

The Civic Machine

Research and political reasoning



AI mediation, persuasion, agenda access and binding public decisions

Created by Yoav Duman

Purpose

The module asks students to connect a real research comparison to a fictional institutional conclusion, trace changes in the object being represented or rated, reason from separate totals for each question and defend a lawful public decision. Its unfamiliar Alder inquiry tests whether students can generate and revise an explanation independently. Treat accurate drafting, preference, agenda access, consequences and binding authority as separately assessable objects. Module authorship: Yoav Duman.

Prepare the assignment

Use the simplified explanations inside the activity during class. The primary articles and further reading are optional for students. Give students time to read, ask questions and return to a passage while answering. Keep the primary sources available for your own checks.

Read the council and Alder keys before class. Keep late records and solutions in instructor materials. Plan when students commit their own account and prediction before each reveal; avoid projecting an instructor key during that commitment.

Check Class mode and the assigned source, data, decision, transfer, workspace and justification settings. Preserve required unfamiliar transfer and method-based help. Tell students which optional work is assigned and what export Canvas will accept.

Try a saved-progress backup, restore and print/export on the devices you expect students to use. Provide a continuity route for unavailable browser storage. Confirm any hosted class arrangement separately; standalone device saving is not a class submission system.

Prepare groups to disagree about lawful options and distributional criteria. Use only the fictional cases, permit private written responses and avoid asking for family political histories. Schedule additional work for the complete inquiry and pilot actual reading, writing and export time.

A first class session with in-app reading

This suggested first class session starts with the in-app explanations. No advance reading is required. It is a planning example, not a measured completion time; use further sessions for the rest of the inquiry.

Phase	Min	Teaching move
Opening question	5	Begin from the fictional notice and save a first idea. No advance reading is assumed.

Phase	Min	Teaching move
Learn from the explanations	15	Read the two in-app explanations together. Ask students to say in their own words what was compared, what was found and one limit. Use the vocabulary help.
First evidence inquiry	15	Write an explanation and alternative, predict a new observation, then inspect the next record. Return to the explanations when needed.
Discuss the analysis	10	Compare evidence and reasons in pairs or small groups. Resolve unclear instructions before evaluating the argument.
Save and plan continuation	5	Save progress and a backup. Continue remaining data, decision, new-case and final-record tasks in further sessions.

Continue the complete inquiry

Continue the full inquiry from the saved stopping point. Read explanations during the activity whenever needed. Keep the data, decision, new-case prediction and reasoned revision, allowing time for questions and discussion. No human student completion time has been measured.

Teaching moves

Ask 'What exact response does this account explain?' Require both rival processes to address that response through an actor and operation. Two events at different stages may coexist without discriminating between causes.

Put document words, dates and the rated question beside each conclusion. Preserve the recorded useful drafting contribution while examining the later public claim; generic distrust of AI should not replace the trace.

When a student invokes a study, ask which comparison or intervention and which measured response do the work. The Tessler human comparison omitted critique; its virtual assembly did not include standard fact-finding or expert testimony. External-rater nondetection on fairness-related criteria is not equivalence.

Require a criterion and a real forgone benefit for every council recommendation. Check annual service and one-time process budgets separately, then ask which actor owes a complaint ruling and which owes a public answer.

Challenge a proposed causal test by marking the varied instruction or procedure, preassignment preference, menu, chosen consequences and response timing. Changing

the menu can change who counts as a winner; changing objective and retention together estimates a package unless components are separated.

Respond to misconceptions

Approval of a common statement establishes support for the route named later.

Have students compare the exact rated object and policy item. CM-06 supplies separate margins without individual linkage; neither margin grants the vendor or office a binding vote.

The 34-point gap tells us exactly how many approvers did not favor Green.

Ask for two different joint tables with the same margins. Their approval-only counts differ. The gap between percentages does not supply the missing overlap.

Minority text weight shows equal political influence, and an attitude shift shows a changed ballot.

Tessler's text-position estimate concerns representation in wording; Lin's measured responses concern attitudes and reported intentions. Neither establishes nomination power, material protection or observed voting.

A popular route or accurate summary settles a fair binding decision.

Separate faithful text, participation, nomination access and lawful choice. Apply the clinic rule and actual authority, and require a defended distributional criterion. Several eligible choices can receive full credit.

The later Alder instruction proves its objective alone caused the omission.

AC-05 establishes a recorded assignment and display choice. It does not isolate objective from length, format or retention instructions. Require an exact component contrast or a bounded instruction-package claim.

Debrief questions

Which exact words and dates support a useful contribution from drafting, and where does the object of the public claim change?

What do your two feasible joint tables leave unresolved, and which linked record would resolve it?

How does a tested detail from each study change your institutional conclusion? Which claim remains outside that evidence?

What lawful alternative is strongest against your council recommendation, and what benefit does your choice forgo?

Which varied feature would discriminate your causal accounts, and how could changing the menu or outcome compromise that comparison?

Which Alder conclusion did you retain or revise, what document warranted that choice, and what genuinely new uncertainty remains?

Assessment within six scoring homes

Score six distinct criteria once each, 0–3, for a maximum of 18: Source contribution; Positive finding and exact trace; For the same result rivals and causal comparison; Lawful feasible decision and accountability; Numerical and representation reasoning; Unfamiliar transfer and explicit revision. Zero is missing or materially inaccurate, one is emerging, two is developed, and three is convincing, specific and bounded. Record evidence and a reason for each score. Use the task descriptors as feedback within these homes, without additional points for reused notes or duplicate transfer reasoning. Mechanical table checks and word floors establish recorded attempts rather than semantic mastery. Policy agreement, grammar, help use, initial correctness and changing one's mind are not scoring criteria; a supported retention can earn full credit.

Source contribution

0 Missing or materially incorrect reasoning for this criterion.

1 A relevant response is present, but a key comparison, link, unit or institutional constraint remains unsupported.

2 Assess this once: does an actual tested comparison/process and measured response from each article change a particular institutional inference? Keep study populations/outcomes separate from all fictional figures; text weight is not agenda power, and attitudes are not ballots. The argument is mostly correct and bounded; an important uncertainty or feasible rival needs development.

3 Assess this once: does an actual tested comparison/process and measured response from each article change a particular institutional inference? Keep study populations/outcomes separate from all fictional figures; text weight is not agenda power, and attitudes are not ballots. The account uses specific evidence, distinguishes a credible rival or limit, and defends a feasible conclusion without overclaiming.

Positive finding and exact trace

0 Missing or materially incorrect reasoning for this criterion.

1 A relevant response is present, but a key comparison, link, unit or institutional constraint remains unsupported.

2 Assess the CM artifact argument once: cite exact words and dated actors, retain the recorded useful drafting contribution and locate the changed public object. Does the learner distinguish this established trace from a merely consistent account of why individuals endorsed? The argument is mostly correct and bounded; an important uncertainty or feasible rival needs development.

3 Assess the CM artifact argument once: cite exact words and dated actors, retain the recorded useful drafting contribution and locate the changed public object. Does the learner distinguish this established trace from a merely consistent account of why individuals endorsed? The account uses specific evidence, distinguishes a credible rival or limit, and defends a feasible conclusion without overclaiming.

Same-outcome rivals and causal comparison

0 Missing or materially incorrect reasoning for this criterion.

1 A relevant response is present, but a key comparison, link, unit or institutional constraint remains unsupported.

2 Assess the original CM rivals and council study design once: actor/operation accounts explain the same step/response and predict different evidence. Does the planned comparison actually vary a defined lawful procedure, record preference before assignment, specify menu/chosen outcome and timing, and confront consequence confounding and post-treatment winner/loser selection? Judge the AC design under transfer instead of counting it here again. The argument is mostly correct and bounded; an important uncertainty or feasible rival needs development.

3 Assess the original CM rivals and council study design once: actor/operation accounts explain the same step/response and predict different evidence. Does the planned comparison actually vary a defined lawful procedure, record preference before assignment, specify menu/chosen outcome and timing, and confront consequence confounding and post-treatment winner/loser selection? Judge the AC design under transfer instead of counting it here again. The account uses specific evidence, distinguishes a credible rival or limit, and defends a feasible conclusion without overclaiming.

Lawful feasible decision and accountability

0 Missing or materially incorrect reasoning for this criterion.

1 A relevant response is present, but a key comparison, link, unit or institutional constraint remains unsupported.

2 Assess the council recommendation once: annual and process costs, relative deadlines, clinic protection, eligible nominations/corrections/challenges, clerk ruling within three working days and council public answer before voting. Require a defended criterion, strongest lawful alternative and real forgone benefit; Green plus access, Blue or dated deferral may each earn full credit. The argument is mostly correct and bounded; an important uncertainty or feasible rival needs development.

3 Assess the council recommendation once: annual and process costs, relative deadlines, clinic protection, eligible nominations/corrections/challenges, clerk ruling within three working days and council public answer before voting. Require a defended criterion, strongest lawful alternative and real forgone benefit; Green plus access, Blue or dated deferral may each earn full credit. The account uses specific evidence, distinguishes a credible rival or limit, and defends a feasible conclusion without overclaiming.

Numerical and representation reasoning

0 Missing or materially incorrect reasoning for this criterion.

1 A relevant response is present, but a key comparison, link, unit or institutional constraint remains unsupported.

2 Assess the numerical product once: independently constructed distinct joint tables match all original margins, explanation distinguishes counts/rates/percentage-point gaps/overlap and preserves the selected base, linked hypothetical fact and later 200-person dossier. Does the library axis use each channel's full base and answer the stated comparison? Arithmetic validation alone is not interpretation mastery. The argument is mostly correct and bounded; an important uncertainty or feasible rival needs development.

3 Assess the numerical product once: independently constructed distinct joint tables match all original margins, explanation distinguishes counts/rates/percentage-point gaps/overlap and preserves the selected base, linked hypothetical fact and later 200-person dossier. Does the library axis use each channel's full base and answer the stated comparison? Arithmetic validation alone is not interpretation mastery. The account uses specific evidence, distinguishes a credible rival or limit, and defends a feasible conclusion without overclaiming.

Unfamiliar transfer and explicit revision

0 Missing or materially incorrect reasoning for this criterion.

1 A relevant response is present, but a key comparison, link, unit or institutional constraint remains unsupported.

2 Assess the AC product once: own initial actor/operation and same-outcome rival, exact omission location, bounded AC-05 update, lawful feasible board recommendation and separated representation/preference/rule-acceptance measures. Does the proposed test hold format/retention comparable, separate factors or name an instruction-package effect, while preserving nomination/challenge rights? Credit reasoned retention as well as change; require an explicit final revised inference or sharper uncertainty. The argument is mostly correct and bounded; an important uncertainty or feasible rival needs development.

3 Assess the AC product once: own initial actor/operation and same-outcome rival, exact omission location, bounded AC-05 update, lawful feasible board recommendation and separated representation/preference/rule-acceptance measures. Does the proposed test hold format/retention comparable, separate factors or name an instruction-package effect, while preserving nomination/challenge rights? Credit reasoned retention as well as change; require an explicit final revised inference or sharper uncertainty. The account uses specific evidence, distinguishes a credible rival or limit, and defends a feasible conclusion without overclaiming.

Worked data and case interpretation

Instructor reference. Keep these calculations and case interpretations separate from learner commitments until the corresponding work has been attempted. All local case data and added teaching examples are fictional; article findings retain their stated sources, units and uncertainty.

Original margins and their difference

Records

CM-06: the same original 100 responders; 80 statement approvals and 46 Green preferences.

Calculation

$80/100 \times 100 = 80\%$; $46/100 \times 100 = 46\%$; $80\% - 46\% = 34$ percentage points.

Interpretation

The difference compares marginal rates. It is not an observed count of approval-only people. 'Did not favor' includes opponents and undecided responders. No citywide probability sample or individual cross-tab is supplied.

Two valid original joint tables

Records

Original CM-06 margins only; cell order is both, approval-only, favor-only, neither.

Calculation

Pattern 1: 46, 34, 0, 20. Pattern 2: 26, 54, 20, 0. Each totals 100; the first two cells sum to 80 and the first plus third sum to 46. In general the overlap x can be any integer from 26 to 46; the remaining cells are $80-x$, $46-x$ and $x-26$.

Interpretation

These are constructed possibilities, not recovered respondent records. Any two distinct feasible overlaps are acceptable, including zero-valued cells. Their different approval-only counts demonstrate what the margins cannot identify. Keep this range in instructor materials until students have attempted the construction.

Separate linked practice fact

Records

The data-infer question adds the hypothetical fact that 11 of the 46 Green favorers did not approve. This fact does not belong to the original CM-06 observation.

Calculation

Both = $46-11 = 35$; approval-only = $80-35 = 45$; favor-only = 11; neither = $100-35-45-11 = 9$.

Interpretation

The extra linkage identifies this practice table. Students must not use it to constrain the two original-case constructions.

Later consultation and faithful texts

Records

CM-D03 and D04 concern a separate later pool of 200 self-selected participants.

Calculation

Green: $120/200 = 60\%$; Blue: $50/200 = 25\%$; undecided: $30/200 = 15\%$. Text X: $164/200 = 82\%$; text Y: $156/200 = 78\%$; endorsement difference = 4 percentage points, or 8 endorsements.

Interpretation

Keep route preferences and text endorsements separate. The counts do not identify why a formulation was preferred, a citywide preference or individual overlap between these responses.

Service eligibility and annual cost

Records

CM-D01, D02 and the exact operating/access table: €1.80 million annual ceiling and a maximum twelve-minute clinic walk.

Calculation

Green plus access: €1.62m + €0.18m = €1.80m, with an 8-minute walk. Blue: €1.74m, leaving €0.06m below the ceiling, with a 9-minute walk. Deferral: €1.20m current service and an 8-minute walk. Green alone: €1.62m and an 18-minute walk, six minutes above the maximum.

Interpretation

Green plus access, Blue and a dated deferral can each be defended with their real tradeoffs. Green alone fails the stipulated access rule even though it fits the annual budget. Cost eligibility does not establish reliability or decide the normative criterion.

Separate process allowance and timing

Records

CM-D05: €18,000 process allowance; audit/hearing €9,000 in four working days, nomination support €6,000 in three, certification €3,000 in two; parallel work permitted; meeting ten working days away.

Calculation

The complete three-activity package costs $€9,000 + €6,000 + €3,000 = €18,000$. Parallel execution can finish those activities in four working days. Adding the deferral hearing to all three would cost €28,000 and exceed the allowance. A deferral hearing plus certification costs $€10,000 + €3,000 = €13,000$, leaving €5,000.

Interpretation

The latter sum demonstrates budget arithmetic only; the whole proposed procedure still needs human review. A clerk ruling is due within three working days of a complaint, and the council owes a public answer before the vote. A late complaint may require rescheduling. A deferral must name a new decision date; its hearing draws on this same process allowance.

Independent library representation practice

Records

Rivermark Library's invented identical item: online 54 yes of 90 responders; paper 36 yes of 45 responders.

Calculation

Online: $54/90 \times 100 = 60\%$. Paper: $36/45 \times 100 = 80\%$. Paper's within-channel rate is 20 percentage points higher. Counts alone are 54 versus 36; combined-sample contributions use a different base, 135.

Interpretation

The question calls for own-channel denominators and a percent axis. Self-selected channels do not establish a causal channel effect or preferences of nonparticipants. These are not Larkhaven records or either article's data.

Review the workspace products

Construct two possible linked-response tables

Inspect all eight cells: nonnegative integer counts, total 100, approval margin 80, favor margin 46 and distinct overlaps. Zero is valid. Then assess the explanation of differing approval-only counts, the needed individual linkage and the selected base. The separate hypothetical eleven-person fact must not constrain these tables. Score this product within Numerical and representation reasoning; arithmetic validation adds no independent grade.

Your authority, cost and deadline ledger

Check each extracted actor, power, duty and dossier deadline. The council binds; residents retain eligible nomination, correction and challenge rights; the transport office reports eligibility; the vendor has no vote. Distinguish the clerk's complaint ruling within three working days from the council's public answer before voting. Check access, annual service cost, the complete process sum and late-complaint timing. Read the connected argument for a criterion, alternative and tradeoff. Completion alone does not certify lawful feasibility.

Define your Alder test after the archive

Require exact varied instructions or factors and comparable inputs, people, length/format and condition retention/display. Several simultaneous changes permit a package claim, or a design that separates components. Evaluate distinct representation, menu-preference and rule-acceptance measures with units and timing, different predictions for the same response, source contributions and protected rights. Read the preserved initial account beside the AC-05 update. Score the entire Alder product once under Unfamiliar transfer and explicit revision.

All seven evidence check keys

The same word, different jobs

Check comparison. Which account preserves the actual source comparisons and their measured objects?

Supported option B: Compare AI versus trained-participant statement ratings in Tessler with opposing assigned advocates' attitudes in Lin; each tests a different job and neither grants council authority.

The source comparisons test different assigned jobs and responses. A statement rating, attitude movement and modeled textual contribution retain distinct meanings; a charter still supplies the authority for the local vote.

Rationale to inspect: Use a specific comparison and measured response from each paper to explain your selection. Identify the link at which one competing account changes the meaning.

Read the research summaries' methods and outcomes. The paired transit illustration is invented.

What does eighty percent support?

Check scope. Which public conclusion uses both results accurately?

Supported option D: Report the two responder outcomes separately: conditional-statement endorsement and Green preference, with citywide backing unresolved.

This keeps the rated objects and selected respondent base distinct. The endorsement contributes shared criteria; neither total supplies a citywide mandate.

All counts and wording are invented teaching records; neither article reports these results.

A label and a data ban

Check transfer. Which response most directly tests that guarantee?

Supported option A: Audit assigned objective, input/output correspondence and item interpretation under the new rule; separately test what the label and profile ban change.

Disclosure and a profile ban concern specified features. The audit still needs evidence about the assigned job, represented input and measured response; it can recognize a benefit without treating the two safeguards as a neutrality test.

Rationale to inspect: Name one actual strategy comparison in Lin and one additional local record needed to assess the proposed rule. Explain why your chosen test measures the feature claimed.

Fictional policy proposal. The studies motivate inspection of the task and measured outcomes; they did not test this exact guarantee.

Locate contributions and responsibility

Check ruling. Which finding identifies the first recorded change into Green backing while preserving the other recorded contributions and limits?

Supported option C: Retain the clinic revision, locate Green backing in CM-08 and record CM-07's input boundary; separately defend the council's unmade binding choice.

The clinic condition survives into the rated text, whereas Green backing enters at publication. Missing input is a separate recorded limit; none of these records supplies the legal or political rule for the binding vote.

Rationale to inspect: Cite two artifact IDs or dates and exact words that locate the changed object. Explain what one other record establishes and what causal or political claim it leaves open.

Fixed invented challenge; do not report its counts as research findings.

Read the objects and the base

Check data-read. Which calculation and description match the display?

Supported option C: 80% approved the statement; 46% favored Green; the totals differ by 34 percentage points.

Both counts use the same base of 100. The 34-point gap compares the two totals; it does not identify exactly which people answered differently.

The display is a teaching exercise, not an estimate from either published paper.

Use a linked follow-up

Check data-infer. As a separate practice assumption, suppose a linked follow-up among the same 100 respondents found that 11 of the 46 who favored Green did not approve the statement. How many statement approvers did not favor Green?

Supported option A: 45 statement approvers did not favor Green.

Of the 46 Green favorers, 11 are outside the statement-approval group, leaving 35 inside it. Of the 80 approvers, $80 - 35 = 45$ did not favor Green. The original bars alone do not supply the extra linked fact.

Rationale to inspect: Identify what the hypothetical eleven counts, then show how the linked fact determines the overlap and requested cell. Distinguish this practice assumption from the original totals.

The display is a teaching exercise, not an estimate from either published paper.

Choose the base and axis

Check data-design. Which replacement values and axis answer that stated comparison?

Supported option D: Plot 60% and 80% on “yes responses / own-channel respondents,” retain 54/90 and 36/45, and use a 0–100% axis.

Each percentage divides yes responses by all respondents in its own channel. Exact counts and bases remain accessible; the chart compares observed rates without estimating a channel effect or citywide preference.

Rationale to inspect: Name each denominator and show why the proposed axis and values answer the intended within-channel comparison. Explain what the count-only poster answers instead.

All Rivermark figures and wording are newly invented for representation practice. They are separate from both local dossiers and the real studies.

Instructor controls and the 5907 panel

Open Instructor controls and enter 5907 for this activity. This opens the local workbench, including task keys and check answers. To manage online classes, connect the saving service and enter 5907 again in the server workbench. Canvas remains a separate login and assignment.

Set the default Class or Fun route and which routes are available. Assignment switches cover research notes, political decision, data, unfamiliar transfer, structured workspaces, written justifications, the representation check and optional frontier availability or requirement. Select full or core response guidance; choose unlimited retries or first recorded answers, and feedback after submission or on request. Method hints, unrelated worked examples, badges and deliberate note reuse can also be assigned.

Save and export the assignment settings, or import a compatible settings file and inspect the staged changes before saving. Learner preferences include reading size, high contrast, reduced motion, art visibility and guided responses. Settings, learner work and display preferences are separate records.

Review the six human criteria and task key separately from completion. Load the learner record or an explicitly submitted snapshot to record a score and evidence for each criterion; do not treat word floors, mechanical checks or supported factual choices as a semantic grade.

The private instructor recovery key is separate from 5907; use it to reopen your owned classes. Instructor sessions and class ownership remain isolated. The server workbench can save defaults, create owned classes, update assignment switches and view explicitly submitted snapshots. Private personal drafts are not silently turned into submissions. Keep local progress and settings backups as well.

Private drafts consent and submission

Online saving is optional. Connecting to a saving service does not upload or submit the device draft. Use Save online personal draft only when you want to create or update that personal copy; keep its private recovery key and your local JSON backup.

A personal online draft is not an instructor submission. To submit through a configured hosted class, find the class with its student submission code, review the destination, apply the assignment switches if requested, complete the Class record and save the current version online. Tick the explicit consent box and use Submit this response snapshot.

Submission sends that version to the identified class. Later edits stay personal until you explicitly submit again. This hosted-class route is separate from the Canvas assignment and its upload confirmation.

Device backups and Canvas

Keep compatible progress JSON and settings backups. A readable TXT or PDF export supports review; the JSON preserves continued work. Saving in the player, saving a personal online draft and explicitly submitting to a hosted class are different actions. Canvas remains a separate assignment upload; verify its requested format and submission receipt.

Research sources

Michael Henry Tessler, Michiel A. Bakker, and colleagues (2024). AI can help humans find common ground in democratic deliberation. *Science*. 386(6719): eadq2852 · DOI 10.1126/science.adq2852

[Complete primary text](#)

Hause Lin, Gabriela Czarnek, Benjamin Lewis, and colleagues (2025). Persuading voters using human-artificial intelligence dialogues. *Nature*. 648: 394–401 · DOI 10.1038/s41586-025-09771-9

[Complete primary text](#)