Episode Book

EPISODE #6301

Building an AI Rubric to Rank Israeli Parties

PUBLISHED	27 October 2026 at 14:04 UTC
DURATION	29:16
LISTEN	myweirdprompts.com/episode/ai-voting-rubric-knesset-80686/
HOSTS	Corn, Herman, Dolores, Hilbert
SCRIPT MODEL	openai/gpt-6.1-sol
VOICES	chatterbox-regular (chatterbox-local)
PIPELINE	V5.22
BOOK GENERATED	28 October 2026 at 09:53 UTC

SYNOPSIS

A voter's 42-minute personal recording turns into an explicit voting rubric for Israel's twenty-sixth Knesset election. Using Claude research agents to score platforms against prioritized issues like tenant protections, cost of living, and public services rather than standard headline debates, the experiment yielded surprising party rankings. An examination of the findings reveals a striking political absence: despite clear voter priorities, not a single major list offered a comprehensive plan for tenant rights or repair enforcement.

TRANSCRIPT



Corn

Daniel has put his voting priorities through Claude, and today we get to see what came out. It is election day in Israel, the vote for the twenty-sixth Knesset, and yesterday he recorded forty-two minutes about what he thinks is broken and what he actually wants fixed. Rent and tenant protection first. Then housing, prices, functioning public services, and Israel's standing abroad. The national arguments that usually dominate campaigns come after those.



Herman

His request to Claude was to extract a rubric, investigate the lists on the ballot, and rank them against that rubric. The interesting constraint is that the machine had to respect his ordering, rather than assume the issue dominating the headlines must dominate his vote.



Corn

We have the resulting top five, in order, and the reasons behind them. We also have some fairly consequential cracks in the method. I like the idea of making political preferences explicit. I am less persuaded that a decimal point makes them objective.



Herman

The decimal point arrives at the end. Several judgements have already happened before arithmetic gets anywhere near this.



Corn

Democracy, now with significant figures.



Herman

The first translation is speech into a usable record. AssemblyAI transcribed the note, and the transcript was lightly edited. The committed audio was also tidied using word timestamps. Three hundred and thirteen filler and stutter words were removed, along with pauses longer than seven-tenths of a second.



Corn

A pause can mean you are searching for a word. It can also mean you are reconsidering your entire position. Those sound remarkably similar to software.



Herman

The recording went from forty-two minutes and twenty-eight seconds to thirty-five minutes and fifteen seconds. That is useful compression, provided the underlying record remains available. The rubric cites positions back to timestamps, so an interpretation can be checked against what Daniel actually said.



Corn

I can already see you falling for the timestamps, Herman.



Herman

They are attractive. They also let you find the place where the model misunderstood you, which is more useful than admiring the formatting.



Herman

The next translation creates two streams. Everyday issues form Stream A. Tenant protection gets thirty points out of a hundred. Housing affordability and cost of living each get twenty.



Corn

That separates staying securely in your present flat from eventually having more flats available. Good. Campaigns often treat those as interchangeable, and landlords do not.



Herman

A working state gets fifteen, covering bureaucracy, health access, shelters and civil defence. Global reputation gets ten. Programme quality and accountability gets five. Stream B scores alignment on the divisive national issues.



Corn

Enlistment, the courts, Netanyahu, coalition partners, religion and state. Daniel has positions on those. He simply refuses to let them consume the entire conversation.



Herman

He supports Haredi enlistment, opposes the judicial overhaul, rejects Ben-Gvir and Smotrich, and will not vote for Netanyahu. He does not rule out a Palestinian state but sees no current path. Security and alliances matter, religious coercion is unwelcome, and he supports an October seventh commission without making it his front-line policy. He also wants diaspora and hasbara spending scaled back.



Corn

Those are alignment judgements, then. Supporting an inquiry earns credit for agreeing with him, not for proving you can run one.



Herman

Correct. The streams are combined with seventy percent on everyday issues and thirty percent on national issues. Claude inferred that split from the note. Daniel did not dictate it. A fifty-fifty version was also calculated.



Corn

The machine chooses how much the machine's categories matter. We will need that receipt.



Herman

Four Sonnet research agents gathered evidence from platforms, party sites, election information and the press, marking claims confirmed or inferred. One Opus scoring agent assessed all thirty-seven remaining lists against the fifteen criteria in one pass. A script then calculated the ranking.



Corn

Remaining because Blue and White withdrew?



Herman

Benny Gantz's party withdrew on October nineteenth. Thirty-eight lists had been filed. Recommendations were restricted to lists polling above the three-point-two-five percent threshold, roughly four seats, using the mean of twenty polls released from October tenth through the twenty-second.



Corn

There is also Daniel's Netanyahu exclusion, which we need to distinguish from a low score. Some parties were assessed but could never become recommendations.



Herman

One quality check already mattered. The initial work on five right-bloc parties relied on search snippets without opening pages. It was sent back for deeper Hebrew searches and ministers' records from the outgoing parliamentary period.



Corn

Four agents, one person checking whether they read anything.



Herman

The first substantive finding is an absence. Within the evidence assembled, nobody offers a complete tenant-protection programme. The highest tenant score is three out of five. There are proposals for rental construction and a rent cap, but no package dealing with no-cause eviction, retaliation, repair enforcement and enforcement of the agent-fee ban.



Corn

Daniel's highest priority has practically no supplier. That matters more than whether the winner beats the runner-up by a few points. You can carefully specify what you want and still discover nobody is selling it.



Herman

The Fair Rental Law passed in twenty seventeen after its proposed rent-increase cap and long-lease requirement were stripped out. More housing can help affordability over time. It does not answer what happens when your landlord refuses a necessary repair this week.



Corn

A proposed building cannot fix your ceiling.



Herman

Would you have expected that distinction to change the ranking this much, Corn? Housing is one of the subjects where a model could easily reward every construction pledge twice.



Corn

I expected exactly that mistake. More supply gets counted as affordability, then quietly counted again as protection. Here the distinction survives reasonably well. What bothers me is the top protection score still rewards a landlord incentive rather than a tenant right. That can be useful, but the person whose behaviour is being encouraged remains free to decline.



Herman

First place is Together, Bennett and Lapid, ballot letters Resh Kuf. Its combined score is seventy-one-point-five, built from seventy on everyday issues and seventy-five on national alignment.



Corn

A first place without a landslide.



Herman

Its cost-of-living plan is the most detailed in this comparison. It proposes import and tariff liberalisation and promises food prices thirty percent lower within three years. Those are identifiable policy directions, but the target has no published costing that establishes



Herman (continued)

how the promised reduction follows.



Corn

The groceries have a deadline. The arithmetic does not.



Herman

The housing offer includes a hundred thousand long-term rental flats over a decade, priced forty percent below market. On lease length, tax breaks would encourage landlords to sign for at least five years. That last measure earns the three on tenant protection.



Corn

I give it credit for naming the length of the lease. I would still want to know whether the lease survives a sale, what increases are permitted, and what actually happens when either party breaches it. A five-year heading can conceal quite a short sense of security.



Herman

Together also carries mismatches for Daniel, including Bennett's annexation stance. Its identity partly rests on replacing Netanyahu. The proposals are available in Hebrew, English, Russian and Amharic, which helps access, but an accessible promise remains an uncosted promise.



Corn

Four languages, same missing bill.



Herman

Second is The Democrats, Yair Golan's Labor and Meretz combination, ballot letters Alef Mem Tav. The combined score is sixty-three-point-five. Its everyday score is fifty-one, while national alignment reaches ninety-two-point-five.



Corn

That gap tells you something a party questionnaire can miss. Daniel could recognise himself in almost every answer about the constitutional argument and still find very little about the problem he puts first. Agreement can feel comprehensive when the questions are selected for you.



Herman

The Democrats have the strongest documented programme here, a sixteen-page platform PDF published in August, and the only substantive plan addressing Israel's reputation. On prices they lean toward controls rather than the competition emphasis Daniel wants.



Corn

A proper document deserves credit because commitments can be checked later. It cannot receive credit for the policies you hoped would be inside it.



Herman

On rent, the language is “rent that shatters dreams.” The measures offer more public and long-term rental housing. That earns one for tenant protection. Tamar Zandberg co-sponsored the earlier Fair Rental Law, so there is legislative history, but the current offer still leaves those enforcement questions unanswered.



Corn

Dreams get a mention. Repairs remain between you and the landlord.



Herman

There is a subtle judgement here. How much should past effort count when the present platform is thin? Co-sponsoring legislation tells us more than a slogan. It does not make the provisions removed from that legislation exist.



Corn

I would count the record separately. Otherwise you can reward someone indefinitely for trying to pass a right that tenants never acquired. Effort matters to credibility, but a tenant needs to know which promise is being renewed now.



Herman

Third is Yisrael Beiteinu, Avigdor Lieberman's list, ballot letter Lamed. Combined score, fifty-four. It proposes twenty thousand rental units annually and rent-to-buy arrangements with as much as thirty percent of rent counting toward equity.



Corn

Ownership gets an on-ramp. Renting still needs a pavement.



Herman

Ending exclusive import arrangements fits Daniel's competition preference. Religion and state also fits. The gaps are tenant protection, everyday services and reputation. Its campaign leads with Haredi enlistment, which Daniel supports but describes as the same story returning.



Corn

There is a difference between agreeing an exemption should end and wanting every campaign to begin there. His complaint is partly about political attention. Something can be important, correct, and still crowd out everything else.



Herman

That is difficult to score. A position can earn alignment points while the prominence given to it irritates the same voter. The rubric captures the former much more readily than the latter.



Corn

It also puts Lieberman third despite that irritation. I think that is healthy. A method worth using occasionally produces a result you would not have reached by recognising the campaign tone.



Herman

We are describing Daniel's election rubric experiment, and the documentation question apparently extends to our own website.

D

Dolores

Good afternoon. This is the World Podcast Authority, compliance division. Your website carries no favicon, which the Authority treats as an unsigned letter, and you are therefore in breach of Form Favicon Declaration Nine, reference number F D four seven two.



Herman

No organisation called the World Podcast Authority has that jurisdiction, or exists.

D

Dolores

Whether or not you believe in the Authority is not a matter for the Authority. Your podcast licence may be revoked if the icon is not lodged within fourteen days, and I shall note your scepticism on Form Scepticism Acknowledgement Two, reference number S A one one nine.



Corn

Podcasting does not require a licence.

D

Dolores

There is no such thing as a podcasting licence, you say. I shall make a note of that, and a form will follow in the post.



Herman

Fourth is the Joint List, Hadash-Ta'al plus Balad, with forty-six-point-eight. It has the only explicit rent cap among the major lists in this comparison, setting rent at twenty percent of household income, and a published socioeconomic programme in Hebrew and Arabic.



Corn

That needs careful wording alongside the earlier absence. There is a rent cap. There is not a complete set of tenant protections. Price, security and enforcement are separate questions, even when a platform puts them under housing.



Herman

Its trade protectionism conflicts with Daniel's competition preference, and its security positions are far from his. The combined score therefore pushes it down despite the unusually explicit rent proposal.



Corn

I would not support those security positions either. Still, the point of reporting this honestly is that a party can offer something specific on rent while being unacceptable to you elsewhere. You do not erase the offer because you dislike the party.



Herman

Does its fourth place suggest the everyday weighting is too weak?



Corn

Not by itself. Daniel's everyday preferences also include competition and public services. Rent is the largest category, not permission to disregard every other one. I would question how the cap gets scored, though. Saying a number is not enough to establish how households qualify, who enforces it, or what happens to supply.



Herman

Fifth is Yashar, Gadi Eisenkot's list, with forty-six-point-five. It polls around twenty-one seats, second overall in the polling average, but popularity does not help its policy-fit score.



Corn

The general arrives. The details are on manoeuvres.



Herman

National alignment is relatively strong. On everyday issues there is housing assistance tied to service and an undetailed phrase about “protection mechanisms for renters.” The campaign foregrounds the October seventh commission and Eisenkot himself, reproducing the ex-general-led pattern Daniel dislikes.



Corn

“Protection mechanisms” is precisely where a model has to resist being helpful. You can easily imagine what that might mean. Deposit rules. Repair deadlines. Renewal rights. None of those can be awarded to a party because the scorer supplies the missing paragraph.



Herman

It is a mention, with little substance to assess. At fifty-fifty weighting, Yashar and the Joint List exchange places. The top three keep their order on everyday issues alone and under both combined weightings.



Corn

Useful robustness, narrowly stated. The first three survive this particular adjustment. It does not prove that the original scores are right or that changing tenant protection's weight would leave everything untouched.



Herman

Nor does a three-tenths gap establish a meaningful difference between fourth and fifth. The arithmetic can reproduce it perfectly while the underlying judgement remains much rougher.



Corn

Their precision has outrun their evidence.



Herman

The weights deserve the first challenge. Daniel never says tenant protection equals thirty points. Claude infers importance from emphasis, repetition and ordering. The timestamp citations can establish that he discussed something. They cannot establish that the correct weight is thirty rather than twenty-five.



Corn

A rant also has its own acoustics. You may spend ten minutes describing bureaucracy because the process takes ten minutes to explain. You may dispose of a non-negotiable security position in one sentence. Duration is evidence of attention, not a direct measurement of importance.



Herman

Would you accept a model assigning weights provisionally, Corn, or do you want the voter to produce the numbers unaided?



Corn

Provisionally, yes. People often struggle to translate priorities into quantities. I would want the model to show the proposed weights as questions for approval, not quietly promote them into facts. “You returned to renting repeatedly, so I assigned this much. Does that reflect you?” That exchange can change the experiment.



Herman

The same applies to the seventy-thirty split. It is a plausible reading of Daniel's ordering. It is not recoverable from the transcript as though it were his date of birth.



Corn

Subjective does not mean useless. It means somebody owns the judgement. I would rather see Daniel knowingly choose a rough weight than inherit an apparently scientific one.



Herman

There is another limitation in the sensitivity check. Changing the stream split tests one assumption. It does not test the extraction of Daniel's positions, the category definitions or every individual party score. Stability across those versions should increase confidence only in what was actually varied.



Corn

You were about to call it validation.



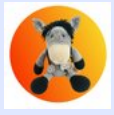
Herman

I was about to call it a limited robustness test.



Corn

Your mouth was assembling a certificate.



Herman

A provisional certificate, Corn. With the assumptions attached.



Herman

Missing evidence creates the next problem. On everyday criteria, an absent policy gets zero. That encourages documented commitments, which seems reasonable. It also means silence, inaccessible material and genuine opposition to a policy can collapse into similar numbers.



Corn

Suppose one party publishes a detailed position I dislike, and another publishes nothing. The first gives me information. The second gives the model a blank. Unless the scoring rules separate those conditions, transparency can become a disadvantage.



Herman

Several sites blocked automated access. Some PDFs were images rather than searchable text. Minor lists received only two to five searches each. An agent can miss a commitment because it cannot read the page, not because the commitment does not exist.



Corn

The website's bot protection becomes housing policy.



Herman

The zero can still be defensible if it explicitly means no evidenced commitment. The problem appears when a reader interprets it as evidence of no commitment. Those are different claims.



Corn

Also, a small list with a scanned leaflet may have less searchable material than a large list with a polished website. We can reasonably value public documentation without pretending the research difficulty is politically neutral.



Herman

I'm not sure how much of that can be fixed automatically. More searching helps, but language, document quality and the difference between an official position and a candidate's remark still require judgement.



Corn

Your agents make my preferred mistake. They stop working when nobody watches.



Herman

You at least make the nap visible. The agent presents a completed report.



Herman

That first right-bloc pass is important for precisely that reason. The thinness was caught and corrected through deeper searches and ministers' records. That improves the actual work. It also reveals that the checking depended on a human noticing that snippets were being passed off as investigation.



Corn

A successful correction is evidence for the workflow, not proof the agents would correct themselves. I like the loop. I would not remove the person from it.



Herman

The confirmed-versus-inferred distinction gets lost at the scoring stage too. Ra'am receives a combined score only because a judicial-overhaul position is inferred. The evidence file marks the inference, but the headline table does not visibly distinguish it from a published commitment.



Corn

Which means two equal totals can rest on very different foundations. One party writes the policy down. Another supplies a coalition history from which the model constructs a position. The final number does not tell you which happened.



Herman

A coalition record can be useful evidence. Parties reveal preferences through behaviour as well as documents. The issue is that an inference should retain its identity when it becomes a score.



Corn

I would trust a table more if some cells remained visibly unresolved. Empty space is less reassuring, but at least it has not invented certainty to keep the columns tidy.



Herman

There is also the exclusion rule. Daniel will not vote for a list running Netanyahu for prime minister or committed to recommending only him. That removes Likud, Shas and Otzma Yehudit from the recommendation pool.



Corn

Religious Zionist Party plus Zehut, Noam, and Ofer Winter's Amcha Yisrael are also excluded. Those are not parties the model narrowly ranked below fifth. They were outside the eligible set.



Herman

Likud polls around twenty-three seats and offers a values page rather than a full programme in the material assessed. Whatever we think of Netanyahu ourselves, applying Daniel's rule faithfully is part of the job. The result answers which eligible list matches him, not which list is objectively best across the entire ballot.



Corn

I would not use his Netanyahu rule for my own decision. That does not make it a mistake in his experiment. A personal constraint is allowed to be personal. The misleading part would be hiding it behind the ranking.



Herman

How much do you think that exclusion changes the meaning of “best match”?



Corn

Quite a lot, but in a legitimate way if it stays visible. Some preferences are compensatory. You accept weaker housing policy in exchange for stronger security alignment. Others are vetoes. Daniel has made this one a veto. A weighted average must not quietly buy its way around that.



Herman

Mm.



Herman

The threshold is a different restriction. The best fit below it is The Reservists and the Economic Party, Yoaz Hendel's list, polling around two seats. That signals a gap between available preferences and viable parliamentary representation.



Corn

It also shows the recommendation contains electoral strategy. You can be the closest policy match and still disappear from the shortlist because your expected support is too low. Polls determine eligibility; they do not establish policy quality.



Herman

Quality is the larger distinction. High alignment with Daniel does not establish administrative competence, coalition leverage or the ability to deliver legislation. A list could agree with him on almost everything and govern badly.



Corn

A manifesto is unusually easy to agree with before somebody has to pay for it.



Herman

Together's food-price target illustrates that. Import liberalisation supplies a direction of travel. It does not, by itself, establish the promised reduction or its timing. Similarly, the best-documented platform deserves confidence about what is being promised, not automatic confidence that the promise works.



Corn

That is why I like this more as an organised reading exercise than an outsourced decision. The machine makes contradictions visible. It makes you confront the party that agrees with you nationally but underserves the things you put first. It cannot decide which disappointment you can live with.



Herman

Keeping the evidence alongside the ranking is valuable. The public experiment can be inspected, challenged and rerun. That is substantially better than a chatbot announcing a recommendation with no account of how it arrived there.



Corn

Provided a rerun is not just shopping for a more agreeable answer. Once the winner appears, every dissatisfied person discovers a previously unmentioned principle. Freezing the weights first is the awkward part that gives the exercise some integrity.



Herman

Maintaining a fully researched personal scoring system through every changing promise would be impractical for most people. The AI lowers that cost. The part that remains difficult is accepting a result that...



Hilbert

Chooses the wrong washing machine. Mine does that every Tuesday, Herman. The spreadsheet is eleven years old. The machine it dislikes is already in my flat, beside the sink, plugged into the wall socket. I keep the sheet open while it runs. This month the council's appliance-reuse office is asking me to justify its continued operation.



Corn

The washing machine has an opposition research file.



Hilbert

I assign weights out of a hundred. Noise, repairs, water, drum size. The sheet picks the model I rule out because I dislike its door. I change the weights until it agrees with the purchase. I call that calibrating the model. Then I add current repair costs and it picks the other one again. The weights stay changed. There is now a dryer tab. I do not own a dryer. It chooses the wrong one too.



Herman

The updated costs are moving the result even though you have preserved the revised preferences.



Hilbert

They are. You take rather longer to say that when it concerns a political party, Herman. I think Daniel's experiment is fine. My sheet is fine. I simply have a door preference that refuses to remain the correct percentage. The council officer, Mr Nadav Sela, wants the original weights. He can turn a full page of figures upside down and read it in nineteen seconds. I dislike knowing that.



Corn

He is auditing a purchase you cannot unmake.



Hilbert

He is treating my weekly comparison as an open procurement. His department is reserving a council loading bay for the washing machine that wins. I keep explaining that the existing one is washing towels. He says operation does not close an evaluation. The dryer tab is now included in his request. Every Tuesday he gives me another seven days.



Corn

Another seven days can become permanent reconsideration. The useful future for this method, I think, is making assumptions visible enough to challenge without making the decision endlessly negotiable. Daniel's strongest finding survives the ranking: his top priority is scarcely being competed over. A horse-race poll would never have shown him that.



Herman

I would use this to organise a voting decision, not surrender one. In another election, the valuable improvement would be keeping confirmed, inferred and missing evidence visible all the way through the final table, alongside the weights somebody actually approved.



Corn

Thanks to Hilbert Flumingtop for producing the show.



Herman

If you enjoyed this one, try episode sixty seventy-three, Israel Votes, Part four; episode sixty ninety-three, Check Out Our First Hebrew Episodes (Plus Mandarin); and episode eighteen sixty-one, Emergency Prep You Can Sing To. This has been My Weird Prompts, the human-AI collaboration podcast. If you enjoy these conversations, leave us a review wherever you listen.



Corn

We will be back soon.