The Cuban Missile Crisis, 16–28 October 1962: an OptaCat reading of the documentary record

OptaCat v0.9.1 --- corpus excomm_1962 (FRUS 1961–63 vol. XI selection), run r02 with sensitivity run r03

September 2026

What OptaCat is

OptaCat turns a body of text into a directed, weighted event graph and asks nine fixed questions of its shape. It works in four frozen stages. **Step 1** splits the sources into Scenes and verbatim Lines and records every named Actor or Artifact that the text itself mentions. **Step 2** merges surface forms that refer to the same thing into canonical identities (and records the merges it refuses). **Step 3** records atomic events — *who acts on whom* on a given Line — and links an event to an earlier one with a *callback* only when the later event would not make sense without the earlier one (a counterfactual test). An event's weight is 1 plus the number of later events that call back to it. **Step 3.5** freezes an analysis policy, hashed, before any number is computed. The analytics then answer: what is the center, the bridge, the rupture, the hidden attractor, the recurring object, the deforming space, the invariant, the perturbator/dark horse, and what everything stands on (load).

OptaCat does not read meaning into the text. It measures the structure of what was recorded and how later records depend on earlier ones. That makes it most useful where its picture diverges from the accepted narrative: the divergence is either an insight or a statement about the record itself.

Key points

- **The record resolves the crisis through the letters, not through the meetings.** Four of the five diplomatic messages in the corpus sit in load tier 1. The 1,280-line evening meeting of 16 October (F021), the largest scene by far, carries zero load: nothing later in the record depends on it.
- **The documented resolution chain is the public one.** The declared outcomes (Khrushchev's 28 October dismantling announcement and Kennedy's reply) rest on exactly three scenes: Khrushchev's 26 October letter, Kennedy's 27 October reply accepting its terms, and the 28 October message. Khrushchev's 27 October Turkey-for-Cuba message is tier-1 load but lies *outside* that support.
- **The secret Jupiter assurance is invisible, and that is a finding about the record.** The only record of the Robert Kennedy–Dobrynin meeting in the corpus is RFK's memorandum of 30 October, which says "there could be no quid pro quo". The scene carries zero load. Historians have known since 1989 (Sorensen's admission) and 1995 (Dobrynin's cable) that a private assurance on the Turkish missiles was given. OptaCat reproduces the US government's own documentary version, and can locate exactly where it is silent.
- **Kennedy is the center, the bridge and the dominant invariant, but his agency coefficient is low (0.24).** He receives far more than he initiates: 155 events target him and 143 originate with him. The record casts him as the point where recommendations converge, which fits the ExComm's advisory design.
- **Khrushchev has the highest agency coefficient among principals (0.76).** In a US-held record he appears almost only as a source, as the author of letters. He is also tier-1 actor load.
- **The dark horses are the people outside the room:** the Mongoose Special Group, the congressional leaders briefed on 22 October (Halleck, Hickenlooper, Fulbright, Russell), and Gromyko. Each appears

in a few scenes but acts with concentrated positive agency. Two of these tracks are well documented escalation risks: the unauthorised CIA agent teams, and the congressional pressure for an air strike.

- **Rupture measures record-type boundaries, not historical turns.** Under the author's metric 25 of 34 transitions saturate, so Q3 is not a finding. Under node-incident churn 9 saturate, and nearly all of those are switches between document types (transcript, memo, telegram). The calmest stretch is the 23–26 October ExComm sequence: the same cast, meeting after meeting.

OptaCat versus the historical record

The historical column summarises the accepted scholarly account, drawing on FRUS 1961–63 vol. XI, May & Zelikow's *The Kennedy Tapes*, Allison & Zelikow's *Essence of Decision*, Fursenko & Naftali's *One Hell of a Gamble*, and the Dobrynin cable of 27 October 1962 (CWIHP *Bulletin* 5, 1995). "Match" means OptaCat reaches the same structural conclusion from the corpus alone. "Divergence (record)" means OptaCat differs from the historians because the corpus does not contain the evidence they used.

#	Historical finding	OptaCat result (this corpus)	Verdict
1	Kennedy was the decisive center; the ExComm advised and he decided.	Center 0.96 (next: McNamara 0.48); bridge 1.00 (next: Rusk 0.30); node invariant 0.81; actor attractor 412 (next: 119). ExComm as an actor has ASC −1.00: it only receives directives.	Match
2	McNamara shaped the quarantine option and was Kennedy's main operational interlocutor.	Center #2, actor load tier 1 (26 dependants, second only to Kennedy). The most invariant relation in the corpus is Kennedy → McNamara (0.53), then McNamara → Kennedy (0.33).	Match
3	The decision for the quarantine over an air strike was made on 20 October (505th NSC).	F034 is tier-1 load (4 dependants). Its callbacks reach into the 21 October NSC meeting (F038) and the 27 October ExComm (F090, via the NATO/Turkey study).	Match
4	The deliberations of 16 October framed the options.	F018 (morning, 16 Oct) is tier 1. F021 (evening, 1,280 lines, 89 events) has load 0: its events are consumed inside the scene.	Partial: framing is visible only for the morning meeting
5	The crisis was resolved by accepting the terms of Khrushchev's 26 October letter and setting aside the 27 October Turkey demand (the "Trollope ploy").	Outcome support = F084 (26 Oct letter) → F095 (Kennedy's 27 Oct reply) → F102 (28 Oct). F091 (the Turkey message) is tier-1 load but not in the outcome support.	Match (for the public track)
6	A secret assurance that the Jupiters in Turkey would be removed, given by RFK to Dobrynin on 27 October, was essential to the settlement.	F096 has load 0 and no path to the outcome events. Its only record in the corpus is RFK's 30 October memo denying a deal ("no quid pro quo").	Divergence (record): the corpus lacks the Dobrynin cable and RFK's handwritten notes
7	Turkey and the Jupiters were the central bargaining object of the last two days.	Turkey is artifact center #3 (after Cuba and Berlin) and recurrence #3. The Turkey trade event (F091) has weight 3 and two inbound callbacks, from F095 and F096.	Match, with weak strength: Step 3 rarely makes artifacts event endpoints
8	Khrushchev held the initiative in the endgame; his letters set the agenda.	ASC 0.76, the highest of any principal. Actor load tier 1 (12 dependants). The heaviest event in the corpus (weight 4) is his 28 October dismantling announcement.	Match
9	Congressional leaders briefed on 22 October (Russell, Fulbright) pressed for an air strike or invasion and opposed the blockade.	Halleck, Hickenlooper, Fulbright and Russell are dark-horse candidates (visibility 0.03; positive ASC). Fulbright's objections draw rebuttals from both McNamara and Kennedy.	Match

#	Historical finding	OptaCat result (this corpus)	Verdict
10	Operation Mongoose ran on during the crisis; CIA agent teams sent without clear authority risked escalation, and the 26 October Special Group meeting suspended them.	The Special Group (Augmented) is the #1 dark-horse candidate (0.16). The F082 events record the suspension and McCone/Harvey's admission of a unilateral CIA decision.	Match: surfaced unprompted
11	27 October ("Black Saturday") was the most dangerous day: the U-2 shoot-down and a looming decision to attack.	F094's loss report has weight 3 and is called back from the RFK–Dobrynin meeting. Rupture at the entry to F096_phone is 0.92. No single day stands out in load.	Partial: danger is not a structural quantity
12	Gromyko's 18 October meeting with Kennedy was a deception episode (he denied offensive weapons).	Gromyko is a dark-horse candidate (0.12; visibility 0.14; ASC 0.33). F018 → F029 is one of only 26 cross-scene callbacks.	Match (structure only)
13	Robert Kennedy was one of the two or three most influential ExComm members.	Actor load tier 1 (8 dependants), ASC 0.35, center #6, bridge #5.	Partial: influential, but less than McNamara, Rusk or Bundy in this record

The corpus and quality assurance

Corpus. The corpus covers 35 Scenes and 3,563 Lines from the FRUS 1961–63 vol. XI selection (plus one S-series memo) for 16–28 October 1962. Step 1 records 745 raw actor forms and 1,460 artifact forms; Step 2 reduces them to 162 canonical actors and 1,167 artifacts. Step 3 records 440 events and 165 callbacks (0.375 per event), with 26 callbacks crossing a scene boundary. Three scenes have no events (F021_postmeeting, F027, F031_workgroups). The corpus declares that it is an *available-document* corpus, not a complete ExComm transcript collection. In particular the 27 October afternoon and evening ExComm tapes are not in it.

Problems found in the delivered files and the correction applied (r02).

1. *Step 3 minted 11 identities.* Thirty-two events used placeholder endpoints (`act_s3_*`, `art_s3_*`) that exist in neither `plot_corpus` nor `plot_dict`. The pipeline therefore dropped them (408 of 440 events survived) and failed QA. Step 3 may not introduce entities, so the fix went back through Steps 1 and 2. Each name was added as a raw surface form taken from the Line text: Hickenlooper, Halleck, Vinson, Saltonstall, Fulbright, Acting Secretary Fowler, Special Group (NSC 5412 / Augmented), General Norstad, Ambassador Hare, "air strike alternative" / "Bundy plan", and "third draft of the President's speech". Each was added to the `dict_ids` of the Lines that contain it and given a canonical root in `plot_dict`, and the events were re-pointed. The two cascade errors (a callback to a dropped event, and two stored weights) disappeared. The full record is in `correction_log_r02.yaml`.
2. *Policy version.* The frozen policy declares `metadata.version: 0.9.0` and fails validation under 0.9.1. Following Appendix D.1, only the version field was migrated. Every coefficient, and the policy id `excomm_1962_v090_default_r01`, is unchanged.
3. *Advisory.* 345 of 880 event endpoints are not named on their own Line but are named elsewhere in the same Scene (for example, an addressee in a dialogue). Another 53, mostly the ExComm as the recipient of a record-of-action directive, are resolved from the document's context. This is acceptable for meeting records, but the Step 3 author should record it as a §6.5 endpoint exception.

QA status.

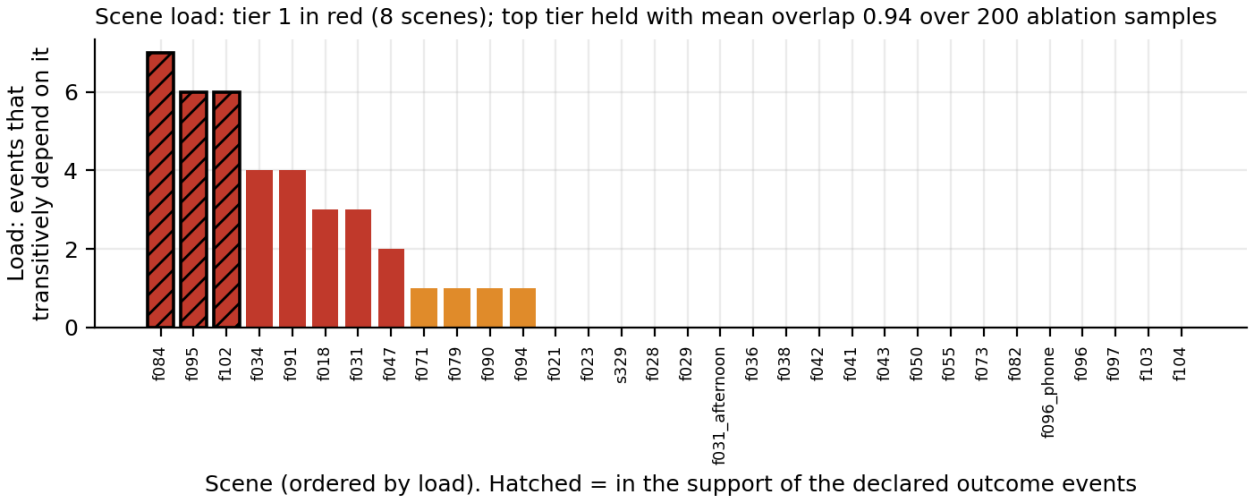
Check	r02 (primary)	r03 (sensitivity)
plot_corpus / plot_dict / plot_events	OK / OK / OK	OK / OK / OK
analysis_policy (schema, validation, consumption: 169/169 fields)	OK	OK
callback density (0.375; concentration 3 max inbound)	OK	OK

Check	r02 (primary)	r03 (sensitivity)
weighted-contact saturation (0.126)	OK	OK
load analysis (all 14 checks)	OK	OK
rupture saturation (edge)	**WARNING: 25/34 = 0.735**	OK: 9/34 = 0.265
overall	WARNING	OK

Sensitivity run r03 (`excomm_1962_v091_sensitivity_r03`) changes two fields only: `rupture.edge_churn_metric: weighted_l1_node_incident`, and `load.outcome_events: [event_f102_dismantle_order_report, event_f104_undertakings]`. It was declared *after* r02 had been read. Its rupture ranking and outcome support are therefore sensitivity results, not frozen findings. All other outputs are identical to r02.

Findings by question

Q9 — What is everything standing on? (load)

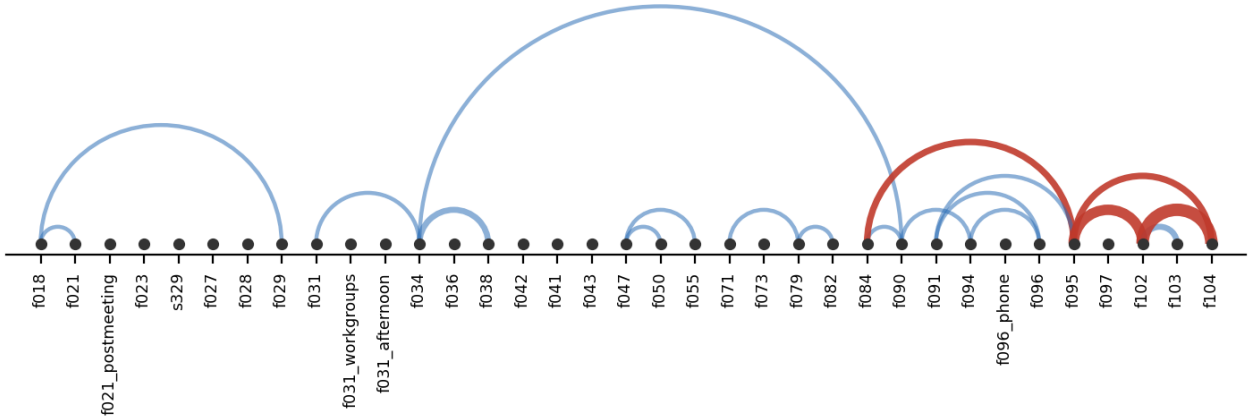


Load is the number of events that depend on a scene through a chain of callbacks: its dependants, taken transitively, with the whole group closed under union. *Tiers* come from comparing load values with strict "greater than" quantile cuts. Scenes with equal load always share a tier, and a tier is not a rank. *Tier stability* is measured by deleting 20% of callbacks at random (200 deterministic samples) and checking how much of tier 1 survives. Here the mean overlap is 0.94, the mean Jaccard 0.74, and 24% of samples reproduce the exact top tier.

Tier 1 has eight scenes: F018 (16 Oct morning), F031 (19 Oct), F034 (20 Oct, 505th NSC), F047 (23 Oct ExComm), F084 (26 Oct letter), F091 (27 Oct Turkey message), F095 (27 Oct Kennedy reply) and F102 (28 Oct). The three highest loads, 7, 6 and 6, all belong to the letter exchange. Load is unrelated to volume: the correlation with events per scene is r = −0.02, and with lines per scene r = +0.01.

Outcome support (r03) is the set of events that the declared outcome events depend on. Here it is four events, 0.9% of the corpus: F084's exchange offer, F095's acceptance of the framework and its assurances, and F102's dismantling announcement. The set is moderately stable under ablation (mean Jaccard 0.75).

All 26 cross-scene callbacks (of 165). Red = the chain that carries the declared outcome events



The arcs show every callback that crosses a scene boundary. Almost all long-range structure in this record sits in the last three days. The red chain is the documented settlement. Its path from F084 to F095 to F102/F104 does not pass through F096, the RFK–Dobrynin meeting.

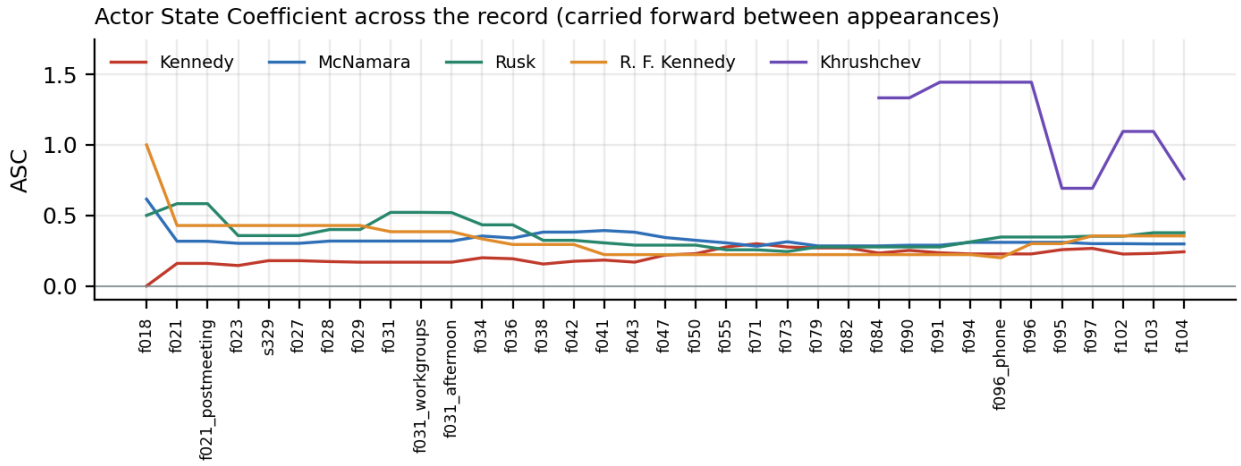
Q1, Q2 — Center and bridge

Center score combines weighted degree in the event graph with the share of scenes in which the node is present. *Bridge score* combines weighted betweenness (shortest paths weighted by the inverse of event weight) with scene diversity.

Actor	Center	Bridge	Actor load (tier)	ASC
John F. Kennedy	0.957	1.000	42 (1)	0.24
Robert S. McNamara	0.476	0.265	26 (1)	0.30
Dean Rusk	0.333	0.300	19 (1)	0.38
McGeorge Bundy	0.292	0.100	20 (1)	0.49
Khrushchev	0.274	0.056	12 (1)	0.76
Robert F. Kennedy	0.262	0.157	8 (1)	0.35
John A. McCone	0.234	0.243	5 (1)	0.21
Maxwell D. Taylor	0.224	0.127	11 (1)	0.42

Bundy and Khrushchev both have zero betweenness: every path through them has a shorter alternative, and Khrushchev's events connect him to Kennedy alone. The artifact centers are Cuba (0.25), Berlin (0.09), Turkey (0.06), the Rio Pact (0.04), Guantanamo and the UN Charter (0.03). Berlin ranks second: the record keeps returning to fear of a Soviet move on Berlin, as the historians also stress.

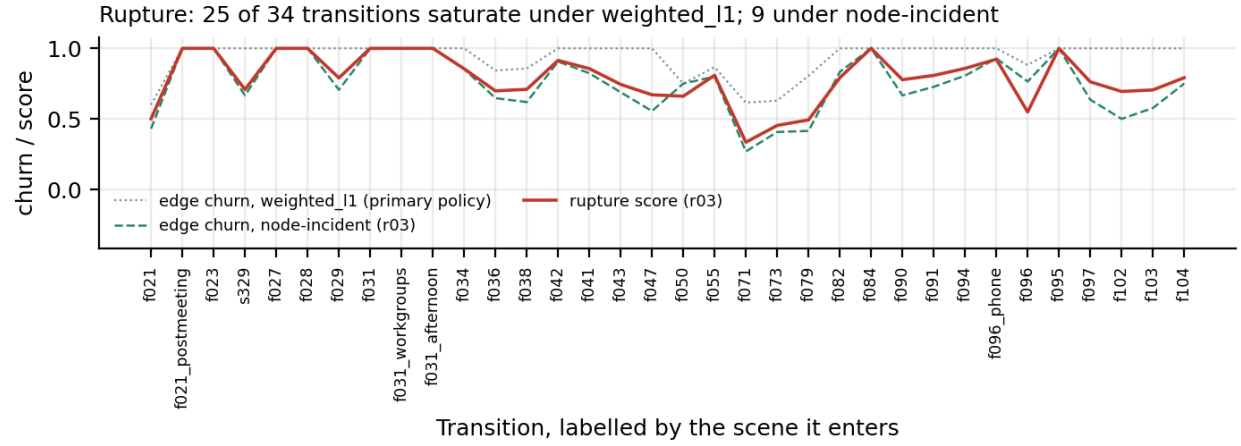
Agency — the Actor State Coefficient



The *ASC (Actor State Coefficient)* is an actor's signed agency ledger divided by the number of events it takes part in. In the primary mode (`source_weight_target_flat`) the source of an event gains the event's weight and the target loses a flat 1. A positive ASC means the actor mostly acts; a negative ASC means it is mostly acted upon. Values above 1 can occur because heavy events count more than once.

Kennedy's low ASC is structural, not a judgement. Recommendations flow to him and directives flow from him, and in the record the first outnumber the second. Khrushchev's line enters only at F084. It peaks above 1.4 while the letters dominate, drops when Kennedy's reply of 27 October addresses him, and rises again with the 28 October announcement.

Q3 — Rupture



Rupture score averages node churn and edge churn between the event windows either side of a scene boundary. *Churn* at 1.0 means the two windows share nothing. *Node-incident edge churn* compares, for each node, the total weight of its incident edges, instead of comparing exact edge keys. Under the author's metric, weighted_l1, 25 of 34 transitions sit at maximum edge churn, and the spec (§15.6) then says Q3 is not a finding. Under node-incident churn, 9 transitions reach 1.0. Almost all of these are genre boundaries: transcript to incidental recording, into and out of the parallel working groups of 19 October, from the Mongoose meeting to the Moscow telegram, and from the RFK memo to the State telegram. The lowest rupture, 0.33 at the entry to F071, begins the 23–26 October ExComm sequence, where the cast holds steady from meeting to meeting. **Reading:** in a mixed-genre documentary corpus, rupture measures changes of recording instrument rather than historical turning points. It becomes meaningful only within a single record type.

Q4, Q5 — Attractors and recurring objects

Only 2 of 440 events have an artifact as their target, so the hidden-attractor question has almost nothing to work with. Its one candidate, "third draft of the President's speech" (score 2.06), is an artifact of the annotation style, not a finding. The *actor attractors* (incoming event weight scaled by visibility) repeat the center ranking. For recurring objects (Q5, scene presence plus callback reactivation) the order is Cuba 0.41, Berlin 0.16, Turkey 0.10 and the Rio Pact 0.07. **Recommendation:** in the Step 3 revision, record interactions with the missiles, the quarantine line, the Jupiters and the letters as actor→artifact events. That would make Q4 testable.

Q6 — Space

The scenes carry dates (`kairos`) but no locus, so Q6 has one undefined locus ("?") and is **not testable** on this corpus.

Q7 — Invariants

Invariant score is the survival rate across scene transitions multiplied by normalised support (a minimum support of 3 applies). Node invariants: Kennedy 0.81, McNamara 0.38, McCone 0.23, Rusk 0.19, Taylor 0.15. Relation invariants: Kennedy → McNamara 0.53, McNamara → Kennedy 0.33, Taylor → Kennedy 0.20, RFK → Kennedy 0.13. The SVB corpus had no relational invariants. This record has one strong operational axis: the President–Secretary of Defense loop.

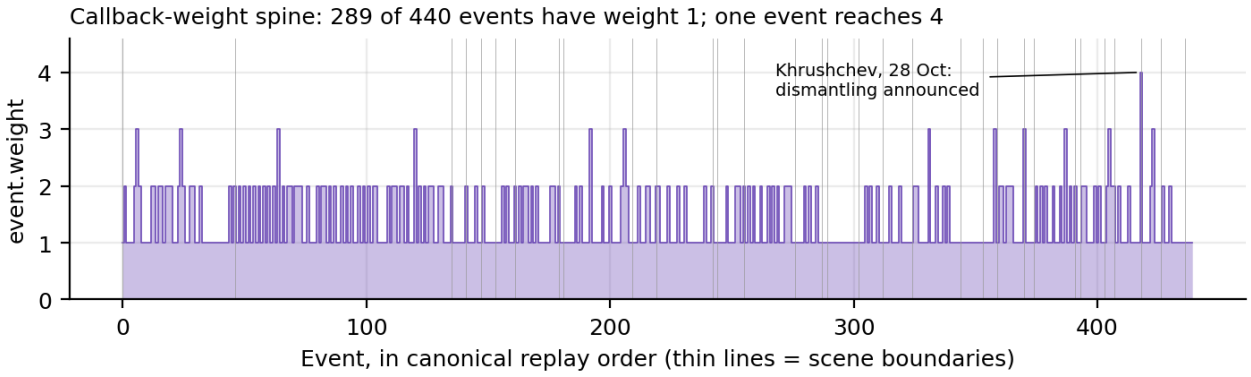
Q8 — Perturbators and dark horses

The *perturbation score* is the normalised change in the field (agency and relations) around an actor's appearances, multiplied by that actor's hiddenness. The *dark-horse score* measures concentrated positive ASC under low visibility; the visibility gate is 0.25 of scenes.

Field perturbators	score	Dark-horse candidates	score	visibility
CIA	0.440	Special Group (Augmented)	0.162	0.06
Fidel Castro	0.437	Halleck	0.144	0.03
Douglas Dillon	0.416	Hickenlooper	0.144	0.03
Maxwell D. Taylor	0.415	Edwin Martin	0.132	0.20
McGeorge Bundy	0.411	Gromyko	0.125	0.14
OAS	0.387	Cline	0.085	0.09
NATO	0.376	Fulbright	0.084	0.03
Llewellyn E. Thompson	0.372	Russell	0.084	0.03

Castro perturbs the field without taking part in a single event: his presence alone moves the actors around him. Thompson, the Soviet expert whose advice underpinned the Trollope approach, appears as a perturbator with only two events.

The callback-weight spine



The *spine* plots each event's weight in replay order. It is almost flat: 289 events have weight 1, 138 weight 2, 12 weight 3, and one weight 4 (Khrushchev's 28 October announcement, called back by the 10th ExComm and by Kennedy's reply). Weight-3 events include McNamara's timing argument and the secrecy answer of 16 October, the air-strike paper handover of 20 October, the NATO/Turkey study, the 26 October exchange offer, the 27 October Turkey trade and the U-2 loss report.

Topological questions report (r02, abridged)

Q	Question	Top result	Status
1	Center	Kennedy 0.957; McNamara 0.476; Rusk 0.333 · artifacts: Cuba 0.249, Berlin 0.094, Turkey 0.060	Finding
2	Bridge	Kennedy 1.000; Rusk 0.300; McNamara 0.265; McCone 0.243	Finding
3	Rupture	weighted_l1 saturated (0.735)	**Not a finding**; r03 reading: genre boundaries
4	Hidden attractor	”third draft of the President's speech” 2.06; structural: none	Not testable in practice (2 artifact targets)
5	Recurring object	Cuba 0.414; Berlin 0.157; Turkey 0.100	Finding (weak)
6	Deforming space	one undefined locus	Not testable
7	Invariant	Kennedy 0.808; Kennedy → McNamara 0.533	Finding
8	Perturbator / dark horse	CIA 0.440 / Special Group (Augmented) 0.162	Finding
9	Load	tier 1: F018, F031, F034, F047, F084, F091, F095, F102; actors: Kennedy, McNamara, Bundy, Rusk, Khrushchev, Taylor, RFK, McCone	Finding (stable: overlap 0.94 / 0.97)

Limits

- The corpus is a FRUS *selection*. Missing sources, above all the 27 October ExComm tapes, the Dobrynin cable and Soviet records, remove exactly the evidence behind the best-known revision of the crisis story. OptaCat's silence on the Jupiter deal is evidence about the corpus, not about history.
- Callbacks are mostly local (139 of 165 stay within their scene). Absolute load values are therefore small (at most 7), and tier 1 is distinguished from tier 2 by a few dependants.
- r03 was declared after reading r02. Its rupture and outcome results are sensitivity analyses.
- ASC and all scores are properties of the recorded structure, not measures of historical influence.

Files

`step_inputs/` holds the corrected r02 inputs and `correction_log_r02.yaml`. `run_r02/` and `run_r03/` hold the full outputs, including `topological_questions_report.md` and `qa_report.yaml`. `charts/` holds the interactive HTML charts (asc, matrix, weights, spine, load, rupture), and `fig/` the static figures, built by `make_figures.py`.