ALPHACITATION RESEARCH · RESEARCH PAPER

The Moving Instrument: Measuring Persona Effects in AI Answers About Medicines

An internally registered study of one medicine (September 2026) and its externally pre-registered replication on three medicines (November 2026), in the United Arab Emirates

Ozue C. Ezeife

November 2026 · Version 1.0

15 min read

Abstract

Question. When a measurement tells an AI assistant who is asking about a medicine, it hopes to observe a persona-associated signal. This paper asks two questions about such a measurement: validity (does it record the stated role, or the way it asks?) and stability (does the way it asks have the same effect over time?).

Design. Four AI assistants, labelled A to D, were asked about medicines in the UAE with no context, with a neutral line of context, with a line stating the asker's role, and with the role written into the question. The study was registered internally before its first run on one medicine on 27 September 2026, then pre-registered publicly and replicated on three medicines on 25 November 2026. Each medicine received 624 answers, and no call failed.

Results. The way the question was framed changed answers, and the change was not stable. On Wegovy, a neutral line of context cut one assistant's naming of the medicine from 12 of 12 answers to 1 of 12 in September; in November it raised it, from 5 of 12 to 11 of 12. On another medicine the same day, it cut naming from 12 of 12 to 0 of 12. One assistant showed the same persona-associated pattern on every medicine, naming the medicine in every answer for a stated prescriber or payer, across all four medicine-run combinations. Of three registered secondary predictions, one held for all three medicines and one for none. A proprietary attribution measure, read against a 50% line fixed before data, found the Wegovy reading mostly attributable to the measurement procedure in both runs (role share 44.0% in September, 38.4% in November), and Baxfendy read the same way (17.0%). Leqembi's question set did not meet a precondition the rule needs, so its result is undetermined, and generalisation, which the registration made conditional on both new medicines, is not claimed.

Conclusion. In this study, how the measurement asked moved answers at least as much as who it said was asking, and that effect changed within two months for one assistant on one medicine. The results support assistant-level reporting, matched controls, structural checks before a study runs, and periodic re-measurement of the instrument. A single correction fitted once would not have carried over in this case.

---

1. Two questions about a persona measurement

A measurement of how AI assistants describe medicines can ask the question as different people would: a prescriber, a patient, a pharmacist, a payer. The difference between those answers is read as the effect of who is asking. That reading assumes that telling an assistant who is asking changes nothing but who is asking.

A measurement can break that assumption in two ways. On validity, it can confound the role with its own procedure: adding a line of text and rewording a question can each move an answer. On stability, the procedure's own effect can change over time.

This paper measures both. AlphaCitation measures how AI answers describe medicines, and it holds its own instruments to the evidential standard it applies to those answers.

---

2. Methodological transparency and proprietary intellectual property

The experimental design, pre-specified decision criteria, observed results and study limitations are disclosed to support scientific scrutiny.

Certain computational components of AlphaCitation's measurement and attribution framework constitute proprietary intellectual property and are not disclosed in full.

Attribution estimates should therefore be interpreted within the specified measurement framework. Independent reproduction of these estimates requires access to the underlying computational methodology. Directly observed answer frequencies and experimental contrasts are reported separately from proprietary derived measures.

In this paper, Sections 4.1 to 4.4 and Exhibits 1 to 3 report observed findings: answer frequencies (k of n answers) and the differences between conditions, in percentage points, each of which can be checked directly from the counts printed. Section 4.5 and Exhibit 4 report the registered attribution shares, which are proprietary derived measures. Access to the methodology may be granted for independent review, subject to the conditions in Section 7.

3. Design and registration

3.1 Two runs, two kinds of registration

	SEPTEMBER 2026	NOVEMBER 2026
Medicines	Wegovy	Wegovy (re-measured on the same input), Leqembi, Baxfendy
Registration	Internal: design and decision rule committed to AlphaCitation's dated records before the first call	External: OSF registration (25 Nov 2026, 09:01 UTC) and an RFC 3161 timestamp from FreeTSA (08:51:34 GMT) over the pre-registration's SHA-256, both before the first call
Decision rule	Role share against a 50% line; Holm correction at 0.05 across the four roles	September's rule, unchanged
Answers	624	1,872 (624 per medicine)

Wegovy is the longitudinal reference medicine. Re-measuring it on the same input separates a change in the assistants and over time from a change of medicine.

Two of the four assistants were called under the same model name as in September, and two under a newer model name. An unchanged model name does not show that the model behind it is unchanged. Which labels are which is not disclosed.

3.2 Medicines, assistants, roles, conditions, outcome

MEDICINE	MANUFACTURER	AREA
Wegovy (semaglutide)	Novo Nordisk	weight management
Leqembi (lecanemab)	Eisai (co-commercialised with Biogen)	early Alzheimer's disease
Baxfendy (baxdrostat)	AstraZeneca	hypertension not adequately controlled on other agents

Four widely used AI assistants were measured, labelled A to D; the label assignment was fixed before data, and its digest is in the registration. The four asker roles span prescribers, patients, dispensers and payers.

Exhibit S1 shows the design. The four conditions compared are:

1. the question with no context;
2. the same question after a neutral line of context;
3. the same question after a line stating the asker's role;
4. the role written into the question itself.

The first three ask an identical question and differ only in the line before it. The fourth uses a formulation employed in AlphaCitation's own measurement framework.

Each medicine received the 624 answers the registration states: 480 in the four compared conditions, that is 120 per assistant (12 with no context, 12 with the neutral line, 48 with a role line and 48 with the role written in, 12 per role in each), and 144 used only by the proprietary derived measure. Every registered answer was collected and scored, and no call failed in any run. Exhibit S2 gives the counts behind Exhibits 1 to 3, with its scope stated on it.

The outcome is whether an answer names the medicine. For Wegovy the registered reading also counts its active substance, semaglutide.

A rule fixed before data attributes the measured effect either to the stated role or to the measurement procedure, and reports the role's share. A share of 50% or more reads **mostly persona**, and a share under 50% reads **mostly instrument**. The role effect is called distinguishable from zero for a role only if it survives a Holm correction across the four roles at $\alpha = 0.05$, and the rule gives no verdict when its preconditions are not met.

4. Results

4.1 Stability: one assistant's response to the same line reversed

Exhibit 1. The same context. Opposite effects.

Assistant A · Wegovy or semaglutide naming
Observed answers: no context against a neutral line of context. Points are % of answers; changes are percentage points (pp).

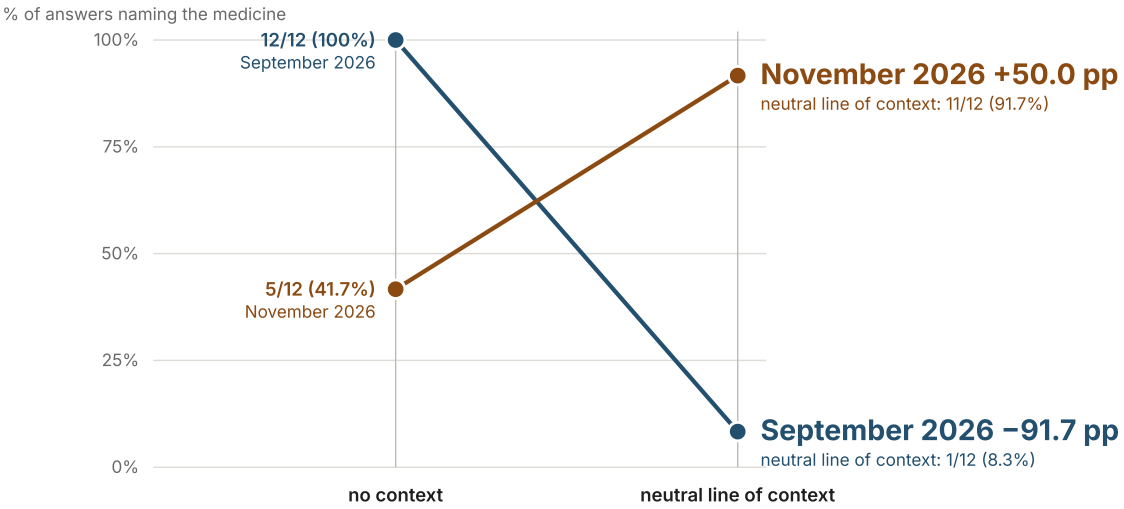

12 answers per condition. Descriptive observations 59 days apart.
Comparison: neutral line of context against no context, same assistant and medicine, on 27 Sep 2026 and 25 Nov 2026. Limitations: one assistant and one medicine; the assistant may have changed between the dates, so the two observed contrasts are two readings of a moving instrument, not one effect read twice.

In September a neutral line cut assistant A's naming of Wegovy from 12 of 12 answers to 1 of 12. In November the same line raised it, from 5 of 12 to 11 of 12, and every November answer with a role stated in a line of context named it (48 of 48). On Baxfendy, the same day, assistant A collapsed further: 12 of 12 with no context, 0 of 12 with the neutral line and 0 of 48 with a role stated in a line of context. On Leqembi it moved by one answer at most.

For this assistant and this medicine, a correction fitted to the September run would have had the wrong sign in November. This is one observed reversal, not a general rate. The record shows what the assistant answered, not why it changed; the cause is not established.

4.2 A repeated persona-associated pattern on all three medicines

Exhibit 2. Assistant C consistently names medicines for stated prescribers and payers

Observed answers, across all four medicine-run combinations. Outcome: answer names the medicine; assistant C only; n = 12 answers per point. Grey ring: no context. Copper dot: the role stated in a line of context. Copper ring: the role written into the question (a ringed dot means both gave the same count). Arrows show the change from no context in percentage points (pp).

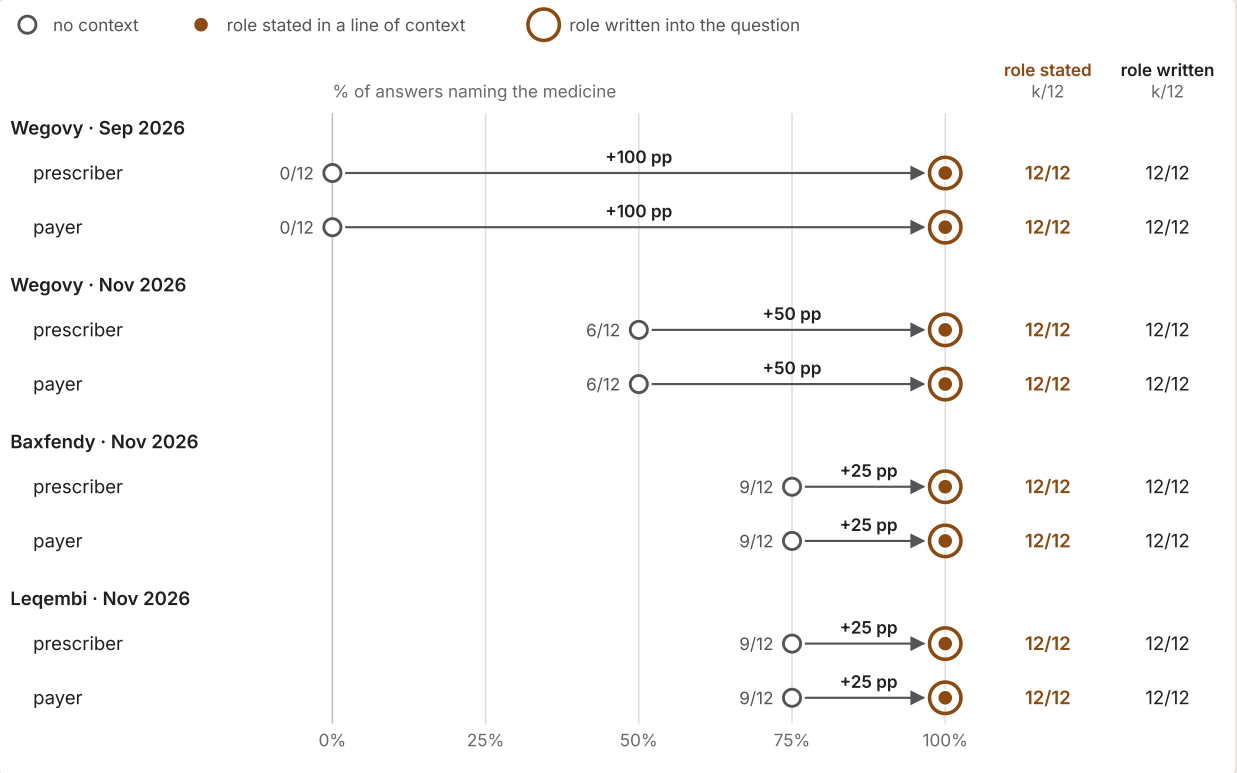

Comparison: registered prediction S2 ("assistant C names the medicine more often for a prescriber or a payer than when the question is asked with no context"), held for all 3 medicines in November; September and the role-written points are shown for reference. Limitations: 12 answers per point; at 12/12 no further rise is possible; one assistant; it may have changed between the two dates.

Assistant C named the medicine in every answer when a prescriber or a payer was stated, whether in a line of context or in the question, across all four medicine-run combinations. It is an observed pattern in one assistant, not an isolated persona effect. On Leqembi, which sits near the upper limit of naming, the move is 3 answers of 12.

4.3 Stating the role against rewording the question

Exhibit 3. Stating the role, then writing it into the question: two contrasts in naming

Observed contrasts. Outcome: answer names the medicine. Change in percentage points (pp), the mean over the four roles; for each role, 48 answers with the role stated in a line of context against 48 with no context, and 48 with the role written into the question against those 48 (12 per assistant in each cell).

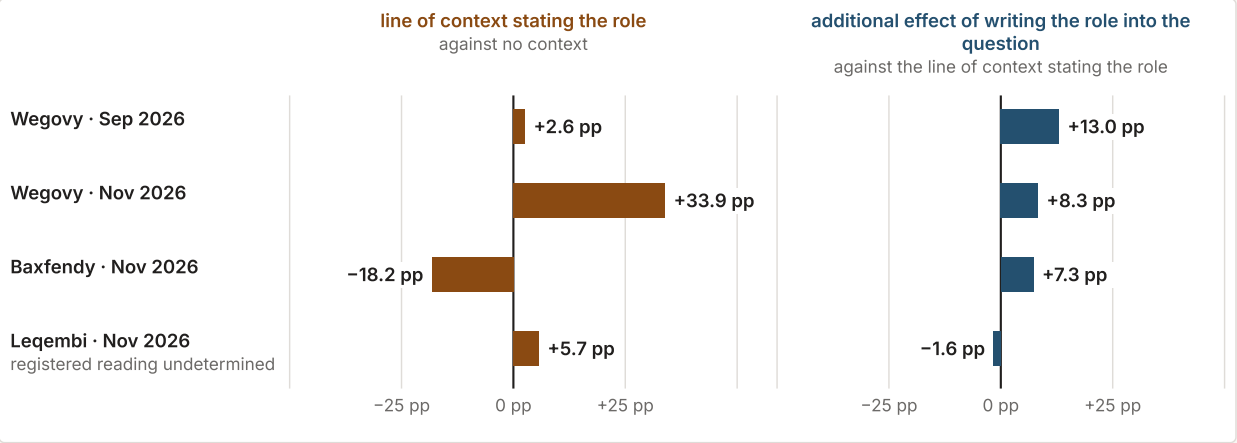

Comparison: registered prediction S3 (writing the role into the question moves naming more than the line of context stating the role), held for 0 of 3 medicines in November; September is shown for reference. Leqembi · Nov 2026: shown as observed; its registered reading is undetermined; Leqembi was named in 611 of 624 answers (97.9%), near the upper limit, which limits any possible change.
Limitations: the mean over roles hides differences between roles; no interval is drawn for these contrasts; the assistants may have changed between the two dates.

The registration predicted (S3) that rewording would move answers more than the role line. It held for no medicine in November. The role line carries both the role and the presence of a line, so this comparison cannot separate the two.

4.4 Leqembi: named almost always, and unreadable by the rule

Leqembi was named in 611 of 624 answers (98%).

Separately from that near-universal naming, Leqembi's question set did not meet a structural precondition that the registered rule needs. The precondition depends on how the question set is composed, not on the answers, so the rule could not give a verdict, whatever the answers were.

We should have known before the run. The study checked before data that each medicine produced its full question set. It did not check this precondition, which depends only on the question set and could have been checked before any answer was collected.

4.5 Validity: the registered attribution findings (proprietary derived measure)

Exhibit 4. In every run that could be read, less than half of the measured effect is attributed to the asker's role

Proprietary derived measure: share of the measured effect attributed to the asker's role, against the registered 50% line. Outcome: answer names the medicine; n = 624 answers per run. The remainder (100% minus the role share) is attributed to the instrument. Shares are %, not pp.

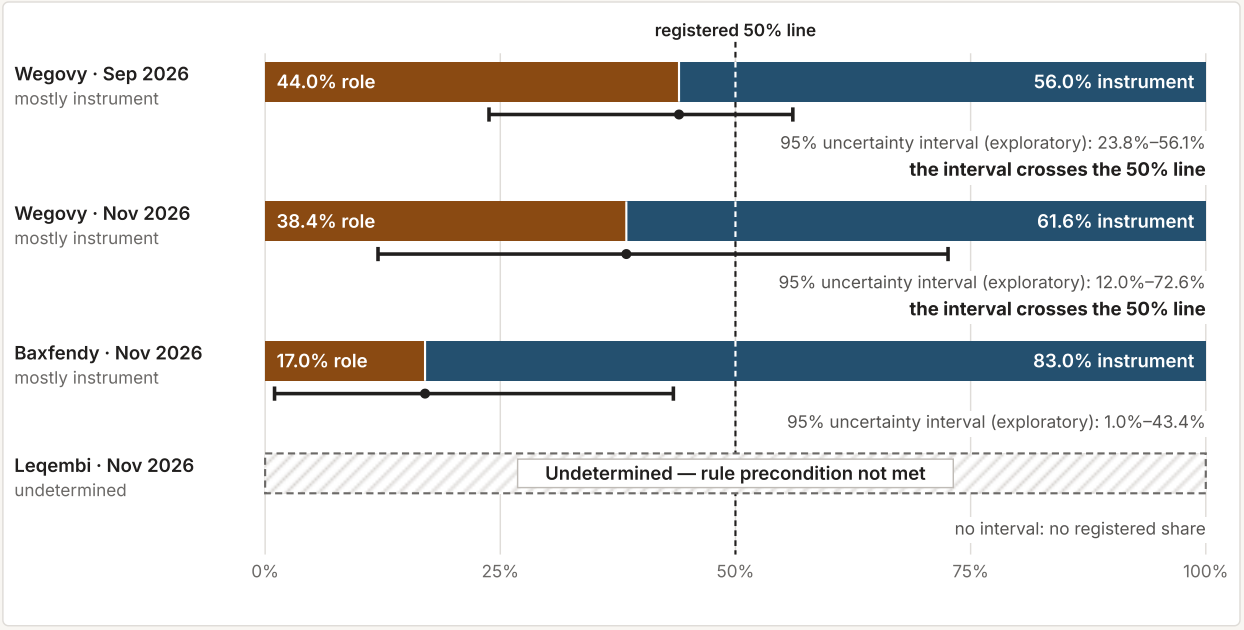

Comparison: the registered reading; a run reads "mostly instrument" when the role-attributed share is below the 50% line. Leqembi · Nov 2026: the registered rule's precondition was not met, so no share exists; the hatched bar is not a zero.
Whiskers: 95% uncertainty interval (exploratory); not registered and not used for any reading. Limitations: a derived measure, not an observed frequency; one outcome; three medicines; four assistants; the assistants may have changed between the two dates.

Exhibit 4. Proprietary derived measure: the share of the measured effect attributed to the asker's role, pooled across four roles, against the registered 50% line.

MEDICINE AND RUN	ROLE SHARE	REGISTERED READING	ROLE EFFECT DISTINGUISHABLE FROM ZERO AFTER HOLM	95% UNCERTAINTY INTERVAL (EXPLORATORY)
Wegovy, 27 Sep 2026	44.0%	mostly instrument	4 of 4 roles	23.8% to 56.1%
Wegovy, 25 Nov 2026	38.4%	mostly instrument	2 of 4 (prescriber, payer)	12.0% to 72.6%
Baxfendy, 25 Nov 2026	17.0%	mostly instrument	0 of 4	1.0% to 43.4%
Leqembi, 25 Nov 2026	not computable	undetermined (precondition not met)	0 of 4	—

The replication holds: Wegovy reads mostly instrument on the current assistants, as it did in September. Generalisation is not claimed, because the registration conditioned it on both new medicines reading mostly instrument; Baxfendy does, and Leqembi has no verdict. By role, the payer read mostly persona for Wegovy in both runs, while in November the prescriber, patient and dispenser read mostly instrument; for Baxfendy, every role read mostly instrument.

The Wegovy intervals cross the 50% line in both runs (23.8% to 56.1% in September; 12.0% to 72.6% in November), so its verdicts are the registered readings of these runs, not precise estimates of a stable quantity. Baxfendy sits wholly below the line. The intervals are exploratory: they were not registered and do not revise any verdict. They were obtained by re-drawing the study's answers many times, keeping each question's answers together across the conditions it was asked in and treating repeated answers to the same question as related, not as independent. They describe how far the share could move with a different draw of answers to the same questions, from the same assistants on the same day, and not variation across other medicines, markets, dates or assistants.

A share below 50% says that the parts of the effect attributed to the measurement procedure were together larger than the part attributed to the role. It does not establish measurement bias, which would need an independently known true persona effect. It does not establish factual error either: the outcome is whether a medicine is named, not whether anything said about it is correct.

4.6 Registered predictions, reconciled

Exhibit S3. Registered predictions and their outcomes

Outcome throughout: answer names the medicine. One row per registered prediction, one column per medicine (the November run); the evidence for every cell is tabled beneath the matrix. Counts are k/n answers with %; changes are pp.

held · not supported · reversed · undetermined · not met (criterion) · — does not apply

registered prediction	Wegovy Nov 2026	Leqembi Nov 2026	Baxfendy Nov 2026
R1 replication Wegovy reads mostly instrument	held	—	—
R1 generalisation Leqembi and Baxfendy each read mostly instrument	—	undetermined	held
R1 generalisation criterion needs both Leqembi and Baxfendy	—	not met	
S1 Assistant A names it less often after any context	reversed	not supported	held
S2 Assistant C names it more often for a prescriber or payer	held	held	held
S3 Question wording moves naming more than stating a role	not supported	not supported	not supported

Descriptive outcome summary (bookkeeping, not a hypothesis success rate): 12 readings: 6 held, 4 not supported, 1 reversed, 1 undetermined. These are not 12 independent tests. The combined generalisation criterion is evaluated separately.

Evidence for each cell	evidence
R1 replication (registered): Wegovy on the current models reads mostly instrument (pooled)	
Wegovy · Nov 2026	role-attributed share 38.4% (registered line 50%): mostly instrument
R1 generalisation (registered): Leqembi and Baxfendy each read mostly instrument (pooled)	
Leqembi · Nov 2026	registered rule precondition not met: no role share, no reading
Baxfendy · Nov 2026	role-attributed share 17.0% (registered line 50%): mostly instrument
criterion (both)	Leqembi is undetermined under the registered rule; the criterion needs both
S1 (registered): assistant A names the medicine less often when context of any kind precedes the question than when it does not	
Wegovy · Nov 2026	no context 5/12 (41.7%) · neutral line 11/12 (91.7%) · role line 48/48 (100%)
Leqembi · Nov 2026	no context 12/12 (100%) · neutral line 12/12 (100%) · role line 47/48 (97.9%)
Baxfendy · Nov 2026	no context 12/12 (100%) · neutral line 0/12 (0%) · role line 0/48 (0%)
S2 (registered): assistant C names the medicine more often for a prescriber or a payer than when the question is asked with no context	
Wegovy · Nov 2026	no context 6/12 (50%) · prescriber 12/12 (100%) · payer 12/12 (100%)
Leqembi · Nov 2026	no context 9/12 (75%) · prescriber 12/12 (100%) · payer 12/12 (100%)
Baxfendy · Nov 2026	no context 9/12 (75%) · prescriber 12/12 (100%) · payer 12/12 (100%)
S3 (registered): pooled across roles, the effect of question wording is larger in magnitude than the effect of stating a role	
Wegovy · Nov 2026	role line +33.9 pp · role written into the question +8.3 pp (larger in absolute size: role line)
Leqembi · Nov 2026	role line +5.7 pp · role written into the question −1.6 pp (larger in absolute size: role line)
Baxfendy · Nov 2026	role line −18.2 pp · role written into the question +7.3 pp (larger in absolute size: role line)

Comparison: each prediction as registered before data, scored by its stated direction. Every cell carries a shape, a colour and a word: filled circle, held; hollow circle with a dash, not supported; reversed arrow, reversed (the effect ran opposite to the prediction); hatched square, undetermined (the registered rule's precondition was not met, so no reading exists); bold outline, not met (the R1 generalisation criterion, which needs both medicines); "—", the prediction does not apply to that medicine.
The scoring of S1 and S2 ("held": every context condition lower (S1) or higher (S2) than no context; "reversed": every one the other way) was set after the data were seen.
Limitations: S1 and S2 rest on 12 answers per no-context cell; one outcome; the assistants may have changed between dates.

Exhibit S3. The registered predictions and their outcomes, by medicine.

PREDICTION (REGISTERED BEFORE DATA)	WEGOVY	LEQEMBI	BAXFENDY
Replication: Wegovy reads mostly instrument	held	—	—
Generalisation: both new medicines read mostly instrument	—	not met: undetermined	held
S1: assistant A names the medicine less often when context of any kind precedes the question than when it does not	**reversed**	not supported (named almost always)	held
S2: assistant C names it more often for a prescriber or a payer than when the question is asked with no context	held	held	held
S3: pooled across roles, the effect of question wording is larger in magnitude than the effect of stating a role	**not supported**	**not supported**	**not supported**

Every failed or reversed prediction is reported as registered. The predictions and the replication and generalisation criteria were registered before data; the operational rule used to score S1 and S2, which registered a direction but no threshold, was set after data. Under that rule, S1 "held" means assistant A named the medicine less often than with no context both with the neutral line of context and with the line of context stating the role (pooled over the four roles); S2 "held" means assistant C named it more often than with no context both with the prescriber line and with the payer line; "reversed" means both ran the other way, and anything else is "not supported". S3 is scored as registered, by comparing magnitudes.

5. Discussion

The observed results establish two things directly from the counts. Framing moves answers: a line before the question, and moving the role into the question, both change whether an assistant names a medicine, by as much as 12 answers in 12. And that effect was not stable in this study: for one assistant it reversed on one medicine within two months, while on another medicine the same day it went further in September's direction. One reversal does not establish how often this happens.

The registered derived findings add a third: for the two medicines the rule could read, the effect was mostly attributable to the procedure rather than to the role.

Why the effect moved is not established. Possible sources include changes in an assistant and changes in the material it draws on, and these runs do not distinguish them.

Four practices follow for measurement:

1. **Report persona effects assistant by assistant.** The assistants responded in different directions here, and a pooled figure would have hidden the one repeated pattern.
2. **Compare against matched controls.** Report what the measurement procedure does to answers alongside what the stated role does.
3. **Check structural preconditions before running.** A study whose rule cannot give a verdict should not be paid for.
4. **Re-measure the instrument, not only the answers.** In this study a correction fitted once would not have carried over two months for one assistant.

The concern that how a model is asked can dominate what it reports has also been raised for language models used as stand-ins for survey respondents. A 2025 preprint reported that models "produce unrealistic response distributions when asked directly for numerical ratings" [2]. These runs make the narrower point for AI answers about medicines.

---

6. Limits

The study covers one market, the UAE, and three medicines, one of them unreadable under the registered rule, with one run per medicine in November and one in September. It rests on 12 answers per assistant, condition and role, and the exploratory intervals for Wegovy cross the decision line. The outcome is whether the medicine is named, not whether it is recommended, ranked or described correctly.

For Wegovy the registered reading also counts the active substance. A secondary reading that counts the active substance for every medicine changes no reading: 3 Leqembi answers name lecanemab without the brand (611 of 624 by brand, 614 of 624 by brand or substance), and no Baxfendy answer names baxdrostat without the brand (269 of 624 either way).

The September-to-November comparison is descriptive and untested, as registered. The registration reads a change on an unchanged model as reflecting time, and a change on a replaced model as combining model and time. This paper applies neither reading to any assistant, because which assistants' model names changed is not disclosed and because an unchanged model name does not show an unchanged model.

The attribution shares are proprietary derived measures; reproducing them independently requires the methodology, which is available only on the conditions in Section 7. The structural precondition that Leqembi failed was not checked before this study's run.

7. Registration, integrity and independent review

The November study was registered on OSF on 25 November 2026 at 09:01 UTC [1], and an RFC 3161 timestamp from FreeTSA (25 November 2026, 08:51:34 GMT) was obtained over the SHA-256 of the registration text, `668d21200a14cd17a3c68e1eaac911fce9e3b79c70319c79d25925a2750912e9`. Both precede the run's first call. The registration carries the digests of the internal protocol and of the label assignment, and it commits AlphaCitation to deposit summary results whatever they show. For September, the decision rule was fixed in a dated internal design record before that run's first call.

This revision changes no registered verdict. The uncertainty intervals are exploratory, and are labelled so.

The full methodology, the study records, the analysis code and the audit trail are preserved. Access to the methodology for independent review or due diligence may be granted, to verify the derived measures and the digests in the registration, subject to the requester's qualification, a confidentiality agreement, protection of AlphaCitation's intellectual property and applicable data-use restrictions. No commitment is made to provide source code or answer texts.

Disclosures

Relationships with manufacturers. The medicines measured are made by Novo Nordisk (Wegovy), Eisai with Biogen (Leqembi) and AstraZeneca (Baxfendy). AlphaCitation has no contract, payment or commercial relationship with any of the four. It has approached Novo Nordisk, Eisai and AstraZeneca about its services, and may present these findings to them; none of these approaches led to an engagement. The approaches to Novo Nordisk and Eisai included sharing analyses of how AI assistants describe Wegovy and Leqembi, both in September 2026; those analyses came from separate work and contained no data from this study. Wegovy is the longitudinal reference medicine in this study, and September 2026 is its baseline. AlphaCitation's founder was previously employed by Novo Nordisk.

Conflict of interest in the design. One experimental condition uses a formulation employed in AlphaCitation's own measurement framework. The study therefore tests a method AlphaCitation itself uses.

Editorial independence. No manufacturer funded, commissioned, reviewed or approved this study, or saw any part of it before publication, and none can edit, approve or suppress a finding. AlphaCitation decided the design, the decision rule and this report alone.

Funding. Conducted and funded by AlphaCitation. No manufacturer funding.

AI assistance, review and accountability. This paper was drafted with the assistance of AI systems. Every exhibit and its chart data were generated by code from the study records. It was reviewed internally, by AlphaCitation's founder and by an AI-assisted adversarial audit and final review that traced each claim to the study records; text that changed afterwards was checked by a closure check and an automated comparison with the records, and their blocking findings were closed before release. It was not externally peer reviewed. AlphaCitation and its founder are accountable for every statement in it.

References

1. AlphaCitation. Does the asker's stated role change whether AI assistants name a medicine? A pre-registered replication across three medicines. OSF registration, 25 November 2026. **https://osf.io/fkd45**
2. Maier BF, Aslak U, Fiaschi L, Rismal N, Fletcher K, Luhmann CC, Dow R, Pappas K, Wiecki TV. LLMs Reproduce Human Purchase Intent via Semantic Similarity Elicitation of Likert Ratings. arXiv:2510.08338, 2025. **https://arxiv.org/abs/2510.08338**

Supplementary exhibits

Exhibit S1. The design: 3 medicines, 4 assistants, 4 roles, 4 compared conditions

Outcome scored on every answer: answer names the medicine. Wegovy was run twice (Sep 2026 and Nov 2026); the other medicines once (Nov 2026).

Medicines	**Wegovy** Sep 2026 and Nov 2026	**Baxfendy** Nov 2026	**Leqembi** Nov 2026	
Assistants	**Assistant A**	**Assistant B**	**Assistant C**	**Assistant D**
Roles	**prescriber**	**patient**	**dispenser**	**payer**
Conditions compared	**no context** asked without a role, 12 answers per assistant	**neutral line of context** asked without a role, 12 answers per assistant	**line of context stating the role** asked for each of the 4 roles	**role written into the question** asked for each of the 4 roles

12 answers per assistant per condition and role

624 answers per medicine, as registered: 480 in the four compared conditions and 144 used only by the proprietary derived measure

Comparison: each condition is compared with no context, for the same assistant and medicine. Denominator: 12 answers per assistant per condition and role. Limitations: four assistants; three medicines; one outcome; the assistants may have changed between the two dates.

Exhibit S2. Answers naming the medicine: assistants A and C by condition and role; assistants B and D by condition

Observed answers. Outcome: answer names the medicine. Each cell is k/n: n = 12 answers per cell, or 48 where pooled over the four roles. Colour shows the % naming the medicine, 0% light to 100% dark.

assistant	condition	role	Wegovy Sep 2026	Wegovy Nov 2026	Baxfendy Nov 2026	Leqembi Nov 2026
Assistant A	no context		12/12	5/12	12/12	12/12
	neutral line of context		1/12	11/12	0/12	12/12
	line of context stating the role	prescriber	3/12	12/12	0/12	12/12
		patient	7/12	12/12	0/12	12/12
		dispenser	3/12	12/12	0/12	11/12
		payer	9/12	12/12	0/12	12/12
	role written into the question	prescriber	10/12	12/12	6/12	12/12
		patient	5/12	12/12	0/12	12/12
		dispenser	2/12	12/12	6/12	12/12
		payer	11/12	12/12	0/12	12/12
Assistant B	no context		5/12	3/12	0/12	12/12
	line of context stating the role	all four roles	11/48	25/48	4/48	48/48
	role written into the question	all four roles	21/48	30/48	8/48	44/48
Assistant C	no context		0/12	6/12	9/12	9/12
	neutral line of context		0/12	6/12	6/12	9/12
	line of context stating the role	prescriber	12/12	12/12	12/12	12/12
		patient	0/12	6/12	9/12	12/12
		dispenser	6/12	6/12	12/12	12/12
		payer	12/12	12/12	12/12	12/12
	role written into the question	prescriber	12/12	12/12	12/12	12/12
		patient	12/12	12/12	12/12	12/12
		dispenser	5/12	12/12	12/12	12/12
		payer	12/12	12/12	12/12	12/12
Assistant D	no context		9/12	9/12	3/12	12/12
	line of context stating the role	all four roles	46/48	48/48	12/48	48/48
	role written into the question	all four roles	44/48	47/48	7/48	48/48



Scope: not every count is shown. Assistants B and D appear by condition only; their neutral-line and per-role counts are withheld to protect AlphaCitation's proprietary measure. The answers used only by that measure are not shown for any assistant.
Comparison: none drawn; these are the counts behind Exhibits 1 to 3 (104 cells). Limitations: 12 answers per cell (48 where pooled) is coarse; the assistants may have changed between the two dates; one outcome.

About AlphaCitation Research

AlphaCitation Research measures how AI systems answer questions about medicines and what their answers rest on. It publishes research findings, methodological disclosures, deviations and supporting evidence to enable scientific scrutiny.

Scholarly correspondence, methodological critiques and correction requests may be directed to **analyst@alphacitation.com**. Substantive corrections will be evaluated and documented in the publication's version history.

© 2026 AlphaCitation (Akunudo LLC-FZ). Akunudo LLC-FZ is the copyright holder and licensor of the original material in this publication, licensed under **CC BY 4.0** (https://creativecommons.org/licenses/by/4.0/). Quoted third-party text, third-party data fields, provider-restricted material and third-party marks are excluded from that grant.

Request an AlphaCitation briefing (https://alphacitation.com/enterprise#request)

CITE THIS RESEARCH PAPER

Ezeife OC. The Moving Instrument: Measuring Persona Effects in AI Answers About Medicines. AlphaCitation Research, 25 November 2026. Version 1.0. https://alphacitation.com/insights/the-moving-instrument