most frequently cited domains in each of the 12 markets; the URLs given in the answer by the models themselves lead to Polish sites in 54.5% of cases, sources from search in 25.0%. In questions about incentives, PAIH appears in 26 of 27 answers with URLs given in the answer and in 3 of 12 with sources from search. The only bilateral chamber in the answers is ahk.pl, and only in one of the twelve markets.
- Two threads show the limits of promotion: the models add the ferry across the Baltic to answers for Swedish companies of their own accord (87 of 348 answers, no question about the ferry), whereas the China–Europe railway appears for Chinese companies mainly when the company asks about it (77 of 98). An argument the model has linked to a decision situation enters the answer by itself; an argument that is merely an association has to be brought in by the question.
- The war does not cost Poland its recommendation in any of the 864 answers, but it changes its content: two thirds of the answers about investment, nearshoring and insurance present Poland as more exposed than the countries it is compared with, and nine in ten of those still recommend it. The models already have arguments that promotion does not need to invent: the European Union (86.9% of answers), energy (75.0%), "the risk can be priced and insured". They lack one: Poland as a base for serving the reconstruction of Ukraine needs no promotion (recommended in every answer); Poland as a safe investment when the question does not name it does.

10.1 What this chapter can say, and what it cannot

Part II measures the answers of six models to 700 company questions from 12 markets (4,200 answers of the study) and to the 96 shortlist questions that do not name Poland (576 answers), collected through the APIs between 25 August and 30 September 2026, as well as two thematic measurements: Poland on request (companies and industries) and Poland in the shadow of war. It is a map of the representation in a dated study, not a forecast of investment decisions and not an assessment of the work of any institution. The conclusions below concern one thing: how Poland looks on this map when a company asks, and what follows for the institutions that promote it economically, as well as for Polish companies that want to appear in the answers by name. The institutions in question are the Polish Investment and Trade Agency (PAIH), ministries, local governments and economic zones, chambers and industry associations.

10.2 The area, not the country

The most important finding of Part II is simple: for companies, what decided everything in Part I does not work. Whether the question is about the region or about Europe in general does not change Poland's visibility (99.6% against 97.7%; chapter 2). What changes it is the decision area, and not in whether Poland appears but in whether it is recommended and whom it stands alongside. In questions about sales markets, incentives, service centres and engineers, Poland leads a regional list with Czechia, Romania and Hungary beside it, and is recommended in 90–93% of answers. In questions about suppliers, contract manufacturing and logistics, the list changes to Germany, Italy, Austria, Portugal and the Netherlands, and the recommendation falls to 80.0%, 69.4% and 62.3% (chapter 4).

For promotion, this means a change of unit. The question "do the models recommend Poland to companies?" has no single answer; the meaningful question is "in which decision area is Poland to be a candidate, and in which is it a backdrop today?". The content meant to reach answers about sales markets can be the content Polish institutions already have, and it already works. Content for the industrial areas is different content: about transit times, certificates, references from Germany. In the regional areas, Poland wins with an argument the models have; in the industrial ones, it loses for want of an argument they do not have.

10.3 Poland loses on geography, not on cost

The refrain of Part II came from reading the 29 logistics answers in which Poland was not explicitly recommended (chapter 4). In 18 of them, the company asks about a single warehouse for the whole of Europe and is given the Netherlands or Germany, because from there deliveries reach the west in two to three days; Poland is the "cost option for Central and Eastern Europe", with a caveat about three to four days to Spain or Portugal. In the remaining 11, a company from a neighbouring country is given its home country as the location and Poland as a market to be served from there. Cost is Poland's argument in every one of these answers. What loses is time and role.

Three things follow for those who promote Polish locations. First, it is hard to compete with Venlo and Duisburg for the role of a pan-European hub in answers that count the days of transit to Lisbon: the models will count them the same way (the report describes the models' reckoning, not logistics costs). Second, there is a point in competing for the role of "second warehouse" and "eastern hub": this is exactly how the models themselves describe western Poland when a company asks about Central Europe (chapter 7, France and the United Kingdom), and that description is worth reinforcing with content about transit times to Berlin, Prague, Vienna and Scandinavia. Third, with the neighbours Poland loses not to the Netherlands but to Brno and western Slovakia, because the model

recommends the company its home country; here the argument is not cost but reach: southern Poland as a location that also serves Czechia and Slovakia, not the other way round.

10.4 In the industrial areas, Poland needs proof of maturity, not of growth

To suppliers, the models describe Poland as a base that is growing: that is how 9 of the 14 answers in which Poland is not explicitly recommended sound, and their point of reference is Germany and Italy, described as reliable (chapter 4). When a company asks for a list of certified Polish subcontractors, a Polish company appears in 5 of 30 answers, even though 28 of them give some names (chapter 5). When it asks for names in a specific industry, it gets a Polish company in 99 of 150 answers: defence 18 of 18, furniture 11 of 12, food retail 10 of 12, but component suppliers 5 of 30.

In the models' answers, "growing" describes a country that has yet to prove itself; "reliable", a country that already has. Industrial promotion that wants to move Poland from the first group to the second does not need more content about growth. It needs content that is proof: lists of certified subcontractors with names and standard numbers, references from customers in Germany and Scandinavia, data on punctuality and quality, presented the way a procurement manager asks for them, in English and in German, at an address the models associate with Poland. The defence industry shows that this can work: it appears with names in every answer about the defence market, and it has names, contracts and documents online. Component suppliers hardly appear by name in the answers (5 of 30); whether their material is missing online, the study did not check.

10.5 Companies on request: the canon does not get in, the registers do

Asked directly about the largest Polish companies, the models give the canon: KGHM in 144 of 144 answers, PGE and PKO Bank Polski in 143, PZU in 139, Orlen and Allegro in 131 (chapter 5). In the answers to company questions this canon barely exists: KGHM in 21 of 4,200 answers, PKO in 15, PZU in one. The Polish companies "in use" are headed by names that do not appear in any answer on request: Krajowy Rejestr Długów (143 answers), BIG InfoMonitor (97), PSE (95), No Fluff Jobs (71), Hays Poland (53). This is the same activation gap as in Part I, only with companies instead of Chopin: a name enters a practical answer when it is linked to the situation the company is asking about: vetting a business partner, recruitment, connection to the grid.

For Polish companies that want to appear in the answers by name, the conclusion is concrete: a model does not bring a company in because it is large, but because it answers the question. KRD appears in 143 answers not as "a Polish company" but as the place where a business partner is checked. A company that wants to be in the answers for foreign procurement managers needs content that answers their decision situation, not the question "who are we?". For institutions promoting strategic sectors, the conclusion is harder: aviation, electronics and pharmaceuticals have neither companies in the answers (0, 0 and 4 of the 144 answers about the largest companies) nor industries (24%, 17% and 7% of the answers about the strongest ones). Here there is nothing to activate, because there is nothing to link. The entities have to be built first: names, places, contracts, described in the markets' languages.

10.6 Polish institutions are present in the models' answers, but search in the market's language rarely leads to them

This sentence recurs in both parts and is the most practical observation of this study. The URLs that the models give companies themselves lead to Polish sites in 54.5% of cases, including 35.0% to official ones; sources from live search do so in 25.0% and 6.9% (chapter 6). PAIH is among the three most frequently cited domains in each of the 12 markets, but in questions about investment incentives it appears in 26 of 27 answers with URLs given in the answer and in 3 of 12 with sources from search.

When the answer is built from what the model has, the Polish agency is almost always in it; when it is built from what search in the market's language finds, it is usually absent.

Three conclusions. First, citability: content meant to reach answers based on search has to be in the market's language or in English, answer the company's question (incentives, procedures, costs, deadlines) and sit at an address the models already associate with Poland, such as the English-language PAIH site, which the models cite of their own accord (chapter 6). Second, chambers: ahk.pl is the only bilateral chamber among the top three domains, and only in the German market; it appears in 52 of 360 answers for German companies, although no question asks for a chamber. Chambers exist for each of the twelve markets; one appears in the answers. Third, registers: KRS (the National Court Register), KRD, biznes.gov.pl and podatki.gov.pl are cited in every market, because they answer specific questions. This is capital that economic promotion has for free and does not need to build, only to avoid spoiling: the addresses must remain stable and the English versions complete.

10.7 The ferry is an argument, the railway an association

Two market-specific threads from chapter 7 show the difference between an argument and an association. The ferry across the Baltic appears in 87 of 348 answers for Swedish companies, although no question asks about the ferry. The models add it themselves when a company asks about a warehouse, deliveries and transit time; in Part I the same ferry was an association the Swedes had with Poland that barely featured in the answers for individuals. The China–Europe railway through Małaszewicze appears in 98 of 348 answers for Chinese companies, but in 77 because the company asked about it itself; outside them it appears in 21 of the 270 answers. The models have the ferry linked to a decision situation; they have the railway linked to Poland, but not to a decision.

The rule is the same one that Part I called "doors": an entity enters a practical answer only when it is linked to the entity the question is about. For the promotion of routes, ports and corridors, this means that it is not enough for a model to "know" about Małaszewicze. It has to have them linked to the question a Chinese company asks: "a warehouse for Germany and Scandinavia, rail or sea, bonded warehouse or not". The link has to come with times, costs and procedures, described in Chinese and in English. The same applies to every argument Poland regards as its own: it is in the answer when it has a door.

10.8 The war: Poland recommended with conditions that the models no longer need to learn

The thematic measurement from October (chapter 8) brought no bad news: in none of the 864 answers to questions naming Poland did the models advise against it as a choice because of the war, and as a base for serving the reconstruction of Ukraine they recommend it in every answer, including when the question does not name it. It produced four conclusions.

First: the reconstruction base is a door already open. Here Poland does not need promotion; it needs content that fills this role (partners, procedures, contracts), because these are what Ukrainian companies ask about (chapter 7).

Second: with safe investment, when the question does not name Poland, this is not the case. When a company asks about safe countries in the region, Poland always appears, but in 37.5% of answers it is placed in the group of riskier countries, because on this scale the models order countries by distance from the front line. Communication about investment security does not need to convince the models that Poland is safe. They say so themselves. It has to give them arguments on a European scale, and these the models already have: the European Union and its funds (86.9% of answers), energy (75.0%), NATO (51.1%).

Third: the risk the models talk about is risk as perceived by a management board, not an insurer's valuation. Two thirds of the answers present Poland as more exposed than the countries it is compared with, and nine in ten of those same answers still recommend it. This is a matter of communication, not of facts: the models' answers themselves propose the formula "the risk in Poland can be priced and insured", and institutions that hold data on the cost of political-risk insurance and on guarantees have a ready argument to put forward in the markets' languages.

Fourth: tone depends on the question, not on the market. Of the 700 company questions in the study, 281 themselves ask about risk, and it is in these that the answers are cool (6.16 against 7.20), and this holds for each of the six models (chapters 1 and 9). Content about risk reaches companies' decisions whether or not Poland prepares it; it is better that Poland does, because work-permit legalisation, energy costs and regulatory change are the topics on which the models build their "but".

10.9 Measure, do not audit

Part I showed that after the study ended, newer versions of the models shifted the share of explicit recommendations of Poland by a dozen or so points within seven weeks, in both directions. Part II rests largely on a single run of 12 questions per area and on a measurement of companies without a second judge and without repetitions (chapter 2). Both lead to the same conclusion as in Part I: the image of Poland in the answers for companies is a flow, not a state. Each thesis (two maps, geography rather than cost, PAIH in search) requires control measurements at regular intervals, for example quarterly, with the same set of questions, markets and models, because the models are updated and the sources change. A one-off audit says how things were. An institution that wants to know whether its content has reached the answers for companies needs a measurement that can be repeated, together with the recorded text of every answer, so as to know not only whether Poland appeared, but in what role.

HOW WE COUNT (CHAPTER 10). All figures come from chapters 1–9 of Part II and from Part I; the chapter introduces no new calculations. The figures for the shortlist (96 questions, 576 answers) and for the study (700 company questions, 4,200 answers) have separate denominators and are not combined; the thematic measurement "Poland in the shadow of war" (October 2026) is a separate set with different model versions. The chapter does not use the recommendation rate per model (annex). The measurement concerns the model's answer as returned by the API.

About the tool and the author

The measuring instrument. Every answer in this report was collected and assessed by a single research application, Semantio.pro. It puts questions to language models (LLMs) through their official APIs, in the language and the hypothetical situation of those asking; it records each answer as the provider returns it and assesses every answer with the same judge, against the same rules. Each of the 1,400 questions was asked on behalf of a specific identity created for the study: a persona from one of the 12 markets who is planning a trip, a course of study, a move or a business decision. The questions were drafted by the research application's question generator, which wrote them in the language of each market from a description of that market and its personas, separately for each decision stage. The same measurement can be repeated on the same questions a week later, a quarter later or after a change of model version. That is what distinguishes it from a single conversation with a chatbot, which can be neither repeated nor counted.

Semantio.pro was built to measure company brands. In this report the brand is Poland, and the method is the same: a dated measurement window, control layers (a repeatability test, a second judge, human verification of factual errors, a newer-model-version test and an open portrait) and thresholds below which we do not interpret differences. The details are in the annex.

What the tool does with this report's concepts. The report introduces a handful of concepts, because without them the results could be neither described nor improved. In Semantio, each of them is a measurement, not a metaphor.

- *Brand semantics*: which entities the models link to the name, and in which situations they mention it. An entity is a named thing that can be recognised in an answer regardless of language or inflection: a person, a place, an institution, an event, a product; for Poland, for example, Chopin, Kraków or the Tatras. Brand semantics adds up to a *semantic map*: market by market, question type by question type, source by source. The twelve portraits in chapter 8 are such a map for Poland.
- *Activation gap* and *entity linking* show which doors are open and which are closed: what the brand holds in its canon, and what actually makes it into practical answers. This is diagnostic work: it shows what needs attention before anyone writes the first word.
- *Citability* shows which sources the answers offer those asking, and in which language; it reveals whether the brand's own websites play any part in the answer at all.
- *Semantic drift* is the reason such measurement cannot be a one-off. Within a few weeks, new model versions can reverse the order of the models by how explicitly they recommend a brand (chapter 7); a one-off audit describes a version that, a quarter later, will no longer be answering anyone.

Work on how a brand appears in model answers has a name in the industry: generative engine optimization (GEO). We prefer to start with brand semantics, because that is what can be measured, and only what has been measured can be improved.

About the author. Grzegorz Miłkowski is the creator of Semantio and of the research methodology the tool applies, and CEO of Brand Semantics P.S.A., a brand semantics think tank that advises boards

on how their brands are understood and recommended by language models (so-called AI models). The author of this report is responsible for its method, numbers and conclusions.

Contact. If you would like to see a similar measurement for your own brand, sector, region or country (across the same or other markets, in a similar or a different format) or to discuss how we can support your business in the world of generative language models: hello@brandsemantics.eu · www.brandsemantics.eu.



Commissioned research, training, collaboration
Katarzyna Bienussa
Head of Market Education, Brand Semantics P.S.A.
k.bienussa@brandsemantics.eu

Methodological annex

Scope and measurement channel. The study covers six generative AI models queried through their official APIs: ChatGPT, Gemini, Claude, Perplexity, Grok and DeepSeek. We measured the language model's (LLM's) answer as returned by the provider's API: a specified model version, with no personalisation, no conversation memory, no account in the provider's app and no memory of earlier questions (each question was sent in a separate, one-off call), and with a maximum answer length set by the research application. The cap was 1,500 tokens for ChatGPT, DeepSeek and Perplexity, 1,024 tokens for Claude (in the main study, the repeatability test and the B6 measurement; 4,096 in the version test, the open portrait and the thematic measurement) and 4,096 tokens for Gemini, while Grok had no cap. Answers that reached the cap break off mid-sentence (see "Truncated answers"). We sent no temperature parameter to any model (the provider's default setting applied). Four models received a system instruction: ChatGPT, Grok and DeepSeek a single general sentence asking for sources and URLs, and Perplexity a single general sentence asking for an accurate, well-documented answer; Claude and Gemini ran without a system instruction. The provider's safety filters were switched off for Gemini ("block none"). Only Perplexity searched the web live; the other five models answered from their own knowledge, and every URL in their answers was written by the model itself (see "Sources and citations"). We recorded the answer text verbatim: for Claude in the main study, the repeatability test and the B6 measurement, as the first text block returned by the API (in a call with no tools and no thinking this is the only block, so the content is complete); from the version test (B8) onwards, as the concatenated text blocks; for Gemini, as the first text block of the answer. We found no resulting loss of content in Gemini's data, although the way the text was recorded does not rule it out. The judge and all automated procedures worked on exactly this recorded text. The study does not reproduce the screen seen by a person using the app (its layout, cards, tables, interactive elements or app-side personalisation). Throughout the report, the subject is the measurement of the model's answer via the API, not what a person using the app sees. The results are probabilistic: a language model may answer the same question differently, so the report gives shares of thousands of answers within a dated window, not a forecast of what a model will say in a single conversation; the spread between repetitions is measured by the repeatability test.

Study material. 1,400 questions: 700 asked from the perspective of individuals (B2C) and 700 from the perspective of companies (B2B). The questions cover 12 markets: China, Czechia, France, Germany, Italy, Norway, Slovakia, Spain, Sweden, Ukraine, the United Kingdom and the United States. Each was asked in the language of its market, which means 11 languages (English for both the United Kingdom and the United States). The questions were assigned to the three stages of the decision path: awareness, consideration and decision. Of all the questions, 82.7% name Poland or a Polish city or region. We count as naming Poland any question that contains the name of Poland, the adjective "Polish" or the name of a Polish city or region, in any grammatical form of the question's language; we checked the classification against the full list of questions. We supplemented the automatic rule with an explicit list of 12 questions containing Polish names in inflected forms that it did not recognise (e.g. "pri Krakove", "Schlesien", "Masuren", "Люблiнi"), and then amended it so that the short form "posen" is recognised only as a whole word, because two French questions contained it inside other words and had been wrongly counted as naming Poland. These two questions were classified blind,

by the same procedure as the rest, by two models (Claude Opus 5.5 and Claude Sonnet 5.5): both classified them as general questions, consistently with each other and with the seven control questions. Of the awareness-stage questions, 242 do not contain the name of Poland, and it is on these that we measure generative visibility. Each question was asked once of each of the six models. This gives 8,400 answers, all of which were assessed.

Personas. The questions were asked by personas: statistically constructed scenarios of individuals and companies created for the study, not real people or companies. For each of the 24 sets (12 markets × individuals / companies) a market card was drawn up in the language of the market: five areas of analysis, four personas described by role and decision (without demographic characteristics and without weights), five comparison countries and a reference domain (poland.travel for individuals, paih.gov.pl for companies). The card was used to write the questions; the model being measured never received it, only the text of the question. Examples of personas from the cards: in Germany, someone planning a weekend away or a city break, and someone managing a production or logistics site; in the United States, someone researching Polish family roots, and a member of a company's management planning expansion in Europe; in Ukraine, someone weighing up a temporary stay and work in Poland, and an exporter entering the Polish or the wider EU market; in China, someone considering studying in Poland, and someone responsible for procurement or manufacturing cooperation; in Czechia, someone visiting Poland across the border, and someone in charge of exports to the Polish market. The full list of personas for all markets (original and translation) is published in the open data of the report; "home country" in the text and "Home country: Name" in the cards and tables denote the persona's country.

Origin of the questions. The 1,200 questions of the main measurement were produced by the research application's question generator in two steps: first, a model with search access (sonar-pro) wrote a short description of the market based on the card, covering how the personas on that card look for information and make decisions, and then gemini-2.5-flash wrote the questions in batches of six, in the language of the market, separately for each decision stage, from the perspective of the assigned persona, phrased either as questions or as instructions to an assistant (78 of the 1,400 questions are instructions, with no question mark). The 200 questions added on 25 September come from a prepared supplementary list covering thinly covered areas. No one reviewed the questions before the measurement. After the measurement, all 1,400 questions underwent a quality review against fixed criteria: language, fit with the segment and the decision stage, steering towards Poland or towards an evaluation, false assumptions, realism, repetition within the set, and sensitive content. The review was carried out by models from a different provider than the generator (Claude Opus 5.5, twelve independent runs, one per market), with the key assessments independently repeated: blind regrouping of repetitions by other runs (a pair was counted as a repetition only if both runs identified it or if its text similarity was at least 0.90), and a blind second assessment of all questions flagged as serious and of a random sample of 10% of the rest (agreement 85%). The author blind-checked a sample of 24 questions (92% agreement with the review, with no differences at the "serious" level) and confirmed every question flagged as serious. The review established two things. First, 317 questions (22.6%; 36% in the company segment, 10% in the individuals segment) repeat the scenario of another question in the same set in different wording. They form 172 groups, in 105 of which the same question carries different decision stages; all of them come from the generator, none from the 200 added questions. Second, 16 questions (1.1%) ask for arguments in favour of Poland (10, e.g. "justify why Poland would be the best choice") or are labelled as decision-stage although they ask for a topic overview or a reading list (6). None of the 242 questions that do not name Poland steers towards Poland to a serious degree. We do not change the set: all 1,400 questions remain in the frozen data and in every number in the report. We did, however, compute a control variant that drops the 16 serious questions and keeps one question from each repetition group (1,071 questions, 6,426

answers): mention rate 94.0% against 95.1%, tone 7.01 against 6.95, recommendation rate 70.2% against 69.6%, recommendation rate in questions that do not name Poland 57.2% against 55.9%, generative visibility 71.5% against 71.3%. All the differences fall below the thresholds of measurement variability (3 percentage points for mention rate, 0.5 points for tone, 5 percentage points for recommendation rate); the same holds for each of the six models (the largest difference being 1.8 percentage points of recommendation rate) and for two stricter variants: weighting each question in a group by the inverse of the group's size, and additionally excluding 28 questions that ask solely for Poland's advantages. Repetitions mainly affect company questions, so for Part II we also checked the table of decision stages: tone and recommendation at each stage change by less than the threshold, and the only number beyond the threshold is the share of risk questions at the decision stage (43.0% in the full set, 35.6% in the control variant), which Part II flags at that table.

Measurement window. The main measurement ran from 25 to 29 August 2026 (1,200 questions). Supplementary runs and the 200 additional questions were carried out from 25 to 28 September 2026. The last answers in the study date from 28 September 2026. The control and additional measurements (the repeatability test, the shortlist, the open portrait, the newer-versions test, the thematic measurements, search mode and the control measurement of position on 2 November) ran from 30 September to 2 November 2026; each is described separately and none changes the figures of the main study.

Model versions.

- ChatGPT: GPT-5.6 Sol, the model on the paid plans (Plus/Pro). From August 2026 people on the free plan were served a smaller model of the same generation (GPT-5.6 Luna).
- Gemini: gemini-3.5-flash, the default model of the research application in the study window (1,252 answers). 148 answers in the main measurement (25–29 August) were served by the fallback model gemini-2.5-flash, which the research application uses when the primary model is temporarily unavailable.
- Claude: claude-sonnet-4-6, the previous model generation available through the API. By then the consumer app was already serving the Sonnet 5 generation.
- Perplexity: sonar-pro.
- Grok: grok-4.3.
- DeepSeek: queried under the alias deepseek-chat, which pointed to DeepSeek-V4-Flash until 9 September 2026 and to V4.1-Flash from 10 September. 1,200 answers (main measurement, 25–29 August) come from V4-Flash and 200 (additional questions, 26 September) from V4.1-Flash.

The answers from the different Gemini and DeepSeek versions relate to different sets of questions, so we do not attribute differences between them to the model version. The changes to Gemini models of 8–9 October 2026 do not affect the study data.

Automated judge. Every answer was assessed by the automated judge in country mode v2.1 (model gemini-2.5-flash). The judge determined:

- the type of mention of Poland: country, city or region only, no mention;
- the tone of what is said about Poland, on a scale of 1–10;
- the recommendation level (high / moderate / low / none);
- other countries and names mentioned;
- whether claims are consistent with the facts.

The judge knows the world only up to the date of its training and does not know about events in 2025; in the thematic measurement B14, 7 of the 10 answers it flagged as contradicting the facts were

correct (e.g. Russian drones over Poland in September 2025). Every error flag is therefore verified by a person, and the report gives only confirmed errors.

Gemini in three roles. The model gemini-2.5-flash played two auxiliary roles in the study and one as a model under study: it wrote the questions in the generator (see "Origin of the questions"), assessed all the answers as the judge, and, as the fallback model, also supplied 148 Gemini answers when the primary model, gemini-3.5-flash, was temporarily unavailable. Three arguments suggest that this does not distort the results. The generator knew none of the models' answers. It wrote the questions from the market card before the measurement, and all models received the same questions. A judge from another provider (claude-sonnet-4-6; see "Second judge") preserved the order of the models on tone and the extremes on recommendation rate (Gemini first, Claude last), and for Gemini's answers the difference between the two judges was the smallest of all the models (0.08 points). We also checked the key result without the 148 answers from the fallback model: the difference between Gemini's tone and the average of all models remains unchanged (+0.51). What cannot be ruled out is that questions written by one provider's model are phrased in a way that the same provider answers more freely; we therefore describe Gemini's lead on recommendation rate as a feature of the August 2026 version, not as a ranking result.

Metrics.

- Mention rate: the share of answers in which Poland or a Polish city or region appears. The denominator is all assessed answers.
- Tone: the mean score on a scale of 1–10, calculated only for answers that mention Poland.
- Recommendation rate: the share of answers mentioning Poland in which the model explicitly recommends Poland (the high level). The headline value is given only for questions that do not name Poland (1,051 answers with a mention). A question that itself points to Poland inflates the recommendation rate.
- Generative visibility (GV): the share of answers that name Poland in response to awareness-stage questions that do not name Poland. For B2C this means 174 questions, or 1,044 answers. Per-market results are published for B2C only, because in B2B the study set has fewer than 8 such questions per market. Visibility in business questions is measured by the separate B6 measurement (Part II).
- Hint type in the question (for GV): awareness-stage questions differ between markets in how strongly they steer towards Poland. We assigned each of the 174 B2C questions to one of four types:
  - points to a region that includes Poland (92 questions): Central or Eastern Europe, the Visegrad Group, "new EU countries", and also "neighbouring countries", "nearby abroad" or a border town, where the persona's home country shares a land border with Poland (Czechia, Germany, Slovakia, Ukraine) or the question names the Baltic explicitly; being neighbours across the sea, without the Baltic being named, does not qualify;
  - concerns a topic strongly associated with Poland (13 questions): e.g. Second World War memorial sites, dental treatment abroad;
  - is general (54 questions): Europe, the EU, "abroad", "outside Western Europe";
  - points neither abroad nor to a region (15 questions): e.g. "which regions are good for a short break", "the culture of this country".

  If a question meets both the region and the topic conditions, region takes precedence. The classification was based solely on the text of the question and the persona's home country. The author prepared the first version; two independent runs, carried out by two different models on the

question list alone, without sight of the study results, then classified the same 172 questions according to the definitions above. The two runs agreed with each other on 98.8% of questions (170 of 172) and with the author's version on 90.1%. For the 16 questions on which both runs agreed on a type different from the author's, the runs' type was adopted; for the 2 questions on which the runs differed, the author's type was kept. The two French questions that entered the set after the correction of the "Poland named in the question" flag were classified in a separate blind round by two models (both: general). GV is reported together with its distribution across these types, and markets are compared within the same type. Questions that point neither abroad nor to a region remain in the GV denominator, because we do not change the definition of a metric after seeing the results.

The business questions were classified by the same procedure: 68 study questions that do not name Poland and 96 questions from the B6 measurement (164 in total). The two blind runs agreed on 97.6% of questions (160 of 164), and with the author's version on 90.2% and 92.7%; for 12 questions the runs' agreed type was adopted, and for the 4 disputed ones the author's type was kept. One rule was clarified for business questions: a named region that is the target market ("a warehouse serving customers in Central Europe") does not count as a region hint if the location the persona is asking about remains open ("where in Europe"). Result: in the study, 35 questions point to a region, 29 are general and 4 do not point abroad; in B6, 39 point to a region and 57 are general. None of the business questions was of the topical type.

- Recommendation gap (RG): in answers to questions that do not name Poland, the difference, in percentage points, between the share of answers naming Poland and the share of answers naming the most frequent alternative country.
- Share of Polish sources (SS): the share of Polish domains (official and other) among citations. Citations come from two modes, which we do not combine into a single inferential figure: "sources from search" (the list of search results returned by Perplexity's API) and "URLs given in the answer" (URLs that the other five models wrote in their text; ChatGPT, Grok and DeepSeek at the request of the system instruction, Claude and Gemini without one). The share of Polish sources shows where an answer points for further detail, not where the model got its knowledge from. We do not compare the number or share of citations across models, because these measure the instruction and the mode, not the model.
- Prominence of Poland in the answer (position and rank): see the separate section "Position of Poland in the answer".
- Model divergence (ED): the deviation of a given model's tone from the mean tone of all models in answers to the same question. Comparing answers to the same questions removes the effect of question selection. Chapter 7 also reports the tone range, the difference between the warmest and the coolest model on the same question (1.8 points on average), and sets it against the range that measurement noise alone would be expected to produce (the standard deviation of tone across repetitions is 0.37 points; see the repeatability test); the observed range is about twice as large.
- Home country as an alternative: the share of answers to questions that do not name Poland in which the persona's home country appears (21.1% overall, highest in Czechia at 64.1%). We do not read it as "advising against the trip": most often the question does not refer to travelling abroad (e.g. Norwegian and Swedish questions with no hint of going abroad), or the model uses the home country as a point of reference (e.g. German questions about neighbouring countries).
- Domains counted per answer (the list of domains in chapter 5; the market cards in chapter 8; chapter 9): we count a domain once for each answer in which it is cited at least once, and give the

number of such answers under the label “answers with a citation”; the share of Polish sources (SS), by contrast, is counted per citation. The two measures are never shown in the same table.

- Thresholds: differences in tone of less than 0.5 points, in mention rate of less than 3 percentage points and in recommendation rate of less than 5 percentage points (for groups of at least about 150 answers) are treated as falling within measurement variability and are not interpreted. The repeatability test provides the justification. We read the recommendation level as a comparison between models, question types and markets, not as an absolute percentage, because it depends on the threshold the judge applies (see “Second judge”).

Baseline measurement and five control layers. The study is a dated baseline measurement (August–September 2026) with five control layers:

1. repeatability test,
2. second judge on a model from another provider,
3. human verification of factual errors,
4. newer-model-version test,
5. open portrait.

The control measurements were run in copies of the research projects, so the study set was left untouched. Answers from the control measurements do not feed into the study metrics. Separately, and outside the control layers, we carried out the thematic measurement “Poland and the war in Ukraine” (B14), described below.

(1) Repeatability test (B2, 30 September 2026). We selected 29 questions from three sets: B2C questions for Germany (10), B2B questions for the United Kingdom (9) and B2C questions for Ukraine (10). On a single day we put them again to each of the six models three times, at intervals of at least 68 minutes, clearing the research application’s cache in the process. This produced 522 answers, assessed by the same judge v2.1 with the same answer length caps as in the study.

- The type of mention of Poland was identical across all three repetitions for 96.6% of question–model pairs. Claude came out lowest (89.7%), and ChatGPT, Gemini and Grok highest (100% each).
- The standard deviation of tone across the three repetitions averaged 0.37 points (from 0.30 to 0.49, depending on the model).
- The recommendation level was identical across all three repetitions for 70.1% of pairs. The share of answers with a high recommendation level in the three repetitions was 61.5%, 62.1% and 59.2% (a range of 2.9 percentage points). On recommendation rate among answers with a mention, Gemini is highest in each of the three repetitions and Claude lowest in two of them and overall (in the first repetition Perplexity is lowest: 48.3% against 50.0% for Claude).

We treated the study answer as a fourth measurement. The mention type matched across all four answers for 95.4% of pairs. The recommendation level in the repetitions matched the study answer in 79.9% of cases. The mean tone of the repetitions differed from the tone of the study answer by 0.52 points on average.

Conclusions:

- differences in tone of less than 0.5 points, in mention rate of less than 3 percentage points and in recommendation rate of less than 5 percentage points are treated as measurement noise;
- results are not compared across measurement windows, because the difference between August and September is of the same order as the threshold;

- the recommendation level is reported only as a distribution within a model or a market, never for a single question;
- the lists of countries named alongside Poland, the most variable element of an answer, are reported only as aggregates.

(2) Second judge (B1, 29 September 2026). We re-assessed a random sample of 226 study answers (37–38 per model; two further answers in the draw were also answers with an error flag and are counted in the error verification) with the same judge prompt but on a model from another provider (claude-sonnet-4-6). Results:

- the type of mention of Poland matched in 100% of answers;
- the main judge rated tone 0.22 points higher on average than the second judge, which is below the 0.5 threshold;
- the recommendation level matched in 78.3% of answers. The divergence ran in one direction only: the second judge sets a higher bar for a "high recommendation level" and assigned the high level to 38.9% of answers where the main judge assigned it to 57.1% (40 answers moved from the high level to moderate, none the other way). For both judges the extremes of the order by recommendation rate are the same (Gemini highest, Claude lowest), while the models in the middle of the field swap places; the order of the segments is the same; per-model values are given in the section "Recommendation rate by model". This is why recommendation rates in the report are read comparatively, not as an absolute percentage of answers recommending Poland.

For Gemini's answers the difference between the judges was the smallest of all the models (0.08 points). There is therefore no sign that the main judge, although built on a Gemini model, favours Gemini's answers, or that the second judge favours Claude's (for these the main judge rated tone 0.17 points higher).

(3) Human verification of factual errors (14 October 2026). The judge flagged 19 answers (0.23% of those assessed) as containing a claim that contradicts the facts; we confirmed 14 of them. Every one of these cases was verified by a person against current sources. We count as a factual error any false claim in an answer, including one about another country (one of the 14 errors concerns Croatia), because it is delivered as part of an answer to a question about Poland.

- 14 factual errors were confirmed: 10 in B2C questions and 4 in B2B. We did not check a random sample of unflagged answers, so the judge's sensitivity is unknown.
- In 3 cases the judge was wrong. Once this was due to the model judge's outdated knowledge (China's accession to the Apostille Convention in November 2023); twice the judge flagged correct sentences as errors.
- 2 cases were not factual claims about Poland (a wrongly matched source, and a sentence that was technically true but misleading) and are not counted.

Claims that the judge marked as "unverified" (85) are not treated as errors. Ten of the 14 confirmed errors occur in answers truncated by the length cap (see "Truncated answers"), but in none of them is the assessed claim a truncated sentence (between 106 and 5,096 characters separate it from the cut point). A truncated answer may, however, have lost a later qualifying sentence; we therefore do not attribute the errors to truncation, but we do not rule out that the full answer would have been more cautious. Confirmed errors per model are given without a ranking, because the numbers are small: ChatGPT 0, Gemini 2, Claude 4, Perplexity 1, Grok 1, DeepSeek 6.

(4) Newer-model-version test (B8, 14–15 October 2026). After the measurement ended, four providers changed their production models; Perplexity and Grok continued to be queried in their study versions. We put the same questions to the new models: ChatGPT gpt-6-luna (the model behind free

ChatGPT), Gemini gemini-3.5-flash-lite, Claude claude-sonnet-5-5 and DeepSeek deepseek-flash. The test covered:

- a random sample of 20% of the study questions (283 questions from all 24 sets: 12 markets × B2C/B2B);
- 24 business questions that do not name Poland;
- 1,228 “new answer – study answer” pairs in total;
- in addition, the 29 questions of the repeatability test (348 control answers, run in the same project copies as the repeatability test).

Results:

- the type of mention of Poland matched in 96.6% of pairs;
- tone differed by −0.07 points on average;
- the recommendation level matched in 76.5% of pairs.

For mention and tone this matches the old model’s agreement with itself after a month (96.3% matching mention types), and the differences in tone between versions (from −0.26 to +0.19 points, depending on the model) fall within measurement variability. The recommendation rate behaves differently. Pair agreement (76.5%) is admittedly close to the old model’s agreement with itself (79.6%), but the shifts are not randomly distributed: in the same 283 questions the share of answers with a high recommendation level changed from 60.1% to 47.7% for ChatGPT (−12.4 percentage points), from 47.0% to 61.8% for Claude (+14.8), from 66.4% to 62.9% for DeepSeek (−3.5) and from 80.9% to 79.9% for Gemini (−1.0). Two of the four changes exceed the 5-percentage-point threshold, and the order of the models by recommendation rate is reversed between versions. The report therefore describes the order of the models by recommendation rate as a feature of the August 2026 version, not of the model brand; mention rate, visibility and tone are robust to the change of version. The answer length caps in the test were the same as in the study for ChatGPT, DeepSeek and Gemini, and higher for Claude (4,096 compared with 1,024 tokens). In pairs where Claude’s study answer was not truncated (129 pairs), the rise in recommendation rate is similar (from 45.0 to 60.5%), so it stems from the version change. In pairs with a truncated answer (154 pairs; from 48.7 to 63.0%), part of the rise may have been due to the length cap (see “Truncation simulation on full answers”). The lists of countries named alongside Poland changed most: list similarity was 0.51, compared with 0.61 for the old model repeated after a month. Comparison countries are therefore treated as tendencies, not as fixed rankings. The new models are as repeatable as the old ones: in both versions the mention type matched across repetitions in 97.7% of cases. The test answers do not feed into the study numbers.

(5) Open portrait (B9, 14–15 October 2026). This is a separate, dated measurement on the production models current on the day of measurement: ChatGPT gpt-6-luna, Gemini gemini-3.5-flash-lite, Claude claude-sonnet-5-5, Perplexity sonar-pro, Grok grok-4.3, DeepSeek deepseek-flash. We asked 10 questions directly about Poland:

- best-known Poles,
- brands and products,
- what Poland is famous for,
- cuisine,
- science,
- culture,
- companies,
- industries,

- places,
- first associations.

The questions were asked in the languages of the 12 markets, twice of each of the six models. This produced 1,440 answers, assessed by the same judge v2.1. Names of people, brands, companies and places were normalised by the same procedure as in the study. The country to which a name is assigned comes from the normalisation, not from the judge's field, because the judge did not assign a country to some Polish names.

The measure is the number of answers in which a given name appears: at most 144 per question, 12 per question within one market and 24 per question within one model. List similarity between the two repetitions averaged 0.44. The top names recur in almost every answer, while the lower places on the lists vary. We therefore publish the top lists with the number of answers, never individual answers. Open-portrait results are not combined with the study numbers.

Business questions that do not name Poland (B6, 30 September 2026; Part II). We asked 96 B2B questions (8 per market) that name neither Poland nor a Polish city, putting each to the six models in their study versions. This produced 576 answers, assessed by judge v2.1. The measurement complements generative visibility in the B2B segment, for which the study set has too few questions per market. B6 results are published only in Part II and are never shown on the same chart as the generative visibility in Part I.

How the activation gap was calculated. The activation gap compares two situations:

- "Poland on request": the share of open-portrait answers in which a given name appears when the model is asked about it directly (e.g. about best-known Poles; 144 answers per question);
- "Poland in use": the share of study answers in which the same name appears in response to a practical question in which it could naturally come up.

The "in use" subset comprises 209 B2C questions in which names of people, works or brands would be a natural part of the answer. We excluded from it the 12 "what is Poland famous for today" questions, because they are portrait-type questions. For people and works the base is 1,116 answers (186 questions × 6 models). For brands it is only 66 answers, so brand results are treated as a tendency.

On both sides we search the full text of the answers with the same multilingual list of names, including their inflected forms. We excluded homonyms, i.e. the names of airports, universities, institutions and programmes named after the person (e.g. Chopin Airport, Nicolaus Copernicus University, the Łukasiewicz Research Network, Marie Skłodowska-Curie fellowships). We checked this method's agreement with the name map on the open portrait. For the top names from the question about best-known Poles both methods give the same number of answers, e.g. John Paul II, Chopin and Copernicus 144 out of 144, Lewandowski 138 out of 138. Chapter 6 uses both measures: the activation gap table is calculated by text search on both sides, and the lists of top portrait names from the entity map after normalisation; for the same name the two may give slightly different numbers (e.g. Solidarity 131 compared with 121 answers), because text search also picks up forms that normalisation did not assign to the entity. The twelve "what is Poland famous for today" questions (excluded from the subset) serve as a control within the study: the same set, the same judge, a different window (September). The "in use" subset contains no questions about sport, film, literature or games; the absence of athletes from practical answers reflects a lack of occasion, not a gap, and is described as such.

The open portrait and the study are separated by date and, for four models, also by version. Grok and Perplexity ran on the same versions in both measurements. For the other four models, the effect of

the version on the presence of names was not measured separately (the version test B8 concerns the mention of Poland, not names). The gap for names from the canon nevertheless ranges from several dozen to almost a hundred percentage points, which is why we describe it as a difference of order, not an exact value.

How the sectors of the economy were calculated (Part I chapter 6 and Part II). We use seven sectors from the list published by the Polish Investment and Trade Agency (PAIH; paih.gov.pl/sectors): automotive and electromobility, ICT, food processing, machinery and household appliances, aviation, electronics, pharmaceuticals and medical devices. A sector counts as present in an open-portrait answer when the answer names the sector in its own language or names a company assigned to that sector in the dictionary. Companies and industries outside this list are assigned to five groups: energy, fuels, mining and chemicals; finance; retail, fashion, e-commerce and logistics; furniture, construction, rail, telecommunications and transport; defence. The main measure is the question about industries (144 answers), and the question about companies is auxiliary; Part I (chapter 6) uses only the question about industries, while Part II also uses the question about companies.

How the company list was built (Part II). The company names in Part II come from a list of "natural companies": 25 core questions and 16 questions on the extended list, i.e. business questions of the study to which a good answer naturally includes company names. The author drew up the list, which was then checked by two independent blind runs of two models on the question list alone; discrepancies were resolved in the same way as in the classification of question types. Companies are counted from the entity map after normalisation (a company from Poland, once per answer; Part II, chapter 5). Company results are given only in Part II, never alongside the results of the B6 measurement.

Sources and citations. Citation lists were analysed at the level of domains. The same domain cited several times in one answer was counted once. This gives 34,073 citations from 11,081 domains. We did not check whether the full URLs exist. A URL written by a model may not lead to any page, which is why the analysis works at domain level. The two citation modes produce different results: among the URLs given in the answer by the five models without search, 49.5% lead to Polish domains (33.5% to Polish official sites; 14,681 citations); among Perplexity's sources from search the figure is 22.4% (6.5% official; 19,392 citations, or 57% of the total). In the company segment the figures are 54.5% and 25.0% respectively, and in the individuals segment 46.3% and 19.9%. The combined value of 34.0% describes only the scale of the material and is not an inferential figure.

The class of each domain was assigned by a language model (gemini-2.5-flash) in a two-question procedure: the owner's country and the type of organisation. A fixed rule derives the class from these two answers.

- More frequently cited domains (2,909 domains, 24,554 citations; figures at the time of classification, after deduplication 24,814 and 9,259, chapter 5) were classified in the first run.
- The long tail (8,172 domains, 9,519 citations) was classified in a second run, with a further run for domains rejected the first time.
- The "Other / unknown" class covers 124 citations (0.4%).
- 582 citations (1.7%) have a low-confidence class and define the uncertainty band of the metric; the band covers both layers of the classification (frequent domains and the tail) and is given alongside the share of Polish sources.

Invalid URLs (cut off or incomplete) were excluded: ChatGPT 22, Claude 4, Grok 32, DeepSeek 41; the other models had none. The number of citations varies between models by a factor of several dozen (Claude 259, Perplexity 19,392). Per-model shares are therefore always given with the number of

citations, and the overall value is always accompanied by an average weighted equally across models. Claude cites sources in 6.6% of its answers, and with the 1,024-token cap its citations, which usually come at the end of an answer, may have been cut off; Claude's citation figures (259; share of Polish domains 60.2%) are therefore read as a lower bound, with the caveat that it received no instruction about sources.

Home-state sources (chapters 8 and 9). We define a home-state source as an official domain of the persona's home country: for each market there is an explicit list of domain patterns (e.g. gov.uk; .gob.es; auswaertiges-amt.de and .bund.de; regjeringen.no, skatteetaten.no, nav.no, udi.no; .gov excluding .gov.pl for the United States; .gov.ua; .gov.cz and mzv.cz; .gov.sk and mzv.sk; regeringen.se, skatteverket.se, migrationsverket.se; .gov.it, esteri.it, viaggiaresicuri.it; .gouv.fr and service-public.fr; .gov.cn), with Polish domains excluded. We count answers to individuals' questions in which at least one such domain is cited; the percentage (from 6.0% in Slovakia to 20.6% in the United Kingdom: 74 of 360 answers) depends on the list, which is why the text gives the extreme values and the domain names with the number of answers, not a ranking of markets.

How the market-specific thread was identified (chapter 8). A market-specific thread is an association with Poland that comes up in the open portrait in a given market markedly more often than in the others (at least 4 of 12 answers and at least three times the average of the other markets). Its presence in the study answers is counted by full-text search with a list of the name's forms, followed by a reading of the context of each hit, because the name alone is not enough: "Chicago" in American answers is the departure airport and the consulate, not the Polish neighbourhood; "Żywiec" in Slovak answers is a town to go shopping in, not the beer. Hits that, on reading, do not relate to the association are excluded, and the portrait says so.

External sources. The list of faces of the "Poland. More than you expected" campaign comes from an August 2026 press release issued by the Polish Tourism Organisation (POT); the list of strategic sectors from the website of PAIH (paih.gov.pl/sectors); model versions and dates from provider documentation and from the version recorded in each answer.

Change in the ChatGPT app after the study. Since 7 October 2026 the ChatGPT app (GPT-6 with the Intelligent UI feature) can present answers with interface elements, such as comparisons, diagrams and tools, that the API does not return. Even before that, the app displayed some elements absent from the API answer (interactive charts, images from the web). The study measured the text of the model's answers from before this change: the last ChatGPT answers in the study date from 28 September 2026, and those in the B2 and B6 measurements from 30 September. The version test and the open portrait (14–15 October) and the thematic measurement B14 (18–19 October) also measured the text of API answers. The effect of the new interface on Poland's visibility was not measured. Source: OpenAI, ChatGPT release notes, entry of 7.10.2026 (help.openai.com/en/articles/6825453-chatgpt-release-notes); openai.com/index/gpt-6-for-everyone.

Scale of the material. The study metrics are calculated solely on the 8,400 answers of the main study. The research programme as a whole covered 14,722 answers:

- main study: 8,400,
- business questions that do not name Poland (B6): 576,
- repeatability test (B2): 522,
- version test (B8): 1,576,
- open portrait (B9): 1,440,
- thematic measurement "Poland and the war in Ukraine" (B14): 2,208.

Search mode (B17) stands separately, outside this total: 1,410 answers from three models, given on 23 and 24 October 2026. It is a supplementary measurement, carried out after the content of the report had been closed, and described below and in Part II, section 6.5.

The sets are not combined into common metrics, because they differ in date and in questions, and B8, B9 and B14 also in model versions. The second judge (B1) re-assessed 226 study answers, and a person verified 19 answers that the judge had flagged as contradicting the facts. Neither check required new model answers.

Truncated answers. Some answers end mid-sentence because they reached the length cap set by the research application (see “Scope and measurement channel”). The API call log shows what share of each model’s calls reached the cap in the main study: about 85% for ChatGPT and DeepSeek, 77% for Claude, about half for Perplexity and none for Gemini and Grok (a cap of 4,096 tokens, or no cap); in the B6 measurement, 81% for ChatGPT, 76% for Claude, 51% for DeepSeek and 20% for Perplexity; in the thematic measurement B14 the shares were lower (ChatGPT 18%, DeepSeek and Claude 47%, Perplexity below 1%). Truncation was not recorded for individual answers; it is recognised by an answer ending mid-sentence or having a length typical of that model’s truncated answers in that market. Truncation is a property of the research application’s cap, not of the model’s behaviour.

Truncation affects the metrics in different ways. Mention of Poland and generative visibility are robust: in questions that do not name Poland, its first mention usually comes within the first few hundred characters of the answer (median 234–706 characters, depending on the model), and in the full answers of Gemini and Grok at most 2.8% of mentions fall beyond the point at which the cap cuts off the answers of ChatGPT, DeepSeek and Perplexity (about 4,300–5,200 characters); for Claude’s shorter cap (about 2,500 characters) the upper bound is 0.6–13.3%. The lists of countries, the question type, the activation gap and the semantic map do not depend on how an answer ends. Tone and recommendation rate may: a qualifying or recommending sentence may come after the cut point. Comparing truncated and full answers within the study does not settle the question, because for ChatGPT and DeepSeek the full answers are a handful of short answers to simpler questions. It is settled by the truncation simulation on full answers described below. Comparisons of recommendation rate and tone between models are therefore read with the caveat that the models had different length caps (1,024, 1,500 and 4,096 tokens, and no cap), and the per-model recommendation rate is given only in the annex (see “Recommendation rate by model”). Two answers from Gemini’s fallback model (26 and 29 August) end in a long string of spaces after brief content (115,556 and 95,510 characters respectively); the judge assessed their content (mention: country; tone 6 and 7), and they remain in the study numbers, where, among 8,400 answers, they do not change any metric.

Truncation simulation on full answers (B16, 22 October 2026). To measure how much the length cap can change the assessment, we took full (untruncated) answers from three models and cut them as the cap had cut answers in the study: to 1,500 tokens (the cap for ChatGPT, DeepSeek and Perplexity) and to 1,024 tokens (the cap for Claude), using each provider’s tokeniser and cutting at a token boundary. The full answers came from Gemini in the study (466 answers: 216 to questions that do not name Poland and 250 to questions that do; we excluded 26 questions without Poland that were served only by the fallback model), Grok in the study (492: 242 and 250) and Claude in the claude-sonnet-5-5 version from the open portrait and the thematic measurement (602 answers with a cap of 4,096; we excluded 5 answers longer than 3,500 tokens and 1 broken-off answer). Each answer was assessed by the same judge v2.1 in four forms: the assessment recorded in the study, the same full answer re-assessed, a copy cut to 1,500 tokens and a copy cut to 1,024 tokens (3,156 assessments). The truncation effect is the difference between the cut copy and the re-assessed full answer, calculated on answers that were actually cut; the difference between the re-assessment and

the recorded assessment of the full answer measures judge drift and noise. We fixed the rule for reading the result before the run.

Result. Mention of Poland and tone are robust to truncation: the difference in tone does not exceed 0.13 points, the loss of mentions when cutting to 1,500 tokens does not exceed 0.6 percentage points, and when cutting to 1,024 it reaches 3.7 percentage points in questions that do not name Poland (Gemini), which, transferred to Claude in the study, gives at most 2.9 percentage points. The recommendation rate is not robust: cutting to 1,024 tokens lowers the share of answers with a high recommendation level by more than 10 percentage points (Claude 5-5 on 395 cut answers: −10.9 percentage points, 95% interval from −15.2 to −6.6), and cutting to 1,500 tokens by 5–10 percentage points (Claude 5-5 −9.3; Gemini −2.6; Grok cannot be assessed at 1,500 tokens, as only 2 of its answers were longer). Judge drift and noise do not exceed 2.4 percentage points of recommendation rate and need no separate caveat. Transferred to the study (the effect multiplied by the share of answers at the cap), this gives an estimate of the impact of truncation on the recommendation rate: Claude up to −8.4 percentage points (up to −6.3 in questions that do not name Poland), ChatGPT from −2.2 to −7.9, DeepSeek from −2.2 to −8.1, Perplexity from −1.2 to −4.3. This is a transfer, not a measurement: the simulation cuts the full answers of Gemini, Grok and Claude 5-5, not the answers of the models actually truncated in the study. The version test provides a check: the rise in Claude's recommendation rate between versions is similar in pairs with a full answer (+15.5 percentage points) and with a truncated one (+14.3), although with a truncation effect of the order of −11 percentage points the truncated group should have risen by about 26. The simulation therefore indicates an upper bound on the impact of truncation rather than its exact value. Consequences for the report: we moved the per-model recommendation rate from the chapter text to the annex, with this caveat; in the chapter text, mention rate, tone and divergence remain per model. Even under the least favourable variant of the simulation, Gemini remains, in its August 2026 version, the model with the highest recommendation rate and Claude the model with the lowest: the difference in recommendation rate in questions that do not name Poland is 26.8 percentage points, and at least 20 percentage points after subtracting the upper bound on the impact of truncation; the order of the models in between is not established. In the B6 measurement the transfer gives Claude up to +8.3 and ChatGPT up to +7.6 percentage points, so the order of the models by recommendation rate in B6 cannot be established.

Recommendation rate by model (with a caveat). All the values below relate to the model versions of the study window (August–September 2026) and are affected by the different length caps (see above); we read them comparatively, not as an absolute percentage of answers recommending Poland (see "Second judge"). Share of answers with a high recommendation level among answers that mention Poland:

model	all questions	questions not naming Poland (n)	regional / general questions (not naming Poland, B2C)	companies / individuals (not naming Poland)	second-judge check: main judge / second judge (B1; whole model sample, n 37–38)	B6 (companies, not naming Poland)
ChatGPT	67.1	52.8 (161)	48.8 / 68.4	58.1 / 50.8	59.5 / 35.1	86.3
Gemini	85.2	70.0 (180)	59.0 / 72.0	82.5 / 64.2	78.9 / 68.4	87.5
Claude	55.2	43.2 (183)	31.0 / 64.0	50.9 / 39.8	42.1 / 21.1	75.3
Perplexity	65.8	52.2 (161)	42.0 / 50.0	69.2 / 46.7	44.7 / 34.2	76.9
Grok	71.9	57.2 (180)	52.3 / 65.2	56.1 / 57.7	59.5 / 35.1	85.4
DeepSeek	72.3	59.7 (186)	49.4 / 57.1	73.2 / 53.8	57.9 / 39.5	83.3

In general questions the recommendation rate is higher than in regional ones for each of the six models (62.3 compared with 47.0% overall). The second judge sets a higher bar for every model but

preserves the extremes: Gemini first, Claude last. In the B6 measurement (96 company questions that do not name Poland, 576 answers) the share of answers with a high recommendation level by market (48 answers each) is: China 72.9; Czechia 87.2; France 83.3; Germany 83.3; Italy 80.4; Norway 95.7; Slovakia 76.6; Spain 68.8; Sweden 95.7; Ukraine 91.7; the United Kingdom 80.4; the United States 75.0. With 48 answers per market, differences of less than a dozen or so points cannot be interpreted, and we do not rank the markets.

Position of Poland in the answer (prominence). For each of the 8,400 answers we established where the first mention of Poland occurs and how many other countries the model named before it. Position is calculated with the same text rule that the report uses to recognise Poland in questions (country name, adjective, language name, Polish cities, in the inflected forms of 11 languages); whether Poland is present in the answer, however, is taken from the judge's assessment (chapter 2), and the text rule finds it in 7,975 of those 7,992 answers (99.8%). Poland's rank is 1 plus the number of countries listed by the judge whose first mention in the text precedes the first mention of Poland; countries are located in the text by name, adjective or main city, and regional groupings ("Benelux") are ignored. A rank is given only when every country on the judge's list could be located in the text (93.0% of answers with Poland); separately, we calculate the rank excluding the persona's home country, which often opens the answer as a point of reference. Truncation of an answer understates only the presence of Poland (an answer cut off before the first mention), not position or rank. The text rule does not know some names in the market languages (e.g. "Katovice", "Masuren"), so in 3.5% of answers with Poland the judge places the first mention earlier than the rule does. We keep the report's rule for consistency with the other numbers.

set of answers	n with Poland	median first mention (characters)	Poland within the first 1,000 characters	rank 1	rank 1 or 2	rank 1 excluding the home country	countries per answer (mean)
questions that name Poland	6,932	32	98.6%	91.8%	99.0%	96.5%	1.9
questions that do not name Poland	1,046	431	78.2%	36.0%	57.7%	41.1%	5.9
— companies	302	480	—	54.7%	—	58.0%	—
— individuals	744	397	—	28.3%	—	34.2%	—

The n values come from the text rule: against the judge's assessments (6,941 and 1,051) it misses Poland in 17 answers (5 of them answers to companies) and finds it in 3 without the judge. Position in a single answer varies, which is why the report does not describe it in the main text. In the repeatability test (30 September; 29 questions, six models, three repetitions) and in the control measurement of 2 November 2026 (50 questions, 31 of them not naming Poland; five models in their November versions; three repetitions on the same day; rank counted with a text rule without the judge, which matches the report rule in 98.9% of the study's answers to questions that do not name Poland), two answers to the same question differed on whether Poland came first in roughly one case in five, and the share of answers with Poland in first place for questions that do not name Poland shifted between runs on the same day by about 10 points; for questions that name Poland, by less than 3 points. The figures in the table therefore describe only a dated measurement and are not a property of a model or of a single question; the position in a single answer is not a result in this report. We treat differences between models as description, not as a ranking, because they depend on the length and style of the answers.

Thematic measurement "Poland and the war in Ukraine" (B14, 18–19 October 2026). Through their official APIs, we put to the six models 16 questions about the safety of travelling, studying,

relocating and investing in Poland in the context of the war in Ukraine: 8 questions from individuals in 11 markets (excluding Ukraine) and 8 questions from companies in 12 markets, in the language of each market, each asked twice at least 24 hours apart. This produced 2,208 answers. Four questions do not name Poland. One individuals' question in nine markets repeats a question from the main study word for word. The models ran in their current production versions (ChatGPT gpt-6-luna, Gemini gemini-3.5-flash-lite, Claude claude-sonnet-5-5, Perplexity sonar-pro, Grok grok-4.3, DeepSeek deepseek-flash), as in the open portrait; Perplexity and Grok ran in the same versions as in the study. The answers were assessed by the same judge (v2.1) as in the study. B14 numbers never form a common metric with the study, B6 or B9, and we never set them against the order of models from Part I, because the model versions differ.

Stance on risk (A–D). We assigned each answer to one of the following classes: (A) Poland is safe, war risk is irrelevant to the decision; (B) safe with specific caveats (e.g. the border zone, a contingency plan, insurance), recommendation "yes"; (C) significant risk, the answer urges reflection or presents Poland as riskier than the countries compared, without advising against it; (D) the answer advises against Poland or recommends another country because of war risk; an answer that recommends Poland and names another country only as an option for those who put risk first is class B or C, not D; (N) Poland appears in the answer, but without an assessment of war risk; (0) Poland does not appear in the answer. The classification, on the full text of the answers and without knowledge of the model or of the judge's assessments, was carried out by Claude Opus 5.5 for all 2,208 answers; independently, with the same instruction, Claude Sonnet 5.5 classified a stratified random sample of 576 answers (26%; 6 for each segment × question × model combination). Agreement between the two runs: 87.5% (κ 0.80). In four questions on the boundary between classes B and C (safe European countries for a few years, investment, nearshoring, war-risk insurance), within-question agreement fell below κ 0.6, and some batches of the first run applied different conventions at that boundary; we recoded these 564 answers in a calibration run (one question = one set of coders, two balanced halves), and this classification is final (agreement with the second coder 85.1%, κ 0.75). The tables give the final classification and, for the four calibration questions, also the range between runs, because the B/C boundary there depends on the instruction, not on the data. In the question about safe European countries for a few years, between half and three quarters of the answers place Poland in the "with caution" group, and in the question about investment almost every answer says "risk higher than in Czechia, but one to be priced and insured". For three of these questions (investment, nearshoring, war-risk insurance; 432 answers) we separately coded whether the answer presents Poland as clearly more exposed to war risk than the countries it is compared with (yes / no / no comparison). We used six independent runs of Claude Sonnet 5.5 for this relative measure, one for each half of each question. Result: 67% of answers present Poland as clearly more exposed (investment 76%, insurance 73%, nearshoring 53%; no comparison 2%), and yet 91% of those same answers fall into class B (Poland recommended with caveats). Class A answers never present Poland as more exposed; class C answers almost always do (22 of 25). By model, the share of "yes" ranges from 53% (Gemini) to 90% (Perplexity), although it depends more on the question than on the model (Gemini: 0% in the nearshoring question, 96% in the insurance question). Halves of the same question coded by different runs differ by up to 18 percentage points (investment: 85 and 67%), so we read the figures for this measure as ranges, not point values. The hard numbers, consistent across all runs, are those for classes D and A. For the questions on energy, workers from Ukraine and a reconstruction base, in which class N predominates, we give only the share of warning answers (C+D) and the judge's tone and recommendation rate.

Arguments. The presence of 13 arguments (including NATO, Belarus, warnings from the home country's foreign ministry, the border zone, drones and incidents, shelters, refugees and energy) is

counted by full-text search with a list of forms in 11 languages; the percentage is the share of answers in which the argument appears at least once.

Factual errors in B14. The same procedure as in the study (judge → human). Of the 10 answers flagged by the judge, we confirmed 1 (a minor one: “2002” instead of 2022); in 7 the judge was wrong, as it does not know about events in 2025; 2 were not factual claims.

Deviations from the B14 plan. Repetition numbers were assigned by the date of the answer (18 and 19 October), not by the order of the runs; each pair has exactly one answer per day, and the shortest interval within a pair is 24.7 hours. The second repetition had a 30-minute pause with no effect on the data, and its window stretched to 11:08–16:33 UTC because of a break in one provider’s API availability and the retries that followed. 49 attempts that ended in a provider error (45 Claude, 4 Grok) were excluded from the files; each was retried once and each has a valid answer.

Search mode (B17, 23–24 October 2026). The measurement tests whether access to a search engine changes the same model’s answer to the same question. The set: 233 study questions that do not name Poland (166 from individuals, 67 from companies; the questions retained in the control variant, see “Origin of the questions”) and 6 questions on which the models that have a search tool made confirmed factual errors (Claude 4, Gemini 2). Models: ChatGPT gpt-6-luna, Gemini gemini-3.5-flash-lite, Claude claude-sonnet-5-5. These are production versions from October, the same in both modes (version recorded in each answer). Each question was asked twice: without the tool, along the research application’s usual path, and along the path with the provider’s search tool (ChatGPT and Gemini with the provider’s search engine, Claude with at most three searches per answer); the median interval between the two answers was 11 minutes, the maximum 3 hours 45 minutes. This produced 705 pairs (1,410 answers, 0 incomplete pairs): 699 core pairs and 6 error pairs. For one company question from Spain the ChatGPT and Claude pairs were produced on 24 October (the run of 23 October had yielded only the Gemini pair for it): the same model versions, the same paths, the two answers 11 seconds apart; all the other answers date from 23 October. With each answer in search mode we recorded the list of queries, the list of sources returned by the tool and a flag from the provider’s metadata showing whether the tool was invoked: ChatGPT 228 of 234, Claude 236 of 237, Gemini 104 of 234 (130 Gemini answers were produced with the tool available but not invoked, and we count them separately). All 1,410 answers were assessed by judge v2.1 (gemini-2.5-flash, the same code as in the study, with no fallback model) on 24 October: 1,406 in a single run and the four added answers the same day, with the same code. Length cap: 4,096 tokens without the tool, 8,192 with it; one Claude answer with search reached the cap. A limitation that has to be read together with the result: the application’s two paths differ not only in the tool. Without the tool, ChatGPT receives a system instruction asking for sources and URLs, and with the tool none, but a higher reasoning effort instead; Claude runs without an instruction when it has no tool, and with the tool receives an instruction to answer concisely and to cite the sources it finds; for Gemini the only differences are the tool and the cap. The difference between the modes therefore measures the answer without the tool against the answer from the search path, not the effect of search alone, which is why we do not compare the number of citations between models. For Gemini, the sources from search are redirect URLs, so we take the domain from the title of the snippet returned by the tool; domain classes are as in the study, each domain counted once per answer. The results below (Part II, section 6.5) are never combined with the study figures. Mention of Poland without the tool / with search: ChatGPT 57.9 / 55.8% (233 pairs), Gemini 73.4 / 70.8% (233; in the 103 pairs in which search was invoked, 73.8 / 71.8%), Claude 84.1 / 74.7% (233); tone among mentions 7.01 / 6.79, 7.58 / 7.61, 6.93 / 6.71; recommendation rate (high) among mentions, with the same caveat as above: ChatGPT 53.3 / 51.5%, Gemini 66.7 / 66.7%, Claude 48.0 / 46.6%. All the differences in recommendation rate and tone are below the thresholds, and the difference in mention rate exceeds the threshold only for Claude, whose search path has a different instruction. In the six error pairs, the

error from the study did not recur in either mode; the pairs were produced on newer model versions than the errors, so we do not attribute this to search. The other eight errors (DeepSeek 6, Grok 1, Perplexity 1) cannot be measured in search mode: DeepSeek and Grok have no search tool in the research application, and Perplexity always searches, so no pair without the tool can be built for it. B17 answers do not feed into any study metric; the files are listed under “Integrity check”.

Brand semantics as a research concept. The report treats Poland the way marketing research treats a brand: as a set of meanings, associations and evaluations that the brand holds in a given mind and in a given culture. This tradition is built on work on brand knowledge and its associations (Keller 1993), on brand concept maps (John, Loken, Kim and Monga 2006) and on the semiotics of meaning in consumption (Mick 1986). What these share is the assumption that a brand’s meanings can be measured through what a given mind says when it is questioned indirectly. The Semantic Brand Score (Fronzetti Colladon 2018) carried this measurement over to large bodies of text and broke it down into three dimensions: the frequency of the name (prevalence), the diversity of its associations (diversity) and its role as a connector between other concepts (connectivity). They inspire three of the report’s measures (without operationalising them): generative visibility, the activation gap together with the entity map, and the “doors” into the answer.

The report’s thesis is that the minds whose brand semantics are worth measuring today include language models. Work from 2024–2025 has already described the image of countries in ChatGPT’s answers (Manor and Segev 2025), the effect of a persona’s nationality on how a model assesses countries (Kamruzzaman and Kim 2025), and recommendations of places for personas from different countries (Andreev et al. 2025). In the sources searched we found no work that combines what this report combines: personas from 12 markets asking in 11 languages, questions that do not name the country being assessed, an assessment of whether the model recommends that country, and six models queried through their official APIs within a dated window.

Concepts used in the report.

Activation gap – the difference between what the models name when asked about the brand directly and what they name in practical answers in which the brand appears: the canon is present in 144 of 144 answers “on request”, and in 0.8–12% of practical answers (chapter 6). At the level of a single market: a market-specific thread that does not make its way into practical answers (chapter 8).

Brand semantics – the set of meanings, associations and evaluations that a brand (a product, a place, a country) holds in a given mind and in a given culture; in the models’ answers it shows up as a set of entities (places, people, events, characteristics) linked with the brand name, together with the links between those entities and the situations the personas ask about. In this report the brand is Poland; we measure brand semantics indirectly, through the answers of six models to 1,400 questions from personas in 12 markets within a dated window, including when the question does not name Poland (chapter 1).

Citability – a property of a source that models give in an answer: it gets there not because it is about the brand but because it addresses the persona’s situation in the persona’s language. In each of the 12 markets the models cite the public authorities of the persona’s home country; Polish official websites reach the answers through their own English-language services for foreigners (chapters 5, 9).

Companies on demand (Part II) – Polish companies named in answers to 25 questions to which a good answer must include the names of Polish companies. Companies in use (Part II): Polish companies named in answers to practical company questions; denominator 4,200 answers. Companies on request (Part II): companies that a model names when the question explicitly asks for Polish companies (open portrait).

Entity – a named thing that can be recognised in an answer regardless of language and inflection: a person, place, institution, event or product. Poland's canon (Chopin, Wałęsa, Copernicus, John Paul II, Skłodowska-Curie, Lewandowski) consists of the entities named most often when the question asks about Poland directly (chapter 6).

Entity linking – the mechanism by which an entity makes its way into a practical answer: the model brings it in only if it has it linked with the entity the persona is asking about, such as a place, a route or a situation. In the text of the report this is called the "doors" (chapter 6): Chopin comes in with Warsaw and Żelazowa Wola, the Tatras with a trip to the mountains; Skłodowska-Curie is linked with Poland, but with nothing that people planning a trip ask about.

Home country – the country of the persona on whose behalf the question was asked, e.g. Czechia for a question asked on behalf of a Czech company. Written in cards and tables as "Home country: Name"; it means the persona's country, not a country the model recommends.

Persona – a scenario of a person or company asking a question, created for the study: country, language, situation and purpose of the question. It is not a real person or company. Examples are given above; the full list is in the open data.

Prominence – the position of Poland in the answer: where the first mention falls and the rank, i.e. the number of countries named before it (see "Position of Poland in the answer").

Risk question (Part II) – a question that itself asks for risks, barriers, pitfalls or typical mistakes to be identified, or that expresses a concern of the company asking; 281 of the 700 company questions, identified in a blind round by two models (agreement 98.6%).

Semantic drift – a change in brand semantics in the answers between model versions or over time, with the same questions. We distinguish it from measurement noise: the differences between repetitions of the same question in the same model version (repeatability test), which set the report's thresholds (tone 0.5 points, mention rate 3 percentage points, recommendation rate 5 percentage points). The newer-model-version test showed drift that was larger than noise and had a direction (chapter 7); measuring brand semantics therefore makes sense only as a dated measurement, and a claim of change requires control measurements on the same set of questions, because models are updated.

Semantic map – a breakdown of entities by market, question type and source: which entities the models link with the brand in a given market, in which questions they name them and from which domains they take them. The twelve portraits in chapter 8 together make up the semantic map of the brand Poland.

Sources from search / URLs given in the answer – the two modes from which citations come: the list of search results returned by the API (Perplexity), or URLs written by the model in the text of the answer (the other models). The report gives them separately (chapters 5, 9).

Truncation simulation on full answers – a test in which the models' full, untruncated answers were cut to 1,500 and 1,024 tokens, as the length cap cut answers in the study, and then re-assessed by the same judge. The difference between the assessments estimates the cap's impact on recommendation rate and tone (see "Truncated answers").

Generative visibility (chapter 2) and market-specific thread (chapter 8) are defined where they are first used.

References cited in the annex.

- Andreev, H., Kosmas, P., Livieratos, A. D., Theocharous, A., Zopiatis, A. (2025). Destination (Un)Known: Auditing Bias and Fairness in LLM-Based Travel Recommendations. *AI*, 6(9), 236.

https://doi.org/10.3390/ai6090236

- Fronzetti Colladon, A. (2018). The Semantic Brand Score. *Journal of Business Research*, 88, 150–160. https://doi.org/10.1016/j.jbusres.2018.03.026
- John, D. R., Loken, B., Kim, K., Monga, A. B. (2006). Brand Concept Maps: A Methodology for Identifying Brand Association Networks. *Journal of Marketing Research*, 43(4), 549–563. https://doi.org/10.1509/jmkr.43.4.549
- Kamruzzaman, M., Kim, G. L. (2025). Exploring Changes in Nation Perception with Nationality-Assigned Personas in LLMs. *Proceedings of EMNLP 2025*. https://doi.org/10.18653/v1/2025.emnlp-main.181
- Keller, K. L. (1993). Conceptualizing, Measuring, and Managing Customer-Based Brand Equity. *Journal of Marketing*, 57(1), 1–22. https://doi.org/10.1177/002224299305700101
- Manor, I., Segev, E. (2025). What ChatGPT "Thinks" About Your Country? Sentiments and Frames of AI Geographies. *Policy & Internet*, 17(3), e70013. https://doi.org/10.1002/poi3.70013
- Mick, D. G. (1986). Consumer Research and Semiotics: Exploring the Morphology of Signs, Symbols, and Significance. *Journal of Consumer Research*, 13(2), 196–213. https://doi.org/10.1086/209060

Integrity check. The checksum of a set is the md5 of the sorted list of the md5 sums of its files.

- Study set: 24 files, set md5 270b1eb5c140b2b0063450b37f01fe8a (checked on 14, 15, 16, 17, 20, 22 and 23 October 2026).
- Business questions that do not name Poland (B6): 14 files (12 markets, question mapping, summary), 23ab8662ec3f6925b9bc5bd5b46164c0.
- Repeatability test (B2): results a9196b1ade35afb94c5253e98aa6ec31, report ff77f59b3794aecfabcdd744446795c9.
- Version test (B8): results e8652bec0fcaecbaf7c389d282c258cf, report 46da3f8c250e45b8cf3f08a1f70e59d5.
- Open portrait (B9): 12 files, set md5 f3e6dd05c8221d0dd5d33535a5f8e68b.
- Thematic measurement (B14): 23 files (12 company markets, 11 individuals markets), set md5 64043388d87c03bb64fcc117e8dac97b; measurement report b3527b07…; A–D classification (coder 1, coder 2, calibration, final class) and coder instructions, kept as files in the working materials (b14_out/).
- Scripts for recognising truncated answers: bce5821f… (rule), 7cf3929b… (position of first mention).
- Tail domain classes: f90dc823b487b8fb9a70c235803a7768.
- List of questions with a Polish name in an inflected form that the automatic rule did not recognise (12 questions): a76bb1660d335beb74d886295e283249.
- Factual error sheet: original 239759dd34557be04edcda044799ed7a, verified version 61b227336407783f4d2ea72204066af2.
- Question-type classification (174 B2C questions): list of questions for blind classification dad0fbd26080994e5f9ede82008442c5; two blind runs 45dcfcada4566c8785b78c81728c0bae and 087d8d531255f8b0ec1f4408a48c1696; comparison 7b5da1f5de6e7b11096329e10f2e3426; final classification v2.5 aa7d33e61bdfb39f64d283c25edccc8e (previous, 172 questions: c71b3f85474f2c5b9b45b5f32cecebc2); blind round on the two French questions: input 33204802…, results 281ad52c… and 5229a187…, comparison 7cc7f1a7….
- Truncation simulation on full answers (B16): assessment results f96073e3…, run report 088ba78f…, analysis script ea8f48f6…, results tables a9bc1a9b….
- Search mode (B17): 24 files (12 markets × individuals / companies), set md5 f525a06a…; combined file 58494e8d…, pairs 7ed3a559…, the providers' raw search-tool responses 7edb0458…, run

reports f0dd3204… and 71ec1b5c…, pair-reading script e41aece3…, supplementary-figures script 9131541c….

- Post-measurement review of the 1,400 questions: rows 777d0984…, repetition groups c3db7abe…, summary 5124210a…, author's blind-sample sheet 08861fc6…, control-variant script 84868de2….
- Position of Poland in the answer (prominence): rows bc9526ab…, aggregates script 519173c1….
- Thematic measurement B14, relative measure (432 answers): results of the six runs 7d00c79d…, 2820db12…, acce620e…, 71889144…, 59a507c2…, 8252efb8…; instruction 3e8b717f…; merge 1c4d596f….
- Open data of the report: scenarios by market (the 1,400 questions of the main study and the questions of the B6, B9 and B14 measurements) 64908f5f…; personas by market 3c90f12b…. Open-data package (zip, 2 files with a README) 0b9cce42….
- Question-type classification (164 B2B questions: 68 from the study and 96 from B6): list for blind classification ddc857bcc0dc307b41c64b1c18a712b9; instruction e0ee8dd05ef8da43fe3d25796b63e2db; two blind runs 509420fcbb322c98b00429b838482cd6 and 7eb7aab75297a414f6cdf54ce2b9c0c3; comparison 9992d9ec2d599b352f2ca89fc8ff01d8; final classification 61f4fe14e97c7e0b6d088b690fe21f4e.

For the six models, the study files contain exactly one row for each question, model and repetition: the answer with content or, where a question was asked more than once, the most recent answer. The research application's database also stores technical attempts that ended in a provider error and were retried, and a small number of earlier answers replaced by a retry, which is why it holds more rows. Every row in the files was confirmed to match the database on 16 October 2026, using row identifiers (Log ID).

Measurements B2, B6, B8 and B9 were carried out in copies of the research projects, so the study set was left untouched (last answer: 28.09.2026). Deviations from the plan: we reclassified tail domains rejected in the first run on 14 October; we carried out the open portrait in two repetitions rather than three; we corrected the "Poland named in the question" flag twice after seeing the data (the list of 12 questions, then the form "posen"), each time recalculating all the numbers and recording the change; the answer length cap for Claude in the main study was lower than for the other models (1,024 compared with 1,500 and 4,096 tokens), which we found after the measurement ended; the questions were reviewed only after the measurement (see "Origin of the questions"), and the outcome of the review did not change any number in the report.