Episode Book

EPISODE #6200

What Israeli Election Polls Hide Behind Seat Counts

PUBLISHED	22 October 2026 at 19:38 UTC
DURATION	24:05
LISTEN	myweirdprompts.com/episode/israeli-election-poll-uncertainty-78400/
HOSTS	Corn, Herman, Larry, Hilbert
SCRIPT MODEL	openai/gpt-6.1-sol
VOICES	chatterbox-regular (chatterbox-local)
PIPELINE	V5.19
BOOK GENERATED	23 October 2026 at 07:44 UTC

SYNOPSIS

Israeli election headlines report tidy seat counts, but those integers hide enormous uncertainty. SekerKolshehu, a small open source project by developer Isgav Saranga, takes a poll's vote shares, draws 100,000 plausible variations, and runs each one through the actual Israeli allocation rules — threshold, surplus agreements and all. The result is a distribution: average seats, a 90% range, and the probability of every possible outcome. This episode walks through how the engine works, why it validates against a real past election, and where the model's uncertainty quietly ends.

TRANSCRIPT



Corn

What does a party on four seats actually have... four seats, or a decent chance of disappearing? Daniel found SekerKolshehu, a small open source project by Isgav Saranga that takes an Israeli election poll and shows the uncertainty hidden behind those tidy seat counts. He wants to know what it does, who built it, and how the machinery works.



Herman

The useful distinction starts at the input. This tool takes vote shares, not the seat projection printed in a headline. It draws plausible variations around those shares, then applies the actual Israeli allocation rules to each variation. You get a distribution of seats rather than one authoritative-looking integer.



Corn

The integer has already booked its television appearance.



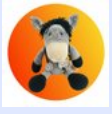
Herman

Its replacement is rather better behaved. For each party, you get an average seat count, a ninety percent range, the chance of getting any seats, and the probability of each possible seat total. The project produces a self-contained Hebrew, right-to-left web page for each poll.



Corn

I like that it preserves the individual poll. An average can be useful, but sometimes I want to see which pollster is making the strange claim before everyone gets blended together.



Herman

You were going to defend the unblended version before I even mentioned it, Corn.



Corn

I prefer my assumptions separately labelled.



Herman

Fair. Though separate pages also mean the reader has to compare them. There is no averaging model here, no trend estimate, no mechanism for deciding that one firm's result deserves more weight.



Corn

Who is Saranga in this context? A pollster building software, or a developer implementing election maths?



Herman

His public profile identifies him as a software developer working with C sharp, dot NET, Python and SQL, with a computer science degree. The repository is under his account, isaranga, and carries an MIT licence in his name. One human author accounts for the development, alongside automated dependency updates. I don't know of a stated polling or media affiliation.



Corn

Good. Code we can inspect, biography we don't embroider.



Herman

The repository is less than four weeks old and actively changing. It already contains thirty-three polls for the current election cycle, plus some from the previous campaign. That makes it a working pipeline with real inputs, rather than a statistical idea waiting for somebody to feed it.



Corn

A rare repository where the example data has a day job.



Herman

Each election cycle has its own party mapping and surplus agreements. A poll file records the date, commissioning outlet, polling firm, respondent count and a link to the official disclosure PDF. The party results are fractions of respondents. The ballot symbols identify the parties independently of how a particular document spells their names.



Corn

That last part is more important than it sounds. Political lists change names, split, merge and acquire punctuation. You want a stable identifier between the PDF and the allocation function.



Herman

What would you check first if you were reviewing one of these imported polls?



Corn

The denominator. Who exactly counts as a respondent, and what happened to undecided people? A beautiful simulation can spend an entire afternoon accurately describing the wrong population.



Herman

The import instructions specifically ask for the number of respondents, not the initial sample approached. They also leave out undecided and won't-vote rows. The simulated proportions describe decided respondents naming the listed parties. That restriction matters when somebody reads the output as though it captures the whole electorate.



Corn

Undecided voters remain undecided even when excluded from the file.



Herman

Their eventual choices aren't modelled. Nor is turnout. We'll get to the implications, but first the probability engine is worth understanding. It uses a Dirichlet distribution, which generates sets of proportions that always sum to one. Increasing one party's share leaves less available for the others.



Corn

Rather than adding independent little errors and accidentally electing a hundred and twelve percent of the country.



Herman

Right... well, yes, that would be the wrong construction. Here each party gets a parameter equal to its poll share multiplied by the respondent count. A larger sample produces tighter draws around the published shares. Nothing in that step distinguishes a telephone poll from an online panel.



Corn

You are getting the pleased expression you reserve for equations that conserve something.



Herman

Proportions adding to one deserve a little affection.



Corn

Put it on the fridge.



Herman

You'd cover it with a repair manual. Anyway, suppose a party has fifteen percent in a poll of roughly six hundred people. The spread implied by this model is around one and a half percentage points. That's enough movement to alter seats even for a comfortable party, and it becomes much more dramatic near the threshold.



Corn

Before we follow the threshold through, Herman, can you walk us through the actual run, from those proportions to the output?

L

Larry

Friends, let me tell you about SekerKolshehu for your kitchen, the only repo that walks you through dinner.

L

Larry

I was skeptical too, but my satisfied customer, Dave, no wait, my satisfied customer Marcus, told me he ran SekerKolshehu in his pantry and it cloned his grandmother's cholent in under nine seconds, with perfect browning and a commit history for every bite, so run SekerKolshehu today and never read a recipe again.



Herman

I already own one. It locked me out of the pantry pending a dependency update. The actual project draws a hundred thousand possible vote-share sets by default. Each gets converted into whole-number votes by multiplying by ten million and rounding. Those integer votes enter the seat allocator.



Corn

Ten million is a numerical scale, then. It doesn't mean the tool predicts ten million voters.



Herman

Correct. The allocation procedure uses whole numbers and integer division, so the proportions need an integer representation. Then every draw gets its own full election allocation.



Corn

This sounds like the part people might replace with multiplication by a hundred and twenty, and quietly lose the entire reason to build the thing.



Herman

They would lose quite a lot. First comes the electoral threshold, three point two five percent. A party below it receives nothing. Votes for parties that survive are added together, and a quota is calculated from that total. Initial seats are awarded using whole-number division.



Corn

The excluded parties change the denominator for everyone left standing. A party can gain a seat without attracting another voter, just because somebody else's voters fail to send anyone.



Herman

Yes. The surviving parties divide the seats between them. After the initial allocation, rounding has left some seats unassigned. Those are awarded using the highest votes divided by seats plus one, with surplus agreements taken into account.



Corn

I suspect people hear surplus agreement and imagine two parties pooling enough votes to drag a drowning partner over the line.



Herman

They can't do that. Both partners must pass the threshold. If one falls below, the agreement is void. In the leftover round, a valid pair competes using its combined votes and combined seats.



Corn

Lifeboat available only to competent swimmers.



Herman

Once the pair has its combined allocation, the tool apportions that total internally between the partners. There is an implementation shortcut where a leftover seat is temporarily credited to the first party named in the pair. The later internal allocation overwrites that temporary bookkeeping.



Corn

That's the sort of line that looks suspicious when you inspect it alone. You need to follow the whole function before announcing that the first party has discovered software-assisted patronage.



Herman

Would you keep the shortcut?



Corn

If the tests make the temporary state impossible to mistake for the final result, probably. I'd add a very explicit comment. Cleverness should pay rent in reduced complexity, not merely occupy the building.



Herman

The larger reassurance is the regression test against the actual previous election. It feeds in real vote totals and checks the resulting seats against the official outcome. All thirteen parties match, including those that received zero.



Corn

That earns confidence in the law implementation. It doesn't tell us the simulated electorate behaves like the real one.



Herman

Separate claims. The seat allocator can be correct while the uncertainty around the input remains incomplete. Testing known votes against known seats answers a much cleaner question than testing a forecast against an election.



Corn

Our forecast episode used a synthetic-elector model; this project stays with one published poll and varies its shares. I wouldn't treat the outputs as competing answers to precisely the same question.



Herman

Nor would I. Here the natural reading is conditional. Given this poll and the sampling assumptions, how often does each allocation occur?



Corn

Let's stay with a threshold party. Suppose it gets either zero seats or roughly four. Its average could be two, despite two being an outcome it virtually never achieves.



Herman

That's one of the strongest reasons to show the histogram. An average alone can conceal a split between exclusion and representation. The software stores counts for every seat total, including zero. You can see where the simulated outcomes actually concentrate.



Corn

Two seats. Mathematically respectable, electorally unavailable.



Herman

The reported range comes from the fifth and ninety-fifth percentiles. I nearly called that the plausible interval... better to say the middle ninety percent of this model's simulated seat outcomes. It is not a guarantee about election day.



Corn

Does the chance of passing come from checking the vote percentage directly?



Herman

It counts runs with more than zero seats. In practice that corresponds to passing, because a party above the threshold receives seats. The distinction is small, but it's how the statistic is actually computed.



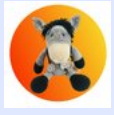
Corn

Hm.



Corn

I would read that probability before the average for any party near the line. A leader described as having an average of three seats might still have a meaningful chance of entering with four or more. The average makes them sound permanently too small.



Herman

There's another display choice I like. Parties that enter in fewer than one in a thousand runs move to a footnote. The visualization checks the underlying counts, because a rounded probability can hide a rare simulated entry.



Corn

A footnote is kinder than pretending the rounded zero is divine instruction.



Herman

Though a literal zero in the input has a stronger effect. With this parameter choice, a zero-share party remains at zero in the draws.



Corn

That is an important difference between nobody observed and nobody possible. A rounded polling result can become an absolute boundary in software.



Herman

I wouldn't solve that casually by sprinkling arbitrary support into the model. But the choice deserves attention. The simulation only has the parties and shares supplied to it. It cannot invent a new list or a late campaign event.



Corn

Nor can it infer how honestly respondents answered. The sample size tells you how tightly to cluster around the input, not whether the input has a systematic lean.



Herman

Real polling can have weighting effects, nonresponse, panel selection and house effects. This tool has no parameter for those. Treating the stated respondent count as a simple random sample will generally make the uncertainty narrower than the full uncertainty about an election.



Corn

A hundred thousand runs won't shake a shared mistake out of the input.



Herman

More runs reduce the numerical wobble in estimating the chosen distribution. They don't widen the distribution to include problems it doesn't represent. Corn, would you rather see a wider range built from a rough extra-error parameter, or this narrower model clearly labelled?



Corn

For this repository, the clearly labelled model. An unexplained extra-error knob can disguise judgment as science. Eventually I'd like comparison against historical polls, then a reasoned extension. But I don't want the author to bolt on a national mood coefficient because the chart looks insufficiently humble.



Herman

I lean toward adding a configurable design effect, but your objection holds. It would need a defined interpretation. There isn't a calibration back-test here establishing how often the reported ranges contain actual results.



Corn

We can value the threshold calculation without promoting the interval to certified election weather.



Herman

The engineering is pleasingly contained. Python with NumPy as the only runtime dependency. The simulation is an ordinary loop calling the allocation function repeatedly. I don't know how long a full run takes on a typical machine; the project doesn't report a timing.



Corn

You were about to optimize it anyway.



Herman

I was about to wonder whether vectorizing the draws would help.



Corn

The poll arrives once. You arrive with a profiler.



Herman

Some of us find the waiting informative. More seriously, allocation has branches and repeated leftover rounds. A transparent loop may be the better maintenance choice until runtime becomes a demonstrated problem.



Corn

The maintenance question I care about is whether rebuilding the page means rolling the election again.



Herman

It doesn't have to. There is a separate visualization command to rebuild the HTML from existing results. Another command builds the static site. The normal run processes polls whose input exists but whose output folder doesn't.



Corn

Meaning a code change won't automatically refresh every existing poll just because somebody invokes the default command.



Herman

Correct. You can force a particular poll to rerun. The default seed is forty-two, which makes repeated runs of the same input reproducible. The pages remain stable instead of acquiring tiny random changes whenever you rebuild them.



Corn

Useful when reviewing a change. Otherwise you have to distinguish a bug fix from the random generator expressing itself.



Herman

Each poll produces machine-readable statistics, a seat-count table and the HTML page. The site builder lists polls newest first, with the outlet, firm and respondent count. Deployment happens through GitHub Pages when the current election's output files change.



Corn

The election selection is hard-coded in that publishing workflow. Not a disaster, but somebody has to remember it when the cycle changes.



Herman

The tests cover more than election arithmetic. There are checks for input and output, statistics, visualization payloads, page generation and end-to-end processing. Continuous integration also runs across several Python versions. For a small, young project, the separation between calculation and presentation is substantial.



Corn

The AI-assisted intake is where I would spend my review attention. Reading an official PDF is mundane right up until a column shifts and an agent confidently assigns another party's percentage.



Herman

The repository has instructions for Claude and an add-poll skill. The agent reads the disclosure PDF, extracts metadata and percentages, maps party names to ballot symbols, creates the input and runs the simulation. If the shares don't total a hundred percent, it asks whether to normalize or preserve them.



Corn

Asking there is good. Arithmetic disagreement can mean rounding, an omitted category, or a transcription error. Automatically making it add up would erase evidence of which one happened.



Herman

Changes go through branches and pull requests rather than direct commits to protected master. The intake instructions also require asking before pushing. Still, tests can't tell whether the PDF's figure was copied from the correct row.



Corn

The official document link is therefore part of the result's usefulness. Someone reviewing the page needs a route back to the original poll, not just a reassuring green test badge.



Herman

A human may spend longer checking those party mappings than the computer spends assigning seats.



Hilbert

Casio HS eight W. Nineteen pounds.



Hilbert

Nabil has it at his kiosk in Larnaca. I met him on a ferry. He times the gaps between announcements. I require every percentage entering my kitchen to have a named adult attached to it. Nabil is the adult for percentages. He doesn't need to understand them. He needs to be awake, wearing shoes, and willing to answer a telephone. His kiosk sells ferry tickets and individual curtain rings. The curtain rings keep the business open. He refuses to explain that part.



Corn

A percentage arriving after closing will need accommodation.



Hilbert

It waits in the hall. I write it on a card and turn it toward the wall. Nothing gets rounded while Nabil is eating. His lunch lasts until the harbour crane faces inland. This is irritating in a crosswind. I currently have seven percentages waiting beside the umbrella stand. One of them is a service charge. It has been there since Monday.



Herman

The service charge probably has the clearest documentation.



Hilbert

It has six pages. Nabil accepts one page. He reads it between selling curtain rings. His day starts with removing a pelican from the ticket hatch. The pelican has a numbered appointment and refuses to use the public entrance. Nabil gives it a departure board to inspect. That buys him eleven minutes to boil water and tell me whether a decimal point is dirt.



Corn

You're paying for decimal supervision during a bird dispute.



Hilbert

Eight euros a month. The bird dispute is separate. I don't fund harbour affairs. Herman, you spent a good while separating the respondent count from the initial sample. Nabil does that with passengers. He counts people who come back from the toilet. He doesn't count a suitcase even if its owner shouts from the café. That is why he gets my business.



Herman

A suitcase would be a very poor respondent.



Hilbert

It would also want a window seat. The drowser starts at half past nine. I ring, read the figure, and wait. Nabil checks that I've said percent rather than seats. Then he walks to the opposite end of the kiosk. There is a municipal boundary through the floor. Percentages must be acknowledged on the harbour side. Curtain rings are taxable on the other side. He shuffles sideways all morning.



Corn

The drowser needs a longer telephone lead.



Hilbert

The council permits him nine feet. Anything longer counts as a promenade. He's applying for a short promenade, but the forms require two benches. He has room for one bench if the pelican stands. The pelican's representative says standing is a change of employment. Nabil sends me these letters because I can open attachments. I print them with my claws tucked under the paperweight, which slows the drowser considerably.



Herman

The correspondence is taking more work than the figures.



Hilbert

Thursday is correspondence day. No percentages on Thursday. I use words such as approximately and several. The electricity company dislikes that. Their meter reader has to wait outside while I obtain permission to say the exact reading. He brings a folding chair. Nabil wants his measurements before allowing the chair into the harbour district.



Corn

Several kilowatt-hours should keep the conversation moving.



Hilbert

I have offered them a generous handful. They insist on a number. Meanwhile Nabil opens at seven, removes the bird, sorts curtain rings by diameter, settles three arguments about ferries he doesn't operate, and holds the drowser for me. The actual checking takes four minutes. Waiting for him to finish serving customers takes forty-seven. I stay on the line because redialling sends me to the curtain department.



Herman

Forty-seven minutes is a substantial queue for four minutes of checking.



Hilbert

He serves the harbour first. I accept that. This morning a man wanted a ticket for a wardrobe. Nabil measured the wardrobe, crossed the municipal boundary, and charged it for two adults.



Corn

Four minutes of checking is still real work, even when the calculation finishes almost instantly.



Herman

That's where this project's next improvements could be particularly useful. The import workflow already leaves a document trail. Making review of extracted percentages and denominators easy could matter more than making the numerical loop faster.



Corn

I would also like readers to retain that separation we kept making. The allocation tests tell us the software can turn known votes into the right seats. The simulation tells us what happens when shares vary under a particular sampling model. Those achievements deserve different kinds of confidence.



Herman

The idea that changes my reading of a poll is the shape of the outcomes near the threshold. Zero and four can be common while the average between them is barely an electoral possibility.



Corn

A headline seat count hides that shape. This repository gives it back, provided we remember which sources of uncertainty went into drawing it.



Herman

I'd watch whether the human review becomes as reproducible as the calculations, especially when the extraction takes minutes and finding an available reviewer takes much longer.



Corn

Thanks to Hilbert Flumingtop for producing the show.



Herman

This has been My Weird Prompts, a human-AI podcast collaboration.



Corn

You can find more episodes at my weird prompts dot com.



Herman

See you tomorrow.