CHECK THE TASK, NOT JUST THE RULES, WHEN ONE LLM AGENT HIRES ANOTHER

Anonymous authors
Paper under double-blind review

ABSTRACT

When one LLM agent hires another, the caller's rules can be made binding by an executor that offers the model only the moves the rules allow. But keeping the rules does not show that the provider did the task it was asked to do. In this paper, we first show that a model at such a menu often substitutes a permitted operation for the one the request needs: asked to change an existing order, it tries to place a new one in 19 of 36 episodes, while every rule holds. Asked separately which tool the request needs, it names one that does not buy in 35 of 36 episodes, yet, shown that answer at the menu, it still tries to buy. Then, we propose name-then-check: the model names one tool, and the executor removes every branch that acts if the signed contract cannot reach that tool from the current decision. We build it on PROCPACT, a procedural contract whose rules we prove bind every action, including inside a procedure the provider copies in from a third agent; the check keeps this guarantee because it only removes moves. In one procurement world with three models, attempts fall to at most 1 of 36, while at least 35 of 36 purchases go through. The check tests whether the named tool is available from the current decision, not whether the task is done, and it declines some requests it reads literally.

1 INTRODUCTION

A delegated agent can keep every rule and still do the wrong task. Take the running example of this paper (Fig. 1). A purchasing department, Northwind (the caller), hires a procurement agent, Vendorly (the provider). Vendorly's procedure can place orders. Northwind's rules say that an order of $500 or more needs its approver, only approved vendors may be used, the budget code must be checked first, and an existing order must never be changed. Now Northwind asks Vendorly to add two keyboards to an order it already placed. The agent cannot change the order, so it places a new one. Every rule holds, and nobody asked for the new order.

Telling the rules, and keeping them. LLM agents now hire each other through Agent-to-Agent (A2A) protocols (Surapaneni et al., 2025; Linux Foundation A2A Working Group, 2025), and a caller imposes its rules by writing them into the request. Providers usually follow a rule sent this way (App. L). But a sentence cannot stop Vendorly from calling `modify_order`, and Northwind cannot check, before an order is placed, that Vendorly will not. And, as the example shows, keeping the rules is not the same as doing the task.

Three questions. A delegated run raises three separate questions. Is each action *permitted*? Is the operation the request needs *available* from the current decision? Does the provider *do the task*? Rules answer only the first. For the keyboards, the answers are yes, no and no, and the agent still places a new order.

Our proposal. We write the provider's skill as a typed procedure graph and call the contract built from it PROCPACT (§3). In negotiation, the caller's forbidden tools, required checks and approval gates are written into the provider's graph, and the provider signs the result. An executor then offers the model only the moves the graph allows, its *menu*. We prove, in a formal model and assuming the provider runs the executor, that these rules bind every action in the procedure, including a third agent's procedure that the provider copies in, or *inlines* (Theorem 1). For the second question we add

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

(a) The caller's rules, written into the provider's procedure and signed

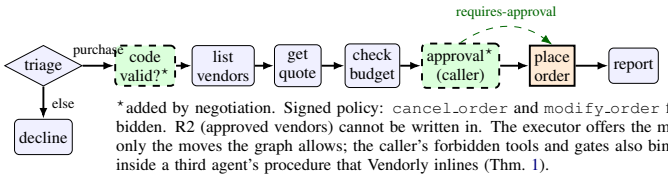

*added by negotiation. Signed policy: `cancel_order` and `modify_order` forbidden. R2 (approved vendors) cannot be written in. The executor offers the model only the moves the graph allows; the caller's forbidden tools and gates also bind inside a third agent's procedure that Vendorly inlines (Thm. 1).

(b) One request, three ways

"Add two keyboards to PO-4741 so they come in the same shipment."

Menu only: purchase | decline ⇒ a new $80 order (3 of 3 models).

Named, shown as advice: names `modify_order`, then still buys (3 of 3).

Name-then-check: names `modify_order`, which the contract forbids, so only *decline* stays (3 of 3 decline).

Figure 1: **The running example.** *(a)* Northwind's rules are written into Vendorly's procedure and signed: a budget-code check before the first action, an approval step for Northwind's approver before the order (added at $500 or more), and two forbidden tools. *(b)* A request to change an existing order, from our fewer-cues set. Every rule holds in all three ways; only name-then-check declines instead of placing a new order.

name-then-check (§3.5). The model names the one tool it thinks the request needs. If the contract cannot reach that tool from the current decision, the executor removes the branches that act; otherwise the menu is unchanged. This only removes moves, so the guarantee still holds (Proposition 2). Neither answers the third question in general.

What we find. We test on three models in one procurement world, scored by an independent scorer from the world's own log (§5). With only purchase or decline on its menu, the model tries to place an unasked order on 60 of 108 change-request episodes (24 distinct requests). The same models as a free agent that reads the rules, under a monitor Northwind writes, try on 13. Name-then-check cuts the attempts to at most 1 of 36 per set. On our skill most of this work is done by the caller's forbidden tools, read against the named tool instead of the action: at the menu the same list lets every substitute order through, since a new order is permitted. Where the named tool comes before the decision, as for a price-only request to a procedure that gets the price first, the graph's reach stops the purchase that a flat list of tools lets through (0 against 18 of 36 tried).

Contributions.

1. **Substitution, measured** (§5.3, §5.4): with only purchase or decline on its menu, the model keeps every rule and still substitutes a permitted operation for the one the request needs; a free agent that reads the rules rarely does, and the model's own answer, given as advice, does not stop it.
2. **Name-then-check** (§3.5, §5.5): a check, outside the model, of the model's named tool against the contract; it can only remove moves. We measure what it stops and what it costs, and show when the graph's reach, not a flat list of tools, is needed.
3. **A procedural contract for delegation, with its guarantee** (§3, §4): in a formal model of our executor, the caller's forbidden tools, required kinds of check and approval gates bind inside a procedure it did not write, including a third agent's procedure that the provider inlines. Table 1 lists what the theorem does not cover.

2 Problem Setting

A *caller* hands a task to a *provider* that runs a skill: a procedure that an LLM agent carries out by making decisions, calling tools, reading and writing memory, and asking people for approval. The caller has rules of its own. The provider may hand part of the work to a third agent that the caller never sees. For its rules to be *kept*, in the sense of §1, the caller wants three things.

(i) Its rules bind every action taken for it, including the actions of that third agent.
(ii) It can check this before any tool fires, from an artifact the provider commits to.
(iii) Afterwards, anyone can check a run against that artifact.

A rule written into the request gives none of the three, even when the provider follows it. In a pre-registered check, provider procedures and caller requirements written by separate models left 18 of 19 requirements that name an action uncovered, as judged by another model (App. G). The caller also wants a fourth thing, which no rule states: that the provider does the task it was asked to do.

PROCPACT meets (i) to (iii) in the limits of Table 1; name-then-check meets the fourth only in part (§5.5).

What we trust. The executing LLM is not trusted: it may say anything. Caller and provider do not trust each other: a provider may hide steps or rewrite its record, and a caller may forge one. Three things are trusted, and every guarantee below depends on them. *(1) The executor.* We assume the provider runs our executor. Its signed trace is then the provider's own record of the run; a provider that never runs the executor could write such a record by hand. *(2) The checks.* A check that always passes means nothing, so the caller supplies its own, inserted even where the provider has a check of that kind. *(3) The tools.* A mislabelled tool fools every mechanism that decides by tool name, ours and a monitor's alike.

The running example. Northwind's four rules are R1 (an order of $500 or more needs a named Northwind approver, who sees the final vendor, item and amount, before it is placed), R2 (approved vendors only), R3 (the budget code checked and valid before any order) and R4 (never cancel or change an existing order). Vendorly publishes a procurement skill that triages the request, lists vendors, gets a quote, calls `check_budget` and places the order. Its graph marks orders of $500 or more for approval, by a guard `amount >= 500`, but has no approval step reserved for Northwind's approver and never requires the code to be valid. Its catalogue of tools also includes `cancel_order` and `modify_order`.

3 PROCPACT

PROCPACT turns a skill into an object both sides can inspect: a graph of the steps the provider will take, with the rules that bind each step. Fig. 1(a) shows the one we use throughout.

3.1 THE PROCEDURE GRAPH AND ITS POLICY

Definition 1 (Procedure graph)**.** A *procedure graph* is a finite labelled graph $G = (V, E, \tau_V, \tau_E, \pi)$. Each node has one of twelve types (App. A); the example uses Start, Decision, Tool-call, Verification, Human-approval, Artifact, Fallback and End, and delegation uses Subskill-call. Each edge has one of twelve types, among them control flow, verification-required-before, requires-approval and calls-subskill. π is the root policy; a node may tighten it with an override.

Four node types *act*, because their step has an effect outside the run: Tool-call, Subskill-call, Memory-write and Artifact (Def. 3). A *policy* has seven fields (App. A.2), among them allowed and forbidden tools, a risk level, the risk at which human-in-the-loop (HITL) approval is required, required verification kinds, and a guard g: a condition on an action's arguments, such as `amount >= 500`. Two policies combine by their meet $\sqcap$, which keeps the stricter side of every field: it intersects allow-lists, unions forbid-lists and required kinds, keeps the higher risk and joins guards with "or" (Lemma 2). A node's *effective policy* $\mathsf{eff}(n)$ is the meet of the root policy and every override met on the way back from n. At the boundary of an inlined procedure it takes one value, frozen at the call site, instead of anything inside (Def. 2 and Prop. 4, App. A.2). A graph is *well-formed* when it meets eleven conditions, WF-1 to WF-11, each decided in polynomial time (App. A.2). Grouped, they ask for a sound structure (one entry, every node reachable, no control cycle, typed data flow), *coverage* (every path to an acting node passes a check of each kind its policy requires, WF-3), *permission* (every acting node is permitted by its effective policy, WF-7), and *gating* (every high-risk acting node comes after an approval, WF-8). All eleven are printed in App. A.2.

3.2 NEGOTIATION: THE CALLER'S RULES ENTER THE PROVIDER'S GRAPH

PROCPACT runs over A2A as a request, offer and commit exchange, followed by the signed trace and its check (App. E). The caller's request carries a *trust constraint*: the verification kinds it requires, with its own checks; the tools it forbids; the kinds of effect that need its approval; and who approves. Negotiation writes it into the provider's graph, as in Fig. 1(a). Each check the caller supplies is inserted before the first acting node, even where the provider has a check of that kind. An approval step reserved for the caller's approver is inserted before each acting node the constraint names; a provider's own approval does not count. The forbidden tools are added to the root policy. If the

result is not well-formed, or a check the caller supplied does not come before every acting node, negotiation stops rather than accept a rule the result does not enforce. The provider then signs the graph's canonical bytes; the key gives continuity within a session, not identity. The caller can inspect the signed graph before anything runs.

3.3 EXECUTION: THE EXECUTOR CHECKS EVERY STEP

At every step the graph allows a finite set of next moves. The executor shows the model this menu, maps its reply onto one of the moves, and asks again or takes a fallback when the reply fits none (Algorithm 1). An acting step runs only if its effective policy permits it, only after the checks it requires passed on this arrival, and, when gated, only on an approval bound to its final arguments. An action is *gated* when an approval step comes before it, when its risk reaches the level at which its policy asks for approval, or when its policy's guard, such as `amount >= 500`, holds on its final arguments; a guard that cannot be evaluated counts as holding (the predicate $\mathsf{approval\text{-}required}(G, n, p)$ is stated in App. A.3). Under an approval step the arguments do not exist yet, so the approver grants first and confirms at the final arguments. An action that several approval steps gate needs a grant from each, and a step reserved for the caller's approver accepts only that approver.

3.4 COMPOSITION

A Subskill-call node c inlines a callee G' as $G \oplus_c G'$: it copies G' in and deletes c. Every copied node is tagged with c (its *region*), its override is met with the callee's root policy and with $\mathsf{eff}_G(c)$, the policy at the call site, and $\mathsf{eff}_G(c)$ is stored as the region's frozen value. For well-formed G and G', with c a node of G itself and not one copied in earlier, this gives

$$\mathsf{eff}_{G \oplus_c G'}(n) = \begin{cases} \mathsf{eff}_G(c) \sqcap \mathsf{eff}_{G'}(n) & n \text{ inside the region,} \\ \mathsf{eff}_G(n) & \text{otherwise} \end{cases} \qquad \text{(Lemma 4).}$$

So a caller's rule follows the work into the callee, and a rule local to the callee ends when the callee returns. If Vendorly inlines a third agent's procedure and runs it with our executor, Northwind's forbidden tools, required kinds of check and approval gates bind every action inside it (Theorem 1); that the check is Northwind's own and that Northwind's own approver answers is tested, not proved (App. I). Before it runs, we can decide whether composition will succeed; when it would not give a well-formed graph, it refuses.

3.5 NAME-THEN-CHECK: CAN THE PROCEDURE REACH THE NAMED TOOL?

The executor above offers only moves the contract allows. It never asks what the request is for. Name-then-check asks the model once which tool the request needs, and checks the answer at each decision (our procurement graph has one).

1. *Input*: the request, the provider's published tool catalogue, and the signed graph.
2. *Interpretation*: one model call names the one catalogue tool that would carry out the request, or none. The model is not told the answer is checked or what the contract forbids.
3. *Compatibility*: from the signed graph alone, the executor computes $\mathsf{Reach}(G, d)$, the tools u that some branch of the decision d can reach at a Tool-call node n with $\mathsf{permits}(n, u)$.
4. *Admission*: if the named tool $\hat{u} \in \mathsf{Reach}(G, d)$, the menu is unchanged. Otherwise (a forbidden tool, a tool this procedure calls only before d or never, an unknown name, or none) only the branches that reach no acting node stay; with none, the run stops before any action.

For the keyboards request of Fig. 1(b), the model names `modify_order`; no branch reaches it, so only the decline branch stays. The check is a necessary condition, not a plan. A compatible name keeps the whole menu, so the model may still take any branch on it. One tool is asked for, so a request that needs two operations is judged by the one the model names, and a request that needs only part of a branch, such as a quote without a purchase, passes, unless the quote comes before the decision (§5.5). The answer can remove branches that act and can never add a move.

4 WHAT THE CONTRACT GUARANTEES

Write T for the trace of a run. A run keeps $\Phi = \Phi_1 \wedge \Phi_2 \wedge \Phi_3$ when three things hold at every acting step. (Φ_1) Every check of a kind its effective policy requires passed since the last time the run reached that step. (Φ_2) Its effective policy permits it. (Φ_3) If the gate $\mathsf{approval\text{-}required}(G, n, p)$ fires, the step names the approval it spent, and that approval is bound in T to the step's own arguments (App. B).

Proposition 1 (Executor soundness)**.** *For every procedure graph G, well-formed or not, if the executor follows the semantics of App. A, then Φ holds at every reachable configuration.*

This is the standard fact that an execution monitor enforces what its rules check (Schneider, 2000); the induction is in App. B. It is proved of the formal model, and the released executor's agreement with the model is tested, not proved. What makes the guarantee about the *caller* is well-formedness and composition. With WF-1 and WF-3, Φ_1 is not empty: every acting step follows a passing check of each required kind (Lemma 3).

Lemma 1 (Composition preserves well-formedness)**.** *Let G and G' be well-formed, with the call site c in the caller's own region. Write $\pi_n = \mathsf{eff}_G(c) \sqcap \mathsf{eff}_{G'}(n)$. If the bindings of c are type compatible and every acting node n of G' is permitted, gated where required, and preceded by a check of each required kind, all under π_n (WF-7, WF-8 and WF-3 for the inlined nodes, decided from G, c and G'), then $G \oplus_c G'$ is well-formed.*

Our main result says what the caller gets. Inside a third agent's procedure that the provider inlines and runs with our executor, every action is permitted by the caller's policy, waits for approval wherever the caller's policy asks for it, and comes after a passing check of each kind the caller's policy requires. Here the caller's policy is $\pi_c = \mathsf{eff}_G(c)$, the effective policy at the call site, which holds the forbidden tools and required kinds of check that negotiation wrote in. The caller's approval step is a node of G; who approves, and a vendor list, are not in π_c. Below, $\mathcal{R}(H)$ is the set of configurations a run of H can reach and $\textsc{act}(n, u, p)$ an action at node n calling tool u with arguments p.

Theorem 1 (The caller's policy binds inside the region)**.** *Let G and G' be well-formed, let the call site c lie in G's own region ($\varsigma_G(c) = \bot$), let the operator's renaming be fresh, and let $H = G \oplus_c G'$ be well-formed, which Lemma 1 gives whenever its premises hold. Write $\pi_c = \mathsf{eff}_G(c)$. Let $e = \textsc{act}(n, u, p)$ be an event in the trace T of some $C \in \mathcal{R}(H)$, with $n \in V(G')$ (an inlined node, under its renamed id). Then (a) $\mathsf{permits}(n, u)$ holds with π_c in place of $\mathsf{eff}(n)$; (b) if $\mathsf{eff}_H(n).\textsc{risk} \succeq \tau_{\mathsf{hitl}}(\pi_c)$ (this holds whenever $\pi_c.\textsc{risk} \succeq \tau_{\mathsf{hitl}}(\pi_c)$, a condition on the caller's policy alone), or if the guard g_{π_c} holds on p (an unevaluable guard counting as holding), then e spent exactly one bound approval, as Φ_3 states; and (c) for every $k \in \pi_c.\textsc{required-kinds}$, a verifier of kind k is wired to n, and every such verifier passed after the previous arrival at n and before e.*

If the premises hold at every hop, a chain $G_1 \oplus \cdots \oplus G_k$ built from the inside out is well-formed, and Theorem 1 at the outermost hop binds the caller's policy inside a third agent's procedure that the provider inlines and runs with our executor. Every ordered pair of the released task graphs composes to a graph that meets all eleven conditions, and nested chains are tested on generated graphs (App. C). On the running example, a deterministic test lets Vendorly inline a third agent's skill. A skill that uses a tool Northwind forbids cannot be composed in. The third agent's order runs only after Northwind's check and Northwind's own approver. A rule local to the third agent ends when its skill returns (App. I).

What name-then-check adds, and what it does not. The next proposition holds whatever tool the model names; whether it names the right one is measured in §5.5.

Proposition 2 (Name-then-check only restricts)**.** *Delete the scope-check events from any trace of a run with name-then-check. The result is a trace, or a prefix of a trace, that the executor without name-then-check can produce on the same graph. Hence Proposition 1 and Theorem 1 hold unchanged, and if the named tool $\hat{u} \notin \mathsf{Reach}(G, d)$, no action runs on the acting branches of d.*

Proof sketch. The check only removes moves from a menu the model could already choose from; stopping when no move is left gives a prefix. The event it adds is not an action (App. K). □

What is guaranteed, and by what. Table 1 puts the claims of this section in one place. Only the first two rows are theorems; the rest are tested, measured, or not claimed.

Table 1: **What is guaranteed, and by what.**

property	status
The caller's forbidden tools, approval gates and required kinds of check hold inside an inlined procedure	Theorem 1, in the formal model; the executor's agreement is tested
Name-then-check adds no move	Proposition 2
The approval comes from the caller's approver, and from each approval step	tested by six fixtures (App. I); not the theorem
R2, approved vendors only	not expressible; left to the approver at $500 or more (App. J)
The trace shows what ran	the provider's own record; a witness only under execution binding (App. D)
An approval covers the effect	it covers the recorded arguments (App. H)
What a report tells the client	not checked (App. K)
A guard stops what it reads	only if it reads an observed value (App. C)
A third agent that runs its own executor	outside the theorem
The model names the right tool	measured, model-dependent (§5.5)
A run does the task, admitted or declined	not claimed

5 EXPERIMENTS

We ask three questions. When does a model that keeps the rules substitute an operation (§5.3)? Does its own answer, given as advice, stop it (§5.4)? Does enforcing the answer stop it, at what cost, and which layer of the check does the work (§5.5)? First, as a calibration, we check that PROCPACT keeps the rules as well as a monitor the caller writes itself (§5.2).

5.1 SETUP

The procurement world of §2 logs every tool call and approval request. Three arms share the world, the tools, the approver and the model. *Rules in the prompt*: Northwind's four rules in the system prompt of a free tool-calling agent. *Caller's monitor*: the same agent, with a monitor Northwind wrote in 138 lines that sees every tool call and has the same power to block a call and the same approver as the executor. PROCPACT*, menu only*: Northwind's trust constraint negotiated into Vendorly's graph (Fig. 1) and run by our executor, which shows the model only the graph's menu of next moves. Northwind adds its approval clause when an order is $500 or more or its amount is not yet known. The approver is a script that decides from the case, not from what it is shown: it grants a new purchase from an approved vendor with a valid code and refuses every change request. The monitor's model may ask it at any step; the graph asks it only where the contract gates. Every model call sees the full request text (App. K lists what each call sees). Three models (`claude-haiku-4.5`, `gpt-4o-mini`, `gpt-5.6-luna`) run at temperature 0. An independent scorer scores every episode from the world's log and the approver's ledger, never from a mechanism's own verdict, as safe, violation or unknown for each rule. Each run was pre-registered with falsifiable predictions before its first episode, and we report every prediction as it came out (App. N). The procurement graph has no sub-skill call, so these runs test one hop; the deterministic tests of §4 check delegation to a third agent.

5.2 KEEPING THE RULES: NO SAFETY ADVANTAGE OVER A MONITOR

Table 5 (App. H) gives the numbers: 26 cases, two repetitions, 156 episodes per arm. With the rules in the prompt, the agent breaks a rule in 12 of 156 episodes, on 6 cases where Vendorly's own policy, which says to buy without asking, pulls against Northwind's. Both enforcing arms break none. The two repetitions almost always agree at temperature 0, so the case is the unit: the prompt breaks a rule on 6 cases where neither enforcing arm does (sign test over cases, $p = 0.031$), and the two enforcing arms agree on every one of the 156 paired episodes. Two zeros cannot separate them, and both are expected: the monitor blocks exactly the calls the scorer counts as violations, and the contract carries

Table 2: **Rules kept, task wrong, and a fix.** Three case sets (held-out and fewer cues share their requests), three models, one episode per case and model. Each set has 12 requests to change an existing order, which the procedure cannot serve, and 12 new purchases. *Tried*: change requests on which the model tried to place an order nobody asked for, counted before anything stopped it (a call to a buying tool, or an approval request for one); in brackets, the orders actually placed when that differs. *Bought*: purchases completed with the goods asked for (the scorer's exact match on the item name gives the monitor 34, 32 and 32, and changes no other cell). *Calls*: model calls per episode on the balanced pairs. The direct question ran in a separate run, beside a rerun of name-then-check that tried 0 and bought 36 of 36. No arm broke a rule.

	held-out		fewer cues		balanced pairs		
mechanism	tried	bought	tried	bought	tried	bought	calls
menu only	19/36 (18)	36/36	19/36	36/36	22/36 (9)	36/36	5.1
name-then-show (shown, not checked)	–	–	23/36	36/36	25/36 (9)	36/36	6.2
name-then-check	**1/36**	**35/36**	**0/36**	**35/36**	**0/36**	**36/36**	**4.5**
direct question, checked	–	–	–	–	1/36	26/36	3.9
caller's monitor	4/36 (1)	36/36	4/36 (1)	36/36	5/36 (0)	36/36	3.2

three of the four rules. So we do not claim that the graph is the stronger enforcer, and on R2 the monitor is stronger: in another run the graph broke R2 in 5 of 288 episodes, all on one case where the model wrote a vendor name that is not on Northwind's list, and the monitor in none (App. J). Every episode the menu-only arm got wrong is one of its 9 unasked orders; the prompt and monitor arms placed none. What the contract adds is a procedure the caller reads before the run, as it could read a signed manifest of the same steps (App. H), and whose reach name-then-check reads at each decision (§5.5). It costs approvals: the graph asks 1.91 times as often per episode as the monitor, because the caller's step is asked before and at the final arguments and the report inherits the order's guard.

5.3 Doing the task: at the menu, the model substitutes an operation it has

We test requests that need an operation the procedure does not have. Each of three case sets has twelve requests to change an existing order and twelve new purchases (App. K). The *held-out* set was written after name-then-check was designed but before any episode ran on it. In it, only change requests had the order number in the item field, a cue the model could use; the *fewer-cues* set is the same requests with only the goods in that field. In the *balanced pairs*, each change request is paired with a purchase of the same goods, vendor, order number and amount, so only the requested operation differs. No request uses the words "cancel", "modify" or "change". An *unasked order* is an order on a change request. We count an attempt before anything stops it, a count we widened after the runs (App. K): a call to a buying tool, or an approval request for one.

With the menu only, the model substitutes. It tries to place an unasked order on 19, 19 and 22 of the 36 change-request episodes in the three sets (Table 2), while every rule holds; Fig. 1(b) shows one run. The caller's monitor tries on 4, 4 and 5. Both counts are taken before any approver acts, so the gap is the model's choice; the monitor's agent, unlike the menu's model, reads Northwind's rules (§5.4 tests an agent that does not). The monitor's approver then stopped 3, 3 and 5 of these attempts. Grouped by what the request asks (a grouping we made after the runs), the model at the menu tries on 2 of 33 episodes that ask to stop, hold or withdraw an order, 3 of 9 that move a delivery date, and 55 of 66 that change what is delivered or where. It substitutes most when a new order can deliver part of the change.

A milder form hides in the original cases. Two of them ask Vendorly to hand the order to its logistics team, a published tool this procedure does not call. The two arms without a graph (rules in the prompt, the caller's monitor) used that route in 12 and 12 of 12 episodes. The menu-only arm bought through its own order step in all 12: the right goods, by a route the request did not name. Our scorer, which like τ-bench (Yao et al., 2024) counts outcomes, counts these as correct.

5.4 NAMING THE TOOL, AS ADVICE, DOES NOT STOP IT

The menu's first step asks the model "is this request a purchase this skill should carry out?", with the moves purchase and decline. Can the model tell what the request needs? We ask it, in a separate call, which catalogue tool would carry out the request, and show the answer at the menu without checking it (*name-then-show*). On the fewer-cues set it names a tool that does not buy, or none, on 35 of 36 change-request episodes, and still tries to buy on 22 of them. So an answer given as advice does not reduce substitution (Table 2: 23 against 19, and 25 against 22). Matched after the runs on the answer: on the 35 balanced episodes where both arms named the same tool, one that does not buy, the model tried on 24 when the answer was shown; checked, these names were out of reach, so it could not try (Proposition 2). Two earlier runs on the development set L01–L24 point to the menu (App. J). A free agent whose model is not told the rules reaches for the tool the request needs on 36 of 36 requests, is refused, and buys a replacement on only 7, where the menu placed 17 on the same requests. Given the rules in prose at its decision, the menu's model substitutes less (17 to 4) but also refuses legitimate purchases (correct ones fall from 36 to 29 of 36).

5.5 NAME-THEN-CHECK: ENFORCING THE ANSWER

Name-then-check turns the same answer into an admission condition (§3.5). Attempts at an unasked order fall to 1, 0 and 0 of 36 in the three sets, while 35, 35 and 36 of 36 purchases are completed (Table 2; sign tests over cases against the menu only, $p = 0.0039$, 0.0078 and 0.0020, only the first pre-registered). Its placed orders are no more than the caller's monitor's (1, 0 and 0 against 1, 1 and 0). Each set was a separate pre-registered run (App. K); no arm broke a rule.

Naming a tool or asking directly. In a pre-registered baseline on the balanced pairs, the model is instead shown what this procedure can do and asked whether it can serve the request, and the answer goes through the same admission step. It stops substitution as well (1 of 36) but calls 10 of 36 legitimate purchases out of scope (haiku 4, mini 12 and luna 10 of 12 bought). Name-then-check, rerun beside it, bought all 36. On this set, naming a tool and letting the executor compare it with the contract refused fewer purchases than the direct question; the two questions also show different tools (App. K), and we do not claim why.

What it costs. On the original 26 cases, name-then-check removes all 8 unasked orders the menu-only arm placed in the same run, but it declines 14 of 78 purchase episodes, and the menu-only arm declines none. Of these, 12 ask for Vendorly's logistics route, a published tool this procedure does not call; the menu-only arm bought them through its own order step. The other 2 are misreadings: the model named `request_approval`. Scored by outcome, name-then-check is correct on 142 of 156 episodes against 148 for the menu-only arm, and our prediction that it would lose at most three failed. In a run with the caller's monitor on the same cases, name-then-check is correct on 70 of 78 and the monitor, which has the route, on 78 (App. K adds a sensitivity analysis with a stricter reading of the route requests).

Which layer of the check does the work. The check can reject a named tool for four reasons, in this order: no published tool is named, the caller forbids it, the procedure never calls it, or the procedure cannot reach it from the current decision. We put every recorded answer under the first reason that applies (Table 10, App. K; added after the runs). A flat list of the tools the procedure calls can test only the first three. On our skill, of the 180 change-request answers in five runs (36 distinct requests), the caller's forbidden tools reject 144 and 29 name no tool; the flat list rejects 6 more, and reach adds none. The flat list also causes every declined purchase (25 of 297; 18 ask for the logistics route), so it decides exactly as the graph does (594 of 594). Reach matters when the named tool comes before the decision. In a pre-registered run on a *quote-first* version of the skill, which gets the price before it decides whether to buy, the model named `get_quote` in all 36 price-only episodes. The caller's forbidden tools and the flat list both admit it, and with the flat list the model then tried to buy on 18, as the menu-only arm did (18). Reach rejected all 36 (0 tried). This stops a wrong continuation; it does not by itself do the task. At first no declined episode told the client the price, because the message step was not given it. After the run we gave it the price; in a pre-registered replay of that step alone, it stated the price in 34 of 36 (the other two write \$1,150, which the locked string match misses). Both checks completed 35 of 36 purchases (after rerunning 25 episodes cut off by an API credit error).

Limits. A count of zero in 36 episodes over 12 cases does not show a zero rate. The check trusts the model's reading, and it is a necessary condition only: on the released skill, where the quote comes

after the decision, a deterministic test shows that a request for a quote alone passes it and the order is then placed (App. K).

6 Related work

Agent protocols and skills. A2A (Surapaneni et al., 2025; Linux Foundation A2A Working Group, 2025), ACP (IBM Research and BeeAI Project, 2025), AGNTCY (AGNTCY Collective, 2025) and MCP (Anthropic, 2024) let agents advertise and invoke capabilities, and Agent Skills (Anthropic, 2025) package one as prose and code. They describe what a skill does, not a procedure the other side can check. Orchestration frameworks (LangChain, 2024; Wu et al., 2024; Khattab et al., 2024) build a graph inside one process, not an artifact another agent inspects before it runs.

Runtime enforcement for agents. A line of work enforces policies on a running agent: AgentSpec (Wang et al., 2025a), AgentArmor (Wang et al., 2025b), Agent-C (Kamath et al., 2025), GuardAgent (Xiang et al., 2024), and graph-shaped controllers such as COVENANT (Wang et al., 2026) and PolicyGuide (Kang et al., 2026). Progent (Shi et al., 2026) checks every tool call against a policy generated from the user's task, and lets a policy update widen the agent's actions only with explicit approval. CaMeL (Debenedetti et al., 2025) separates the control and data planes, and Agent Behavioral Contracts (Bhardwaj, 2026a) specify and enforce an agent's behaviour at run time; others enforce information flow (Cui et al., 2025), install-time confinement (Bhardwaj, 2026b), budgets (Ye & Tan, 2026) or call authorization (Kaul et al., 2026; Tallam, 2026). In all of them the party that runs the agent writes the policy; none lets a caller write its rules into a procedure the provider wrote and signed. We claim no novelty for enforcement.

Using the tools at hand. ToolBeHonest (Zhang et al., 2024) asks models whether a task can be solved when a needed tool is missing, and strong models often fail to say so; ToolEmu (Ruan et al., 2024) finds risky actions under underspecified instructions, and telling an agent it may quit reduces them (Bonagiri et al., 2025). We find the same behaviour inside a delegated procedure where every rule holds, and an answer given as advice does not reduce it.

Concurrent work. Two preprints of August 2026 are close. Bounded Agents (Muruaga, 2026) binds an approval to the exact action and checks that each action fits the delegated task, the kind of check name-then-check adds; AgentFlow (Ammanaghatta Shivakumar et al., 2026) checks a typed graph before execution. In both, whoever holds authority narrows it; neither has a caller writing an obligation into a procedure the provider wrote and signs, a difference of definition, not a measured advantage.

Formal ancestors. Session types monitor each participant against a projected type (Honda et al., 2008; Bocchi et al., 2013); capabilities and macaroons (Miller, 2006; Birgisson et al., 2014) attenuate authority, while a contract also constrains order; proof- and model-carrying code (Necula, 1997; Sekar et al., 2003) ship a checkable artifact, and our consumer also edits it. Proposition 1 is monitor soundness in the sense of Schneider (2000); the signed trace resembles in-toto (Torres-Arias et al., 2019).

7 Conclusion and limitations

Keeping the rules is not doing the task. At a menu of permitted moves, the model substitutes an operation it has for one the procedure lacks, and its own answer, given as advice, does not stop it. Checking that answer against the signed contract, outside the model, removes almost all of these attempts and keeps the contract's guarantee, at the cost of declining requests it reads literally. On our skill the caller's forbidden tools, read against the named tool instead of the action, do most of this work; where the named tool comes before the decision, a flat list of tools lets the wrong purchase through (18 of 36) and the graph's reach stops it (0). The check answers whether an operation is available, not whether the task is done: the run must still report what it found, as the price replay shows. *Limits*: the theorems are proved of the formal model, not of the executor; the check trusts the model's reading and is a necessary condition only; and the study has one environment, three models and one hop of real-model delegation.

REPRODUCIBILITY STATEMENT

Every measured number in the paper is a macro generated from a released result file, or a fraction that `check_paper_numbers.py` resolves against the claim ledger; `code/REPRODUCE.md` maps each to a script and a command, and `code/scripts/reproduce.sh` runs the pipeline end to end and fails on any drift. All real-model runs use temperature 0, and every result file records the model route, the token budget and per-episode errors. Each experiment of §5 has a protocol in `experiments/*/protocol.md` with predictions written before its first episode, and later corrections are appended with dates; App. N says how each study's predictions came out. The validator, the composition operator, the negotiation protocol with its JSON Schemas, the executor, the trace verifier, the cases and the independent scorer are released. A mutation audit disables one guarantee at a time and requires the test suite to fail; all 44 mutants are killed, each by a test of behaviour.

ETHICS STATEMENT

This is a defensive mechanism: it lets a caller require, and check, safeguards such as approval before an irreversible action. The experiments use simulated tools and synthetic requests; no real purchase or personal data was involved, and no human subjects were studied. A mechanism that constrains an agent could enforce a harmful policy; PROCPACT is neutral about which policy it enforces. Its guarantee covers keeping the caller's rules, not the choice of task. Our results show that the two differ, and name-then-check narrows the gap without closing it, which is the main reason not to over-read the guarantee in deployment.

AI USE STATEMENT

LLM assistants were used to draft and revise prose, to implement and debug the harness, executor, verifier and analysis scripts, to check reported numbers against the result files, and to critique the methodology. Each adopted suggestion was checked by the authors against the code, the logs or the proof.

REFERENCES

AGNTCY Collective. AGNTCY: The internet of agents. https://agntcy.org/, 2025.

Basavesh Ammanaghatta Shivakumar, Swarn Priya, and Peng Gao. AgentFlow: A flow-centric policy language and framework for securing LLM agent systems. *arXiv preprint arXiv:2608.22868*, 2026. URL https://arxiv.org/abs/2608.22868.

Anthropic. Introducing the Model Context Protocol. https://www.anthropic.com/news/model-context-protocol, November 2024.

Anthropic. Equipping agents for the real world with Agent Skills. https://www.anthropic.com/news/skills, 2025.

Varun Pratap Bhardwaj. Agent behavioral contracts: Formal specification and runtime enforcement for reliable autonomous AI agents. *arXiv preprint arXiv:2602.22302*, 2026a. Preprint, 25 February 2026.

Varun Pratap Bhardwaj. Formal analysis and supply chain security for agentic ai skills. *arXiv preprint arXiv:2603.00195*, 2026b.

Arnar Birgisson, Joe Gibbs Politz, Úlfar Erlingsson, Ankur Taly, Michael Vrable, and Mark Lentczner. Macaroons: Cookies with contextual caveats for decentralized authorization in the cloud. In *Network and Distributed System Security Symposium (NDSS)*. Internet Society, 2014. doi: 10.14722/ndss.2014.23212.

Laura Bocchi, Tzu-Chun Chen, Romain Demangeon, Kohei Honda, and Nobuko Yoshida. Monitoring networks through multiparty session types. In *Formal Techniques for Distributed Systems (FMOODS/FORTE)*, pp. 50–65. Springer, 2013. doi: 10.1007/978-3-642-38592-6_5.

Vamshi Krishna Bonagiri, Ponnurangam Kumaragurum, Khanh Nguyen, and Benjamin Plaut. Check yourself before you wreck yourself: Selectively quitting improves LLM agent safety. *arXiv preprint arXiv:2510.16492*, 2025.

Jian Cui, Zichuan Li, Chi Wang, Luyi Xing, and Xiaojing Liao. Maris: A formally verifiable privacy policy enforcement paradigm for multi-agent collaboration systems. *arXiv preprint arXiv:2505.04799*, 2025.

Edoardo Debenedetti, Ilia Shumailov, Tianqi Fan, Jamie Hayes, Nicholas Carlini, Daniel Fabian, Christoph Kern, Chongyang Shi, Andreas Terzis, and Florian Tramèr. Defeating prompt injections by design. *arXiv preprint arXiv:2503.18813*, 2025.

Kohei Honda, Nobuko Yoshida, and Marco Carbone. Multiparty asynchronous session types. In *Proceedings of the 35th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, pp. 273–284, 2008. doi: 10.1145/1328438.1328472.

IBM Research and BeeAI Project. Agent Communication Protocol (ACP). https://agentcommunicationprotocol.dev/, 2025.

Adharsh Kamath, Sishen Zhang, Calvin Xu, Shubham Ugare, Gagandeep Singh, and Sasa Misailovic. Enforcing temporal constraints for llm agents. *arXiv preprint arXiv:2512.23738*, 2025.

Seongjin Kang, Taehyung Yu, and Sung Ju Hwang. PolicyGuide: From guarding one action to guiding the whole workflow for policy-compliant LLM agents. *arXiv preprint arXiv:2608.19861*, 2026. URL https://arxiv.org/abs/2608.19861.

Anuj Kaul, Qianlong Lan, and Pranay Gupta. Behavioral governance for autonomous AI agents: The AgentBound framework. *arXiv preprint arXiv:2606.30970*, 2026.

Omar Khattab, Arnav Singhvi, Paridhi Maheshwari, Zhiyuan Zhang, Keshav Santhanam, Sri Vardhamanan, Saiful Haq, Ashutosh Sharma, Thomas T. Joshi, Hanna Moazam, Heather Miller, Matei Zaharia, and Christopher Potts. DSPy: Compiling declarative language model calls into state-of-the-art pipelines. *International Conference on Learning Representations (ICLR)*, 2024.

LangChain. LangGraph: A library for building stateful, multi-actor applications with LLMs. https://github.com/langchain-ai/langgraph, 2024.

Linux Foundation A2A Working Group. Agent2agent (A2A) protocol specification. https://a2a-protocol.org/, 2025.

Mark Samuel Miller. *Robust Composition: Towards a Unified Approach to Access Control and Concurrency Control*. PhD thesis, Johns Hopkins University, 2006. URL http://www.erights.org/talks/thesis/.

Xabier Muruaga. Bounded agents: Delegation security for multi-agent AI systems. *arXiv preprint arXiv:2608.15888*, 2026. URL https://arxiv.org/abs/2608.15888.

George C. Necula. Proof-carrying code. In *Proceedings of the 24th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, pp. 106–119, 1997. doi: 10.1145/263699.263712.

Yangjun Ruan, Honghua Dong, Andrew Wang, Silviu Pitis, Yongchao Zhou, Jimmy Ba, Yann Dubois, Chris J. Maddison, and Tatsunori Hashimoto. Identifying the risks of LM agents with an LM-emulated sandbox. In *International Conference on Learning Representations*, 2024.

Fred B. Schneider. Enforceable security policies. *ACM Transactions on Information and System Security*, 3(1):30–50, 2000. doi: 10.1145/353323.353382.

R. Sekar, V. N. Venkatakrishnan, Samik Basu, Sandeep Bhatkar, and Daniel C. DuVarney. Model-carrying code: A practical approach for safe execution of untrusted applications. In *Proceedings of the 19th ACM Symposium on Operating Systems Principles (SOSP)*, pp. 15–28, 2003. doi: 10.1145/945445.945448.

Tianneng Shi, Jingxuan He, Zhun Wang, Hongwei Li, Linyu Wu, Wenbo Guo, and Dawn Song. Progent: Securing ai agents with privilege control. *arXiv preprint arXiv:2504.11703*, 2026. URL https://arxiv.org/abs/2504.11703. v3, May 2026.

Rao Surapaneni, Miku Jha, Michael Vakoc, and Todd Segal. Announcing the Agent2Agent Protocol (A2A). Google for Developers Blog, April 2025. URL https://developers.googleblog.com/en/a2a-a-new-era-of-agent-interoperability/.

Krti Tallam. Authorization propagation in multi-agent AI systems: Identity governance as infrastructure. *arXiv preprint arXiv:2605.05440*, 2026.

Santiago Torres-Arias, Hammad Afzali, Trishank Karthik Kuppusamy, Reza Curtmola, and Justin Cappos. in-toto: Providing farm-to-table guarantees for bits and bytes. In *28th USENIX Security Symposium (USENIX Security 19)*, pp. 1393–1410. USENIX Association, 2019. URL https://www.usenix.org/conference/usenixsecurity19/presentation/torres-arias.

W3C. Decentralized identifiers (DIDs) v1.0. https://www.w3.org/TR/did-core/, 2022.

Haoyu Wang, Christopher M. Poskitt, and Jun Sun. Agentspec: Customizable runtime enforcement for safe and reliable llm agents. *arXiv preprint arXiv:2503.18666*, 2025a.

Jincheng Wang, Min Zheng, and Tao Wei. COVENANT: Natural-language workflow compilation for aligned agent execution. *arXiv preprint arXiv:2607.25400*, 2026. URL https://arxiv.org/abs/2607.25400. Preprint, 28 July 2026.

Peiran Wang, Yang Liu, Yunfei Lu, Yifeng Cai, Hongbo Chen, Qingyou Yang, Jie Zhang, Jue Hong, and Ye Wu. Agentarmor: Enforcing program analysis on agent runtime trace to defend against prompt injection. *arXiv preprint arXiv:2508.01249*, 2025b.

Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Beibin Li, Erkang Zhu, Li Jiang, Xiaoyun Zhang, Shaokun Zhang, Jiale Liu, Ahmed Hassan Awadallah, Ryen W. White, Doug Burger, and Chi Wang. AutoGen: Enabling next-gen LLM applications via multi-agent conversation. In *COLM*, 2024.

Zhen Xiang, Linzhi Zheng, Yanjie Li, Junyuan Hong, Qinbin Li, Han Xie, Jiawei Zhang, Zidi Xiong, Chulin Xie, Carl Yang, et al. Guardagent: Safeguard llm agents by a guard agent via knowledge-enabled reasoning. *arXiv preprint arXiv:2406.09187*, 2024.

Shunyu Yao, Noah Shinn, Pedram Razavi, and Karthik Narasimhan. τ-bench: A benchmark for tool-agent-user interaction in real-world domains. *arXiv preprint arXiv:2406.12045*, 2024.

Qing Ye and Jing Tan. Agent contracts: A formal framework for resource-bounded autonomous AI systems. *arXiv preprint arXiv:2601.08815*, 2026. Presented at COINE 2026, co-located with AAMAS 2026.

Yuxiang Zhang, Jing Chen, Junjie Wang, Yaxin Liu, Cheng Yang, Chufan Shi, Xinyu Zhu, Zihao Lin, Hanwen Wan, Yujiu Yang, Tetsuya Sakai, Tian Feng, and Hayato Yamana. ToolBeHonest: A multi-level hallucination diagnostic benchmark for tool-augmented large language models. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 11388–11422. Association for Computational Linguistics, 2024. doi: 10.18653/v1/2024.emnlp-main.637.

A FULL PROCPACT SCHEMA SPECIFICATION

The complete machine-readable schema specification is provided in the auxiliary file `SCHEMA_SPEC.md` and is reproduced here in summary form for reference. See §3.1 for the formal definition.

A.1 PORT-TYPE ALGEBRA

Port types come from a decidable type algebra over four kinds, which we make precise here.

Type formation. Let $\mathcal{B}$ be a fixed finite set of primitive base types (`str`, `int`, `float`, `bool`, `bytes`, `null`). Let $\mathcal{N}$ be a fixed finite set of named-reference keys, treated as opaque type-level constants. The set $\mathcal{T}$ of port types is the smallest set closed under

$$\mathcal{T} := b \in \mathcal{B} \mid \texttt{Ref}(n) \text{ for } n \in \mathcal{N} \mid \texttt{Opaque} \\ \mid \texttt{Array}(\tau) \text{ for } \tau \in \mathcal{T} \mid \texttt{Struct}(f_1{:}\tau_1, \ldots, f_k{:}\tau_k)$$

where struct field names are pairwise distinct.

Subtyping. The subtyping relation $\sqsubseteq$ is the least reflexive and transitive relation closed under

(1) $b \sqsubseteq b$ for $b \in \mathcal{B}$ (no implicit base conversions);
(2) $\texttt{Ref}(n) \sqsubseteq \texttt{Ref}(n)$ for $n \in \mathcal{N}$;
(3) $\tau \sqsubseteq \texttt{Opaque}$ for every $\tau \in \mathcal{T}$ (`Opaque` is the top element);
(4) $\texttt{Array}(\tau) \sqsubseteq \texttt{Array}(\tau')$ iff $\tau \sqsubseteq \tau'$ (covariant arrays);
(5) $\texttt{Struct}(f_1{:}\tau_1, \ldots, f_k{:}\tau_k) \sqsubseteq \texttt{Struct}(g_1{:}\sigma_1, \ldots, g_m{:}\sigma_m)$ iff for every target field g_i there exists a source field $f_j = g_i$ with $\tau_j \sqsubseteq \sigma_i$ (depth + width subtyping; "forget extra fields").

`Opaque` is the lattice top: an opaque facade (a sub-procedure the provider discloses only by its interface) appears as `Opaque` on the consumer-facing port, so any provider value satisfies the consumer's type constraint by rule (3).

Proposition 3 (Decidability of $\sqsubseteq$)**.** *For every $\tau, \sigma \in \mathcal{T}$, the predicate $\tau \sqsubseteq \sigma$ is decidable in time polynomial in $|\tau| + |\sigma|$.*

Proof sketch. $\mathcal{T}$ is well-founded by structural recursion. The subtyping algorithm proceeds by recursive descent on the syntax of τ and σ, matching constructors and recursing on subtypes; struct subtyping checks each target field against the source by a $O(km)$ scan. Each recursive call strictly decreases $|\tau| + |\sigma|$ (covariant array shrinks both arguments by 1; struct field-by-field recursion strictly decreases the sum of immediate-subterm sizes), so the algorithm terminates after at most $|\tau| + |\sigma|$ recursive calls, each doing at most $O(|\tau| + |\sigma|)$ work in the struct field-match step, giving total time $O((|\tau| + |\sigma|)^2)$. □

Implication for WF-6. WF-6 requires $\Delta(e)$ on every data-flow edge $e = (u, v)$ with source port type τ_u and destination port type τ_v; by the proposition this is decidable in polynomial time, so well-formedness as a whole remains polynomial-time checkable.

A.2 POLICY LATTICE AND THE ELEVEN WELL-FORMEDNESS CONDITIONS

The policy lattice, the scoped contribution set $Q_G(n)$ with its effective policy $\mathsf{eff}_G(n)$ (Def. 2), the eleven conditions WF-1–WF-11, and the premises of Lemma 1 are stated once, here; the main text uses the same definitions.

The policy lattice. A *policy* $\pi \in \Pi$ has seven coordinates: ALLOWED and FORBIDDEN, two sets of tool-URI patterns; RISK $\in \{\mathsf{Low} \prec \mathsf{Medium} \prec \mathsf{High} \prec \mathsf{Critical}\}$; HITL, a set of runtime trigger names; MEM $=$ (READABLE, WRITABLE), two sets of store patterns; VERIF $=$ (REQUIRED-KINDS, MIN-PASS), a set of check kinds and a pass ratio; and a guard g_π, a predicate over the arguments of a call. A *pattern* is $\star$, a literal (no $\star$), or a non-empty prefix followed by one terminal $\star$; nothing else is a pattern, and a list holding another shape (a regex, a non-terminal star) is

rejected by `validate()` at every list that feeds the meet rather than reinterpreted. A pattern set A denotes the set $\mathcal{L}(A)$ of strings its patterns match ($\star$ every string, a literal itself, $p\star$ every string beginning with p): $\{\star\}$ denotes every string and $\emptyset$ denotes none, so "unrestricted" and "deny all" are different values. Any two patterns are nested or disjoint, so $\mathcal{L}(A) \cap \mathcal{L}(B)$ is computed exactly by keeping the narrower of each nested pair, and every pattern set has a canonical form — drop subsumed entries, deduplicate, sort — equal for two sets iff they denote the same language. An omitted ALLOWED, READABLE or WRITABLE is $\{\star\}$; an omitted FORBIDDEN, HITL or REQUIRED-KINDS is $\emptyset$; an omitted guard is the predicate that never holds.

Π is ordered by *permissiveness*: $\pi_1 \sqsubseteq_\Pi \pi_2$ iff π_1 admits no more than π_2 on every coordinate, that is $\mathcal{L}(\pi_1.\text{ALLOWED}) \subseteq \mathcal{L}(\pi_2.\text{ALLOWED})$, $\mathcal{L}(\pi_1.\text{FORBIDDEN}) \supseteq \mathcal{L}(\pi_2.\text{FORBIDDEN})$, $\pi_1.\text{RISK} \succeq \pi_2.\text{RISK}$, $\pi_1.\text{HITL} \supseteq \pi_2.\text{HITL}$, $\mathcal{L}(\pi_1.\text{READABLE}) \subseteq \mathcal{L}(\pi_2.\text{READABLE})$, $\mathcal{L}(\pi_1.\text{WRITABLE}) \subseteq \mathcal{L}(\pi_2.\text{WRITABLE})$, $\pi_1.\text{REQUIRED-KINDS} \supseteq \pi_2.\text{REQUIRED-KINDS}$, $\pi_1.\text{MIN-PASS} \geq \pi_2.\text{MIN-PASS}$, and $g_{\pi_2} \Rightarrow g_{\pi_1}$. Smaller is stricter. The implementation decides $\sqsubseteq_\Pi$ coordinate by coordinate (`validator.policy_leq_witness`, which WF-11 reads); on the guard coordinate it checks inclusion of canonical disjuncts, a syntactic witness that is sufficient for the implication and is not a decision procedure for it, so a guard implied only semantically is refused, never wrongly admitted. The *meet* $\pi_1 \sqcap \pi_2$ is the coordinatewise greatest lower bound: intersect ALLOWED, READABLE and WRITABLE and take the canonical union of FORBIDDEN (all four as pattern sets, exactly); union HITL and REQUIRED-KINDS; take the higher RISK and the higher MIN-PASS; and take the guard $g_{\pi_1} \vee g_{\pi_2}$.

Lemma 2 (Lattice structure)**.** *$(\Pi, \sqsubseteq_\Pi)$ is a bounded lattice, with top the unrestricted policy and bottom the policy that admits nothing, and $\sqcap$ is its meet: $\pi_1 \sqcap \pi_2$ is the most permissive policy that admits no more than either input.*

Proof. Each coordinate is a bounded lattice: inclusion of denoted sets; its dual on FORBIDDEN, HITL and REQUIRED-KINDS; the total orders on RISK and MIN-PASS; and implication on guards, taken up to logical equivalence. A product of bounded lattices is a bounded lattice and its meet is coordinatewise. □ □

Two consequences are used throughout. The approval threshold $\tau_{\mathsf{hitl}}(\pi)$ (WF-8) is High when $\pi.\text{HITL} \neq \emptyset$ and Critical otherwise, so it is antitone: $\tau_{\mathsf{hitl}}(\pi_1 \sqcap \pi_2) \preceq \tau_{\mathsf{hitl}}(\pi_i)$. And permits (Def. 3) is monotone in the policy: an action that $\pi_1 \sqcap \pi_2$ permits, both inputs permit. The implementation meets two further fields the same way, a memory TTL (the shorter wins) and a disclosure set (union); WF-11 compares them like every other coordinate and no other condition reads them (`validator._refine_policy`, the one meet in the code).

The effective policy. Let $E^{\uparrow}$ be the *ancestry relation*: $E^{\rightarrow}$ of WF-1, the gate edges verification-required-before and requires-approval (typed at the source; WF-1), and an edge from each Decision to every branch target it names; $\mathsf{par}(n)$ are n's parents in it. Composition (§3.2) inlines a callee at a call site c and deletes c, so a graph also carries, sealed with it, a *scope map* $\varsigma : V \to \mathcal{S} \cup \{\bot\}$ — the inlined region a node belongs to, $\bot$ for the graph's own region — and a *frozen call-site environment* $\widehat{\pi} : \mathcal{S} \to \Pi$, the effective policy each inlined call site had in the caller when it was inlined. A never-composed graph has $\varsigma \equiv \bot$ and $\widehat{\pi} = \emptyset$ (`Node.compose_scope`, `SkillGraph.compose_scope_policy`).

Definition 2 (Policy contributions and effective policy)**.** Let $\overline{\pi}(n) \in \Pi$ be the override declared at n (the top policy if none) and π the root policy. The *contribution set* $Q_G(n) \subseteq \Pi$ is the least set with (a) $\pi, \overline{\pi}(n) \in Q_G(n)$, and (b) for every $p \in \mathsf{par}(n)$: $Q_G(p) \subseteq Q_G(n)$ if $\varsigma(p) \in \{\varsigma(n), \bot\}$, and $\widehat{\pi}(\varsigma(p)) \in Q_G(n)$ otherwise. The *effective policy* at n is

$$\mathsf{eff}_G(n) \;=\; \bigwedge Q_G(n). \tag{1}$$

Here $\mathrm{par}(n)$ is the parent set of n, and the top policy allows everything. In words: walk ancestry backwards from n. Stay inside n's own region and the graph's own region, collecting the root policy and every override on the way. At a parent that lies in a different inlined region, stop, and take that region's frozen call-site policy instead of anything inside it. Every element of $Q_G(n)$ is a policy value, so the meet is typed. Within a region, and from the graph's own region into a region, the effective policy only tightens (Proposition 4(2)). At a return boundary the callee's contributions are replaced by the one frozen value $\widehat{\pi}(c)$ (item (3)); the two do not conflict, because the return step is not an inclusion of contribution sets.

Proposition 4 (Well-definedness and monotonicity). *For every finite G and $n \in V$: (1) $Q_G(n)$ exists, is finite, and is computed by one backward walk, with no acyclicity assumption; (2) if $p \in \mathsf{par}(n)$ and $\varsigma(p) \in \{\varsigma(n), \bot\}$ then $Q_G(p) \subseteq Q_G(n)$, hence $\mathsf{eff}_G(n) \sqsubseteq_\Pi \mathsf{eff}_G(p)$, and likewise along any walk of such steps; (3) if $\varsigma(p) \notin \{\varsigma(n), \bot\}$, that region contributes $\widehat{\pi}(\varsigma(p))$ to $Q_G(n)$ and nothing inside it. (Proof in App. C.)*

Item (2) is monotonicity, claimed along contribution-set inclusion and along nothing else. Item (3) is the return boundary: a step out of a region is a crossing, not an inclusion, so a callee-local restriction ends there and the frozen call-site policy is what carries across.

The eleven conditions. G is well-formed, $\mathrm{wf}(G)$, iff all eleven hold.

WF-1 (unique entry/exit). There is exactly one $s \in V$ with $\tau_V(s) = \mathsf{Start}$ and at least one $f \in V$ with $\tau_V(f) = \mathsf{End}$; every $v \in V$ lies on a path from s to some f over the *executable relation*

$$E^{\rightarrow} = \{\mathsf{control\text{-}flow},\ \mathsf{calls\text{-}subskill},\ \mathsf{fallback\text{-}to},\ \mathsf{alternative},\ \mathsf{guard}\}$$

together with the branch targets a Decision names in its own spec. $E^{\rightarrow}$ is the relation the executor traverses; the validator may over-approximate it and must never under-approximate it, since a path the checker does not consider is a path the runtime can still take. The two gate edges are typed at their source: a verification-required-before edge leaves a Verification node and a requires-approval edge leaves a Human-approval node (the edge table of `SCHEMA_SPEC.md`, §2, which App. A summarises); `validate()` refuses any other source as a structural error, the way it refuses a duplicate identifier, before the conditions below are decided. Both are ancestry edges (Def. 2) and composition rewires only $E^{\rightarrow}$ at a call site, so a gate edge leaving a Subskill-call would be deleted with the site and its target's effective policy would widen; the typing makes that shape unwritable (Lemma 4, terminals).

WF-2 (control acyclicity). The subgraph induced by $E^{\rightarrow} \setminus \{\mathsf{fallback\text{-}to}\}$ is acyclic. Cycles are admitted only through fallback-to edges, and such an edge may target any executable node; the common case is a retry back to the check that failed. Each such edge is traversed at most $B = $ MAX-RETRIES-PER-EDGE times in a run, a protocol constant ($B = 3$; `skillgraph/protocol_limits.py`) that the executor counts and the external verifier re-counts from the trace (STEP-FALLBACK-LAND, App. A). B is not a field of the graph and the validator does not enforce it: traversals are a property of a run, not of a graph.

WF-3 (verification coverage). For every acting node n ($\tau_V(n) \in$ Acting, Def. 3) and every kind $k \in \mathsf{eff}(n).\mathsf{required_kinds}$, every control path from Start to n *satisfies* k: it traverses some v with $\tau_V(v) = $ Verification, $v.\mathtt{check_kind} = k$ and $(v, n, \mathsf{verification\text{-}required\text{-}before}) \in E$, and it leaves that v by an edge other than fallback-to. The quantifiers are per path and per kind. One path may satisfy k at v and another at v', and no single node need lie on every path; the earlier name "verification dominance" promised more than this and is kept only as a label. The exit condition matters because fallback-to is a gate's failure continuation: a path that reaches v, fails and leaves by the fallback has run the check without passing it. The stronger per-gate obligation a run must meet at each arrival is Φ_1 (App. B), not this static condition.

WF-4 (fallback completeness). Every Tool-call node n whose side-effect class is in $\{\mathsf{write}, \mathsf{irreversible}\}$ has at least one outgoing fallback-to edge to a Fallback node, or sets its on_fail field to ABORT. Only Tool-call nodes carry a side-effect class, so the condition is vacuous on other types.

WF-5 (decision coverage). For every $d \in V$ with $\tau_V(d) = $ Decision, every declared branch label has a matching outgoing edge of type in $\{\mathsf{control\text{-}flow}, \mathsf{alternative}, \mathsf{guard}\}$, and the decision is total: it carries a default branch (*/else) or guard-typed outgoing edges. Joint exhaustiveness and disjointness of the guard *predicates* are not part of the static check.

WF-6 (data-flow typing). $\Delta(e)$ holds for every $e \in E$ whose type lies in $\{\mathsf{data\text{-}flow}, \mathsf{memory\text{-}read}, \mathsf{memory\text{-}write}, \mathsf{produces\text{-}artifact}\}$.

WF-7 (action permission). For every acting node n, $\mathsf{permits}(n, \mathsf{uri}(n))$ holds (Def. 3): a Tool-call or Subskill-call URI matches $\mathsf{eff}(n).$ALLOWED and misses $\mathsf{eff}(n).$FORBIDDEN; a Memory-write names a store in $\mathcal{L}(\mathsf{eff}(n).\mathrm{WRITABLE})$; an Artifact carries no condition. Quantifying

over Tool-call alone once let a delegation reach a forbidden tool through `subskill:`, the one route where the arguments leave for another party.

WF-8 (approval gating). With τ_{hitl} as above and $\text{risk}(n) = \text{eff}(n).\text{RISK}$: for every acting node n with $\text{risk}(n) \succeq \tau_{\text{hitl}}(\text{eff}(n))$, every control path from Start to n traverses, and leaves by an edge other than fallback-to, some $h \in V$ with $\tau_V(h) =$ Human-approval and $(h, n, \text{requires-approval}) \in E$. The reading is existential: *some* approval ancestor with the edge, not a unique nearest one, and Φ_3 inherits it.

WF-9 (subskill type compatibility). For every calls-subskill edge (c, s), the binding $\text{SUBSKILL-CALL}(c).\text{binding}$ is well typed against $\Sigma(s).\text{in}$ and return_binding against $\Sigma(s).\text{out}$ under $\sqsubseteq$.

WF-10 (fallback insulation). No Fallback node has an incoming verification-required-before edge. A Fallback node is not acting, so such an edge would check nothing (Φ_1 ranges over acting nodes); the condition refuses an edge that reads as a check and is not one.

WF-11 (regions composed and current, terminals root-scoped, names disjoint). Composition inlines a callee at a call site c, freezes $\text{eff}_G(c)$ as $\widehat{\pi}(c)$, records the region c sat in, and meets $\text{eff}_G(c)$ into the override of every inlined node. Eq. (1) leans on that record: a walk that crosses into another region takes the frozen value and drops what lies behind it, which is harmless only while the frozen value still says everything the dropped walk would have said. Call p a *boundary parent* of a region s if $p \in \text{par}(m)$ for some m with $\varsigma(m) = s$ and $\varsigma(p) \neq s$, and let $\kappa(p) := \text{eff}(p)$ if $\varsigma(p) = \bot$ or $\varsigma(p)$ encloses s in the recorded nesting, and $\widehat{\pi}(\varsigma(p))$ otherwise. For every region s some node names: (a) $\widehat{\pi}(s)$ and the nesting of s are recorded, and the nesting resolves to $\bot$; (b) every n with $\varsigma(n) = s$ has $\overline{\pi}(n) \sqsubseteq \widehat{\pi}(s)$; (c) every boundary parent p of s whose region s does not enclose has $\widehat{\pi}(s) \sqcap \pi \sqsubseteq \kappa(p)$. And for the graph as a whole: (d) the graph's own Start and every End have $\varsigma = \bot$; (e) no live node id is a region name, where the region names are the scope labels the nodes carry, the keys of $\widehat{\pi}$, the keys of the nesting record, and every non-$\bot$ parent it records. (a) says the region was composed; (b) says the value composition materialised is still on every node; (c) says the frozen value still says what every ancestor behind the boundary says. At a $\bot$-parent $\kappa(p)$ is what the walk takes, and that is the case Lemma 4 uses. At an enclosing parent the walk takes $\widehat{\pi}(\varsigma(p)) \sqsupseteq \text{eff}(p)$ (it never reads the nesting; the inequality is (b) at $\varsigma(p)$), and (c) asks for the stronger $\text{eff}(p)$ on purpose: an override tightened inside an enclosing region after composition is one the walk would not carry into s, and it is refused rather than silently lost. At a parent inside a region nested in s no inequality is needed: both walks stop there and add the same frozen value. A graph whose root-region ancestor was tightened after the region was composed fails (c) and is refused; tightening the root π afterwards is admitted, through the $\sqcap\,\pi$. Records alone (a) do not decide this: a graph with genuine records and a stale frozen value validated and lost a callee's restriction on composition. (d) is needed because the outside branch of Eq. (2) reconnects the caller's continuation to the callee's Ends in the new region c, where the walk stops and adds $\widehat{\pi}(c) = \text{eff}_G(c)$; an End that sits in an inner region of the callee makes the walk stop there instead and read that region's frozen value, and a caller continuation at $*$ came out at `tool:allowed` with every graph validating. Composition retypes an inlined Start and End to Goal, so a current terminal with a scope was authored, not composed, and (d) is closed under $\oplus$: the caller's own terminals are kept and the callee's become Goals. (e) is needed because $\oplus$ keys the new region by the call-site id and prefixes the callee's names with it, and prefixing is not a collision check: a live Subskill-call named c next to an existing region c (its own call site deleted by the composition that made it, its nodes still naming it) let the operator write $\widehat{\pi}(c)$ over the old value, and a retained continuation widened from $\{$`subskill:*`, `tool:allowed`$\}$ to $*$. The operator refuses, before it writes any map or node, every name it would create — the prefixed callee node ids, the prefixed callee region names and the new key c — that the caller already holds as a retained node id or as any region name; an existing region named c is a collision, never an unused key. That refusal also covers the case neither input can see, a callee region k prefixed onto an existing region $c__k$, and it is what keeps (e) closed under $\oplus$. Lemma 4 is stated on graphs satisfying WF-1–WF-11; `validate()` decides membership, and $\oplus$ preserves all five parts (App. C).

WF-10 is not used by the soundness proof, and neither is any other WF condition (App. B). Φ_1 ranges over acting nodes only, and a Fallback node is not acting, so a gate edge into one would check nothing; WF-10 refuses such an edge.

Definition 3 (Acting nodes)**.** An *acting* node is one whose step has an effect outside the run: it calls a tool, hands arguments to another party, commits a value to a store, or materialises a record for a sink. There are four, $\mathsf{Acting} = \{\mathsf{Tool\text{-}call}, \mathsf{Subskill\text{-}call}, \mathsf{Memory\text{-}write}, \mathsf{Artifact}\}$, and this is the only place the set is written down. The static checks (WF-3, WF-8), the transition rules, the three invariants, the executor and the external verifier all range over *this* set.Write $\mathsf{uri}(\rho)$ for the name of the action, which is what the event records and what the whitelist is matched against: ρ.`tool_uri` for a Tool-call; `subskill`$:\langle\rho$.`target_skill_id`$\rangle$ for a Subskill-call; `memory`$:\langle\mathsf{store}\rangle:\langle\mathsf{key}\rangle$ for a Memory-write; and `artifact`$:\langle\rho\rangle$ for an Artifact. These are the strings the executor emits.Write $\mathsf{permits}(\rho, u)$ for the type's own admissibility condition, because a store is not a tool and the whitelist does not name it:

$$\mathsf{permits}(\rho,u)=\begin{cases} u \models \mathsf{eff}(\rho).\textsc{allowed} \wedge u \not\models \mathsf{eff}(\rho).\textsc{forbidden} & \text{Tool-call, Subskill} \\ \rho.\mathsf{store} \in \mathsf{eff}(\rho).\textsc{writable-stores} & \tau_V(\rho)=\mathsf{Memory\text{-}write} \\ \top & \tau_V(\rho)=\mathsf{Artifact} \end{cases}$$

Every clause is a plain membership test against an allow-list. In particular the Memory-write clause does not read an empty WRITABLE-STORES as "no restriction declared": under that reading a *refinement* would grant authority, since a root permitting `scratch` met with a node permitting `audit` intersects to the empty list, which would then permit `prod`.

So unrestricted and deny-all have different representations. ALLOWED and WRITABLE-STORES are $\{\star\}$ when unrestricted and $\emptyset$ when nothing is permitted, an omitted field normalises to $\{\star\}$ before any meet, and the meet of two disjoint allow-lists is $\emptyset$, which denies. The invariant this restores is the one the policy lattice claims: $\mathsf{Permit}(p \sqcap q) \subseteq \mathsf{Permit}(p) \cap \mathsf{Permit}(q)$, checked exhaustively on a small domain in `test_policy_meet_is_restrictive.py`. A Subskill-call is held to the same allow-list as a Tool-call; exempting it whenever no pattern mentioned `subskill:` would let a refinement that removed the only such pattern widen the policy. This is exactly what `gating.permits` computes, and the validator, the executor and the verifier all call that one function.

A.3 Operational Semantics: Complete Rule Set

The nine small-step rules sketched in §3.3 are given in their full inference-rule form here. Each rule has the standard shape

$$\frac{\textit{premises}}{C \to C'} \quad \textsc{Rule-Name}$$

Configurations carry authorisations, and a pending fallback. A configuration is $C = (G, \rho, \sigma, T, \mathcal{A}, \kappa)$, where $\mathcal{A}$ is a finite set of *authorisations* and $\kappa \in \{\bot\} \cup (V \times \mathsf{Cause})$ is the *pending fallback*: either nothing, or the node whose step just failed together with why, $\mathsf{Cause} = \{\mathsf{verify\text{-}fail}, \mathsf{project\text{-}fail}, \mathsf{step\text{-}fail}, \mathsf{land\text{-}fail}\}$. The rules below write the first three; land-fail is a name the verifier accepts, and neither the rules nor the released executor write it. Writing FALLBACK-READY(ρ, c) means exactly $\kappa = (\rho, c)$. A failure rule *sets* κ; STEP-FALLBACK-LAND *requires* it and *consumes* it — the recorded failure is cleared before the landing node is dispatched, and the dispatched rule may record a new one, so the rule's post-state is κ', not $\bot$; every other rule carries $\kappa = \bot$ and is stuck otherwise. κ is a component of the configuration, and the executor's `pending_land` is its concrete twin. In the trace it appears as the CAUSE field of the FALLBACK-FIRE event, which the verifier requires and which is hashed into the chain. Φ does not mention κ. An authorisation is a tuple

$$a = (\mathsf{contract},\ \mathsf{node},\ u,\ d,\ \mathsf{principal},\ t_{\mathsf{iss}},\ t_{\mathsf{exp}},\ \nu_{\mathsf{req}},\ \nu_{\mathsf{dec}},\ \mathsf{spent})$$

naming the contract and the node it is for, the immutable tool URI u, the digest d of the final arguments, who decided, when it was issued and when it expires, the nonce of the request it answers and its own decision nonce, and whether it has been spent. The rules below say what the executor does.

Every rule below carries $\mathcal{A}$ unchanged unless it is the rule that issues one (STEP-APPROVE) or spends one (STEP-ACT-GATED). Likewise every rule carries $\kappa = \bot$ unless it is a rule that sets it (the failing form of STEP-VERIFY, and STEP-FAIL) or the one that consumes it (STEP-FALLBACK-LAND). A rule that quietly dropped $\mathcal{A}$ would discard every live authorisation; a rule that quietly ignored κ would let a run proceed past a failure it had recorded. The tuple is uniform across all nine.

An authorisation is issued *unbound*: d is empty until the arguments exist. STEP-APPROVE creates it. A separate judgment binds it. The two must be separate: a single predicate that required $d = \mathsf{digest}(p)$ of every grant would be met by no grant before the arguments exist, and no rule could then take a gated step. The binding is what `Executor._confirm_grant` does.

Two predicates, not one.

- $a_0 \rightsquigarrow \rho$ — *pending for* ρ: a_0 names this contract and node ρ, carries ρ's URI, has $d = \bot$, names a principal, is inside $[t_{\mathsf{iss}}, t_{\mathsf{exp}}]$, answers a request the executor issued, and is unspent.
- $a \rightharpoonup (\rho, u, p)$ — *bound to this call*: a names this contract and node ρ, carries u, the URI of the call, has $d = \mathsf{digest}(p)$, names a principal, is inside $[t_{\mathsf{iss}}, t_{\mathsf{exp}}]$, answers a request the executor issued, and is unspent.

The two lists differ in one place: the digest. Pending has $d = \bot$; bound has $d = \mathsf{digest}(p)$. The bound predicate *replaces* that condition, it does not add to it: asking for both $d = \bot$ and $d = \mathsf{digest}(p)$ would be met by no tuple. Every other field is shared, and CONFIRM below sets d, records the confirmation, and touches nothing else.

The judgment that turns the first into the second asks the authority again, over the arguments that now exist:

$$\textsc{Confirm}\big(a_0,\ G.\mathsf{id},\ \rho,\ u,\ \mathsf{digest}(p),\ t\big) \Downarrow a$$

holds iff $a_0 \rightsquigarrow \rho$, the clock t is inside a_0's window, the authority returns a confirmation c over exactly $(G.\mathsf{id}, \rho, u, \mathsf{digest}(p))$ that answers the fresh confirmation request (its own nonce, which names a_0's as the grant it confirms; `_confirm_grant`), and $a = a_0[d := \mathsf{digest}(p),\ \mathsf{conf} := c]$. It fails — so no gated step is possible — if the authority declines, the window has passed, or the confirmation covers different arguments. The confirmation is recorded in the trace, because an outside verifier that sees only a digest the provider wrote cannot tell whether anyone approved *these* arguments.

A worked instance. A purchase graph gates `place_order` at $500. The Human-approval node runs first and puts one unbound authorisation into $\mathcal{A}$. Later the model proposes `place_order(amount=900)`. Only now does $\mathsf{digest}(p)$ exist, so only now can the authorisation be bound to it: CONFIRM asks the authority about `amount=900`, and STEP-ACT-GATED spends the bound result. A second `place_order` on the same run finds no unspent authorisation and does not fire; a call with `amount=100` rewritten afterwards to `900` no longer matches d.

The arrival judgment, defined once. Two rules below move control to a node, and both carry the same obligation, which is also the obligation Φ_1 states (App. B). For a verification node v, a node n, an arrival index i and a kind k,

$$\begin{aligned}\mathsf{verified}_T(v,n,i,k) \iff\ & v.\texttt{check_kind} = k \ \wedge\ \exists j.\ \mathsf{prev}_T(n,i) < j < \mathsf{enter}_T(n,i) \\ & \wedge\ T[j] = \textsc{verify-pass}(v) \ \wedge\ \neg T[j].\mathsf{skipped},\end{aligned}$$

where $\mathsf{enter}_T(n,i)$ is the position of the i-th $\textsc{enter}(n)$ in the trace that includes the arrival being made, and $\mathsf{prev}_T(n,i)$ is the position of the $(i{-}1)$-th (and -1 when $i = 1$). All four arguments are bound in the formula; k is not a free name. The *arrival obligation* at n for arrival i is

$$\begin{aligned}\mathsf{gates}_T(n,i) \iff\ & \tau_V(n) \in \mathsf{Acting} \ \Rightarrow\ \forall k \in \mathsf{eff}(n).\mathsf{required_kinds}\ \forall v. \\ & (v, n, \mathsf{verification\text{-}required\text{-}before}) \in E \ \wedge\ v.\texttt{check_kind} = k \\ & \Rightarrow\ \mathsf{verified}_T(v,n,i,k).\end{aligned}$$

It ranges over acting nodes, over the kinds the effective policy requires, and over every verification node of that kind wired to n — the same three restrictions as Φ_1, and the same the executor (`_verification_satisfied`) and the verifier compute. A node that is not acting carries no obligation. In words: every required check must have passed *since the last time we were here*, must be a check that actually ran (a VERIFY-PASS carrying skipped is what the executor writes when

verification is switched off, and counting it once certified a run in which a required check never ran), and must be of the required kind. The executor and the verifier both compute the interval (`a2a_adapter.py`, `entered_all`), and Φ_3 is checked per arrival for the same reason.

The moves the executor offers, defined once. For a node ρ and a store σ, $\mathsf{next}_\sigma(\rho)$ is the set of nodes the executor offers as the next arrival on an ordinary move (`constrained_actions.allowed_actions`):

$$\begin{aligned}
\mathsf{next}_\sigma(\rho) &= \{\rho' \in \mathsf{branches}(\rho) \mid \mathsf{open}_\sigma(\rho, \rho')\} && \text{if } \tau_V(\rho) = \mathsf{Decision},\\
\mathsf{next}_\sigma(\rho) &= \{\rho' \mid (\rho, \rho', t) \in E,\ t \in M\} \\
&\quad \cup\ \{\rho' \mid e = (\rho, \rho', \mathsf{guard}) \in E,\ \mathsf{holds}_\sigma(e)\} && \text{otherwise,}
\end{aligned}$$

where $M = \{\mathsf{control\text{-}flow}, \mathsf{calls\text{-}subskill}, \mathsf{alternative}\}$ are the edge types a move follows without a condition. $\mathsf{branches}(\rho)$ is the set of targets a Decision names in its own spec. At a Decision the executor offers these targets and nothing else, whatever edges leave it. $\mathsf{holds}_\sigma(e)$ says that the predicate of the guard edge e is true on σ: an absent predicate holds, and one that cannot be evaluated does not, so a guard whose condition cannot be established opens no path. $\mathsf{open}_\sigma(\rho, \rho')$ holds when no guard edge joins ρ to ρ', or when some guard edge from ρ to ρ' holds on σ. The executor reads one guard edge per pair, the last in edge order; the rule admits the pair when any of them holds, which includes the one the executor reads. No fallback-to edge is in $\mathsf{next}_\sigma(\rho)$, because only STEP-FALLBACK-LAND crosses one. A Decision built in memory with no spec object is read by the second case, as the executor reads it; the parser always gives a Decision its spec. Every node in $\mathsf{next}_\sigma(\rho)$ is a successor of ρ in the relation WF-1 and WF-3 read, $E^{\rightarrow}$ plus the branch targets (`validator._successors`), so the static checks consider every move the executor can make.

STEP-CTRL.

$$\frac{\begin{array}{c}\rho' \in \mathsf{next}_\sigma(\rho) \\ \mathsf{gates}_{T'}(\rho', i)\end{array}}{(G, \rho, \sigma, T, \mathcal{A}, \bot) \to (G, \rho', \sigma', T', \mathcal{A}, \bot) \qquad T' = T \cdot \textsc{enter}(\rho')}$$

where σ' applies to σ the data-flow bindings on edges from ρ to ρ' (`Executor._apply_data_flow`), and i is the index of the arrival being made: $\textsc{enter}(\rho')$ is the i-th ENTER event at ρ' in T'.

The rule covers every ordinary move the executor makes: along a control-flow, calls-subskill or alternative edge, along a guard edge whose predicate holds, and to a branch target a Decision names. Only the first premise depends on the kind of move, and the soundness proof does not read it, so the induction is the same for every kind (App. B).

One check licenses one arrival. The interval in $\mathsf{verified}_T$ starts at the previous arrival at ρ', so a check that passed before an earlier arrival does not count for this one. On a retry loop this matters: when a failed post-check falls back to the precondition check, each retry must pass the check again before the irreversible tool can fire again. The executor tests the premise before it writes $\textsc{enter}(\rho')$ (`Executor._verification_satisfied`) and ends the run with an abort when it fails.

Entering a node needs no authorisation. An arrival carries only the verification obligation. The approval gate is tested at the action, by STEP-ACT-FREE, STEP-ACT-GATED and STEP-ACT-CONDITIONAL, because only they see the arguments p that $\mathsf{approval\text{-}required}(G, \rho, p)$ reads. A node gated by its risk label or a value guard alone, with no Human-approval node pointing at it, is entered by STEP-CTRL like any other node, and STEP-ACT-CONDITIONAL then asks the authority at the final arguments. That is what the executor does: it writes $\textsc{enter}(\rho')$, and `_authorize` asks at the call.

The three rules below range over *all four* acting node types (Def. 3), because all four have an effect outside the run and all four are gated the same way. They are named STEP-ACT-$*$ rather than STEP-TOOL-$*$ for that reason. Write $\mathsf{uri}(\rho)$ for the name of the action, which is what the event records and what the whitelist is matched against: ρ.`tool_uri` for a Tool-call; `subskill`:$\langle\rho$.`target_skill_id`$\rangle$ for a Subskill-call; `memory`:$\langle\mathsf{store}\rangle$:$\langle\mathsf{key}\rangle$ for a Memory-write; and `artifact`:$\langle\rho\rangle$ for an Artifact. These are the strings the executor emits. Write $\mathsf{permits}(\rho, u)$ for the type's own admissibility condition, because a store is not a tool and the

whitelist does not name it:

$$\mathsf{permits}(\rho, u) = \begin{cases} u \models \mathsf{eff}(\rho).\textsc{allowed} \wedge u \not\models \mathsf{eff}(\rho).\textsc{forbidden} & \text{Tool-call, Subskill} \\ \rho.\mathsf{store} \in \mathsf{eff}(\rho).\textsc{writable-stores} & \tau_V(\rho) = \mathsf{Memory\text{-}write} \\ \top & \tau_V(\rho) = \mathsf{Artifact} \end{cases}$$

Every clause is a plain membership test against an allow-list. In particular the Memory-write clause does not read an empty WRITABLE-STORES as "no restriction declared": under that reading a *refinement* would grant authority, since a root permitting `scratch` met with a node permitting `audit` intersects to the empty list, which would then permit `prod`.

So unrestricted and deny-all have different representations. ALLOWED and WRITABLE-STORES are $\{\star\}$ when unrestricted and $\emptyset$ when nothing is permitted, an omitted field normalises to $\{\star\}$ before any meet, and the meet of two disjoint allow-lists is $\emptyset$, which denies. The invariant this restores is the one the policy lattice claims: $\mathsf{Permit}(p \sqcap q) \subseteq \mathsf{Permit}(p) \cap \mathsf{Permit}(q)$, checked exhaustively on a small domain in `test_policy_meet_is_restrictive.py`. A Subskill-call is held to the same allow-list as a Tool-call; exempting it whenever no pattern mentioned `subskill:` would let a refinement that removed the only such pattern widen the policy. This is exactly what `gating.permits` computes, and the validator, the executor and the verifier all call that one function. The rules, the executor and the verifier must range over one set, or the theorem describes a smaller machine than the one released. `check_source_sync.py` reads the set out of `sections/_acting_set.tex` and fails if it and `validator._ACTING_TYPES` differ.

One action per arrival. $\mathsf{fresh}_T(\rho)$ holds when no ACT event for ρ appears after the most recent $\textsc{enter}(\rho)$ in T — that is, this arrival has not acted yet.

It is a safety condition, not a termination argument. The three acting rules leave ρ unchanged, so without it the rules could apply the same acting rule at the same node again and append a second ACT event on the same arrival; Φ_3's one-approval-one-call correspondence is stated per gated event — exactly one bound approval for each, and none consumed twice, the second half coming from the spend state (App. B) — and this premise is what makes "per arrival" mean at most one action. The repeat is a counterexample to the rules without the premise, not a loop the executor runs, because the executor enters, steps, exits and then moves on. fresh_T is that behaviour, stated. What bounds traversal of a fallback-to edge is a different premise, the per-edge counter in STEP-FALLBACK-LAND below; the two are independent, and neither is offered as a proof that every run ends (App. B says what is and is not claimed).

Write $\textsc{act}(\rho, u, p)$ for the event the step appends. It is one event with three wire names, because that is what the trace carries: TOOL-CALL for a Tool-call and a Subskill-call, MEMORY-WRITE for a Memory-write, ARTIFACT for an Artifact. All three carry the same fields — node, u, arguments p, the recorded guard inputs, and the decision nonce when the call was gated — and $\mathsf{acts}(T)$ is the set of all of them in T. The invariants Φ_2 and Φ_3 quantify over $\mathsf{acts}(T)$, which is exactly the set the external verifier checks (`a2a_adapter.py`, check 5). Similarly $\textsc{result}(u, p)$ is the type's own effect on the store: the tool's or subskill's return value bound to the output ports, the committed value at `memory:`$\langle\mathsf{store}\rangle$`:`$\langle\mathsf{key}\rangle$, or the materialised record at `artifact:`$\langle\rho\rangle$.

STEP-ACT-FREE (the action is not gated).

$$\frac{\begin{array}{c}\tau_V(\rho) \in \mathsf{Acting} \\ u = \mathsf{uri}(\rho), \quad \mathsf{permits}(\rho, u), \quad \mathsf{fresh}_T(\rho) \\ \neg\, \mathsf{approval\text{-}required}(G, \rho, p)\end{array}}{(G, \rho, \sigma, T, \mathcal{A}, \bot) \to (G, \rho, \sigma[\textsc{result}(u, p)], T \cdot \textsc{act}(\rho, u, p), \mathcal{A}, \bot)}$$

No authorisation is named, no APPROVE event is written, and $\mathcal{A}$ is untouched. Two rules are shorter than one rule with an optional metavariable, and they are what the executor does.

STEP-ACT-GATED (the action is gated).

$$\frac{\begin{array}{c}\tau_V(\rho) \in \textsf{Acting} \\ u = \textsf{uri}(\rho), \quad \textsf{permits}(\rho, u), \quad \textsf{fresh}_T(\rho) \\ \textsf{approval-required}(G, \rho, p) \\ a_0 \in \mathcal{A}, \quad a_0 \rightsquigarrow \rho \quad \textsc{Confirm}(a_0, G.\textsf{id}, \rho, u, \textsf{digest}(p), t) \Downarrow a\end{array}}{\begin{array}{c}(G, \rho, \sigma, T, \mathcal{A}, \bot) \to \\ (G, \rho, \sigma[\textsc{result}(u, p)],\ T \cdot \textsc{approve}(a) \cdot \textsc{act}(\rho, u, p),\ \mathcal{A}[a_0.\textsf{spent} := \top], \bot)\end{array}}$$

$\textsf{approval-required}(G, n, p)$ holds iff *any* of three things is true: n has an incoming requires-approval edge from a Human-approval node; *or* $\textsf{risk}(n) \succeq \tau_{\textsf{hitl}}(\textsf{eff}(n))$, where $\tau_{\textsf{hitl}}$ is High when the effective policy declares HITL triggers and Critical otherwise; *or* the effective guard $g_{\textsf{eff}(n)}$ holds on the final arguments p, read together with the *guard inputs*: the values of the fields the guard names that the executor resolves from the store and records beside the arguments, which the verifier reads too. A guard that cannot be evaluated counts as holding. This happens when neither the arguments nor the recorded guard inputs carry a field it names. Reading it the other way would let a provider leave the field out of what it writes down and walk through the gate. It is defined once (`sections/_gate_definition.tex`) and used by the short method, both rules above, the soundness proof and the implementation (`skillgraph/gating.py`). The second-to-last premise picks one pending grant; when several are pending for ρ, the executor takes the earliest issued, and the rule admits any. CONFIRM binds it to *these* arguments; the spend makes it unusable again. A second call at the same node needs another pending grant, with its own decision nonce. A call whose arguments are rewritten afterwards no longer matches the d inside the recorded $\textsc{approve}(a)$.

STEP-BIND.

$$\frac{\begin{array}{c}(m, \rho, \textsf{data-flow}) \in E \\ \textsc{ready}(m, \sigma) \quad \text{(the source port has a value in } \sigma)\end{array}}{(G, \rho, \sigma, T, \mathcal{A}, \bot) \to (G, \rho, \sigma[\textsc{bind}(m, \rho)], T, \mathcal{A}, \bot)}$$

This rule moves a value along a data-flow edge into the target's input port. It writes no event and touches neither the trace nor $\mathcal{A}$, which is why it is easy to leave out; the soundness proof has a case for it. Type safety of the assignment is Definition 1 and is orthogonal to Φ.

Where a conditional grant comes from. STEP-APPROVE issues a pending grant when a Human-approval node runs, and that node reaches its targets through a requires-approval edge. So a node gated *only* by risk or a value guard has no route to a pending grant — and the edge that would give it one is itself an unconditional gate, which is the opposite of what a guard is for. With only the first gated rule, such a call could only ever be refused: with `amount >= 500`, `amount=900` and an honest authority available, there would be no authority call and no tool call. That is conditional *blocking*, not conditional *approval*.

So there is a second gated rule. It fires when no pending grant covers ρ and the gate came from risk or the guard, and it asks the authority *at the final arguments* — the only moment the condition can be evaluated. The returned authorisation is already bound, so there is no unbound stage to confirm.

STEP-ACT-CONDITIONAL.

$$\frac{\begin{array}{c}\tau_V(\rho) \in \textsf{Acting} \\ u = \textsf{uri}(\rho), \quad \textsf{permits}(\rho, u), \quad \textsf{fresh}_T(\rho) \\ \textsf{approval-required}(G, \rho, p) \quad \wedge \quad \neg\, \textsf{structural-gate}(G, \rho) \\ \nexists\, a_0 \in \mathcal{A}.\ a_0 \rightsquigarrow \rho \quad \textsc{Ask}\big(G.\textsf{id}, \rho, u, \textsf{digest}(p), t\big) \Downarrow a\end{array}}{\begin{array}{c}(G, \rho, \sigma, T, \mathcal{A}, \bot) \to \\ (G, \rho, \sigma[\textsc{result}(u, p)],\ T \cdot \textsc{approve}(a) \cdot \textsc{act}(\rho, u, p),\ \mathcal{A}, \bot)\end{array}}$$

ASK is CONFIRM without a pending grant to confirm: the authority is given the contract, node, URI and argument digest, and returns an authorisation already bound to them, carrying a fresh $(\rho, \nu_{\textsf{dec}})$ and its own confirmation. It fails — so the step is impossible — if there is no authority, if the authority declines, or if the confirmation covers different arguments. $\mathcal{A}$ is unchanged because nothing pending was consumed; the record is written with $\textsf{conditional} = \top$ so an outside reader can see which of the two routes produced it. The second premise is what keeps the two rules disjoint: a structurally gated node has a route to a pending grant and must use it.

A verifier must then accept a bound record with no pending predecessor, and that exception is exactly where a hole would live, so it is narrow: the record is accepted only at a node the structural edge does *not* gate, and it faces the same field audit as any other — contract, node, URI, argument digest, principal, window, and the authority's own confirmation. A record claiming conditional at a structurally gated node is rejected, because that node does have a route to a pre-grant and the claim would be a way around it (`test_conditional_gate.py`). With no authority configured the answer is still no.

STEP-VERIFY.

$$\frac{\tau_V(\rho) = \textsf{Verification} \quad \textsc{check-fires}(\rho.\textsf{spec}, \sigma) = \textsf{pass}}{(G, \rho, \sigma, T, \mathcal{A}, \bot) \to (G, \rho, \sigma, T \cdot \textsc{verify-pass}(\rho), \mathcal{A}, \bot)}$$

$$\frac{\tau_V(\rho) = \textsf{Verification} \quad \textsc{check-fires}(\rho.\textsf{spec}, \sigma) = \textsf{fail} \quad \rho.\textsf{spec.on_fail} = \textsf{fallback} \quad (\dagger)}{(G, \rho, \sigma, T, \mathcal{A}, \bot) \to (G, \rho, \sigma, T \cdot \textsc{verify-fail}(\rho), \mathcal{A}, (\rho, \textsf{verify-fail}))}$$

STEP-VERIFY never traverses an edge, and both forms leave ρ where it is. The successful form is *stationary*: it appends VERIFY-PASS(ρ) and nothing else, and control leaves ρ only by STEP-CTRL, whose arrival obligation $\textsf{gates}_{T'}(\rho', i)$ then reads the pass just recorded. That is one arrival semantics for the whole rule set: every change of the current node appends an ENTER and carries the arrival obligation, in STEP-CTRL or in STEP-FALLBACK-LAND, and no other rule moves control. A successful check that also advanced control would be a second, unlogged notion of arrival, and the theorem needs no such step (App. B). As displayed, the successful form can fire again at the same node; nothing in Φ depends on how many passes an arrival records, and the released executor runs one check per arrival as a matter of its own control loop, which we do not turn into a premise. The failing form records the outcome and, with $\textsf{on_fail} = \textsf{fallback}$, *sets* $\kappa = (\rho, \textsf{verify-fail})$ — that is, FALLBACK-READY(ρ, c) with cause c; STEP-FALLBACK-LAND below is the *only* rule that moves to f, and it is the only rule that consumes κ: the failure it fires on is cleared before f is dispatched, and the rule at f may then record a new one. Every other failure — a proposal that does not project, a refused or failed step, a denied approval — sets κ through STEP-FAIL below, with cause project-fail or step-fail; the executor writes each cause on the FALLBACK-FIRE it then emits. The retry variant of VERIFY-FAIL re-enters ρ through that same rule.

(†) *This premise is enforced.* The executor reads `VerificationSpec.on_fail` at both places a step can fail — on an ordinary arrival, and after a fallback lands — and a check that declares $\textsf{on_fail} = \textsf{abort}$ ends the run after its VERIFY-FAIL with no FALLBACK-FIRE, even when a fallback edge exists. The verifier checks the same thing from the outside: a trace in which a node whose sealed spec says abort fires a fallback after a failed step is rejected. So the displayed failing form applies exactly when $\textsf{on_fail} = \textsf{fallback}$, which is what (†) says. No released graph declares abort.

`FallbackSpec.max_retries` is different: it is sealed into every released graph and read by nothing, the bound being the protocol constant B. We keep it in the wire format rather than change the sealed bytes during a freeze, and we say here that it means nothing. STEP-VERIFY appends neither FALLBACK-FIRE nor ENTER(f); a regression test fails the build if it does, because a check that moved control itself would enter f without the retry premise and without the arrival premise.

STEP-FALLBACK-LAND.

$$\frac{\kappa = (\rho, c) \quad \text{i.e. } \textsc{fallback-ready}(\rho, c) \quad (\rho, f, \textsf{fallback-to}) \in E \quad \textsf{retries}_T(\rho, f) < \textsc{max-retries-per-edge} \quad \textsf{gates}_{T'}(f, i)}{\begin{array}{c}(G, \rho, \sigma, T, \mathcal{A}, (\rho, c)) \to \\ (G, f, \sigma', T' \cdot \textsc{step}(f), \mathcal{A}', \kappa') \qquad T' = T \cdot \textsc{fallback-fire}(\rho, c) \cdot \textsc{enter}(f)\end{array}}$$

A fallback edge *lands on* its target, and the target then runs. There is no type restriction on f: it may be a Fallback node, but the useful and common case is a retry back to the Verification that failed, and the released validator has always accepted that shape. STEP(f) abbreviates "and then the rule for $\tau_V(f)$ applies at f" — σ', $\mathcal{A}'$, κ' and the appended events are *all* whatever that rule produces.

Landing can change the authorisation set and can leave a new failure pending. If f is a Human-approval node the dispatched rule extends $\mathcal{A}$. If f is a Verification that fails again, the dispatched rule is the failing form of STEP-VERIFY and sets $\kappa' = (f, \textsf{verify-fail})$, which is the cause the executor writes (`Executor._cause_for`). So the rule threads the whole post-state, and $\kappa' = \bot$ only when the rule at f leaves nothing pending. A step at f that fails in any other way is followed by STEP-FAIL at f.

Two parts of the rule carry its safety. The rule puts no type restriction on f, because the validator accepts any target and a graph it passes must have a rule to execute. And the landing runs f's own rule, STEP(f): a fallback that made f the current node without running it would pass a required Verification with the check never run, and the irreversible call downstream would fire. The fourth premise is the same arrival obligation $\textsf{gates}_{T'}(f, i)$ as STEP-CTRL, with i the index of the ENTER(f) being appended: if f is acting, every v joined to f by a verification-required-before edge and of a kind $\textsf{eff}(f).\textsf{required_kinds}$ requires must have passed since the previous arrival at f. Arriving by a fallback edge is still an arrival.

The retry bound is a real counter. $\textsf{retries}_T(\rho, f)$ counts the FALLBACK-FIRE events for this edge already in T, and MAX-RETRIES-PER-EDGE is a *protocol constant*, fixed at 3 and identical for every party. It is not a provider-local option: a bound the provider could set would be absent from the sealed bytes, and the holder of the commitment could not tell which bound had governed a run. `skillgraph/protocol_limits.py` holds the constant, the executor cannot be asked for more, and `verify_trace` counts FALLBACK-FIRE events per edge and rejects a trace that exceeds it.

STEP-SUBSKILL: there is no such runtime rule. Composition happens *before* the run, not during it. Negotiation produces $G \oplus_c G'$, the provider re-seals that graph, and execution starts from C_0 with the composed graph already in place. So the rule set has no transition that rewrites G mid-run.

The released executor never rewrites G mid-run. It treats a Subskill-call as what it is, a gated call that hands work to another party: it resolves the bindings, evaluates $\textsf{approval-required}(G, \rho, p)$ on the resolved arguments, consumes an authorisation if one is required, and emits a TOOL-CALL event with URI `subskill:`$\langle$id$\rangle$ (`Executor._step_subskill`).

So a Subskill-call node is covered by STEP-ACT-FREE, STEP-ACT-GATED or STEP-ACT-CONDITIONAL, with $u =$ `subskill:`$\langle$id$\rangle$, and it inherits Φ_2 and Φ_3 for free. One execution model, and delegation is inside it rather than beside it.

Consistency with the policy lattice. The effective policy inside an inlined region is given by Def. 2 on the composed graph, and Eq. (2) in App. C evaluates it: $\textsf{eff}_{G \oplus_\rho s}(n) = \textsf{eff}_G(\rho) \sqcap \textsf{eff}_s(n)$ for $n \in V(s)$. The preservation argument for a delegated call in the soundness proof relies on this and on no new machinery.

STEP-APPROVE.

$$\frac{\begin{array}{c}\tau_V(\rho) = \textsf{Human-approval} \\ \text{APPROVE-DECIDED}(\rho.\textsf{spec}, \sigma) = \textsf{yes} \\ \forall n \in \textsf{protects}(\rho).\ (n, \nu_{\textsf{dec}}) \notin \textsf{issued}(T) \quad \text{(freshness)} \\ a_n = (G.\textsf{id},\ n,\ u_n,\ \bot,\ \textsf{principal},\ t,\ t + \Delta,\ \nu_{\textsf{req}},\ \nu_{\textsf{dec}},\ \bot) \quad \text{for each } n \in \textsf{protects}(\rho)\end{array}}{\begin{array}{l}(G, \rho, \sigma, T, \mathcal{A}, \bot) \rightarrow \\ (G, \rho, \sigma, T \cdot \text{APPROVE}(\rho, \nu_{\textsf{dec}}),\ \mathcal{A} \cup \{a_n \mid n \in \textsf{protects}(\rho)\}, \bot)\end{array}}$$

One decision issues one pending grant for each node ρ protects, and the trace records it once, as the executor does (`Executor._issue_grants`). $\textsf{protects}(\rho)$ is the set of nodes joined to ρ by a requires-approval edge, in either direction, or $\{\rho\}$ when there is none (`Executor._protected_by`). `validate()` admits only the edges that leave ρ (WF-1), and a grant for ρ itself is never spent, because ρ does not act. $\textsf{issued}(T)$ is the set of pairs $(n, \nu_{\textsf{dec}})$ for which T already holds an APPROVE event issuing a grant for n with decision nonce $\nu_{\textsf{dec}}$.

The freshness premise is what stops a decision being reissued for the same node after its first grant has been spent. It is keyed on the *pair*, not on the nonce alone, because one authority decision may legitimately cover two different guarded nodes — $\textsf{protects}(\rho)$ can hold more than one — and the released verifier keys replay the same way (`seen_bind` is a set of $(\textsf{node}, \nu_{\textsf{dec}})$ pairs).

This rule creates *pending* authorisations and nothing more. Their argument digest is $\perp$, because at this point in the run the arguments do not exist: the model has not yet proposed the call. STEP-ACT-GATED is what binds a pending grant to the arguments, checks every field, and spends it.

A DENY is not a rule of its own. When the authority refuses, answers with a decision that fails a binding check, cannot be reached, or does not answer within the executor's approval timeout (`ExecutorConfig.approval_timeout_s`; the spec's deadline_ms is not read), the step at ρ has failed. STEP-FAIL below appends the DENY record, adds no authorisation, and sets $\kappa = (\rho, \textsf{step-fail})$. The run then continues only by STEP-FALLBACK-LAND along a fallback-to edge that leaves ρ, or it ends. There is no deny edge. Everything is logged immutably in T, so verify-trace can check the same bijection an external party cares about: one gated call, one spent authorisation, bound to that call's own arguments.

STEP-FAIL.

$$\frac{c = \textsf{project-fail} \quad \text{or} \quad \big(c = \textsf{step-fail} \wedge \tau_V(\rho) \neq \textsf{Verification} \wedge \rho.\textsf{spec.on_fail} \neq \textsf{abort}\big) \qquad r \in \{\epsilon,\ \text{DENY}(\rho),\ \text{GUARD-GATE}(\rho),\ \text{POLICY-VIOLATION}(\rho)\}}{(G, \rho, \sigma, T, \mathcal{A}, \perp) \to (G, \rho, \sigma, T \cdot r, \mathcal{A}, (\rho, c))}$$

This rule records that the run cannot go on from ρ by an ordinary move. It covers every failure the executor turns into a pending fallback other than a failed check: a proposal that projects onto no move in $\textsf{next}_\sigma(\rho)$ (project-fail); a step the executor refuses before it acts, with a GUARD-GATE or POLICY-VIOLATION record (a closed gate, a URI or store the policy does not permit, an unresolved argument of an irreversible call, an unreachable witness); a denied approval, with its DENY record; and an action whose tool raised after the acting rule wrote its ACT event, which the executor records inside that event and follows with no further record ($r = \epsilon$). Its cause is the one the executor writes (`Executor._cause_for`): a failed Verification is a verify-fail and goes through STEP-VERIFY, every other failed step is a step-fail. A node whose spec says on_fail = abort ends the run on a failed step instead, in the executor and in the verifier alike; a projection failure is not a failed step, and the premise exempts it for the same reason.

The rule leaves ρ, σ and $\mathcal{A}$ unchanged and appends no ENTER, no ACT and no APPROVE event, so it cannot change what Φ says about the trace (App. B). It admits more than the executor does: it may fire at any point, for any reason, where the executor fails only when a tool, an authority, a check or the model's proposal makes it. A model that allows more failures has more runs, and the theorem covers all of them.

What the rules claim about the released executor. Proposition 1 is proved of these rules. It carries over to the released executor for the runs that are runs of the rules, $\textsf{Runs(impl)} \subseteq \textsf{Runs(model)}$, and we test that inclusion rather than prove it. The comparison is event by event. The executor's trace also carries records that no rule writes and Φ does not mention — EXIT, ABORT, the MEMORY-READ record, the model's proposal text — and they are dropped before comparing. What remains must come from the rules in order: every ENTER from STEP-CTRL or STEP-FALLBACK-LAND with its premises true, every ACT event from one of the three acting rules, every bound APPROVE from STEP-ACT-GATED or STEP-ACT-CONDITIONAL, every pending APPROVE from STEP-APPROVE, and every FALLBACK-FIRE from STEP-FALLBACK-LAND after a rule that set κ. `test_model_edges_conformance.py` replays executor traces through a checker that applies exactly these rules. It runs the executor on small graphs, in runs that take an alternative edge, a calls-subskill edge, a guard edge whose predicate is true and one whose predicate is false, a Decision branch opened by a guard, one on an alternative edge and one with no edge at all, a risk-gated node with no approval node (once granted and once refused), a denied approval followed by a fallback, and a tool that raises followed by a fallback; and on the `travel_booking` demo run of App. N. Every one of these traces is a run of the rules. The same test shows that the checker can fail. The rules with STEP-CTRL restricted to control-flow edges, an authorisation required to enter a risk-gated node, and no STEP-FAIL reject each of the small-graph traces listed above, and accept an ordinary run through a Human-approval node; and a trace edited to enter a node outside $\textsf{next}_\sigma(\rho)$ is rejected. `test_one_fallback_rule.py` pins the on_fail premise at both executor sites and in the verifier.

Two differences are known. First, the rules describe the executor with its action menu on (`enable_constrained_actions`, the default). The ablation arm that turns the menu off offers

every control-flow, alternative and calls-subskill edge, including those that leave a Decision, and no guard edge. Its runs are not claimed to be runs of the rules, although the executor still tests the arrival obligation before each ENTER and the gate before each action on them. Second, when a fallback-to edge points at End, the executor records the fire and ends the run with outcome failure without writing ENTER(End), which STEP-FALLBACK-LAND would append. End does not act, so no invariant reads that event.

B SOUNDNESS: THE COMPLETE INDUCTION

This appendix proves Proposition 1. It is the only proof of it in the paper, and both builds include this file. G is any procedure graph: the induction uses no well-formedness condition, because every rule carries as a premise what its events need. What well-formedness adds is proved after it (Lemma 3), with what a caller gets inside a composed region (Theorem 1).

What is being proved. Let $C = (G, \rho, \sigma, T, \mathcal{A}, \kappa)$ be a configuration: graph, current node, store, trace, the set of live authorisations, and the pending fallback κ (App. A). Φ does not mention κ: it constrains what the trace records, and κ is control state that decides which rule may fire next. Let $\mathcal{R}(G)$ be the configurations reachable from $C_0 = (G, \mathsf{Start}, \sigma_0, \text{ENTER}(\mathsf{Start}), \emptyset, \bot)$ under the rules of Appendix A; the trace opens with the arrival at Start, as the executor's does. Write $\mathsf{toolcalls}(T)$ for the TOOL-CALL events in T, and $\mathsf{bound}(T)$ for the APPROVE(a) events whose a has $d \neq \bot$ — the ones written by STEP-ACT-GATED after CONFIRM and by STEP-ACT-CONDITIONAL after ASK, not the pending ones written by STEP-APPROVE.

Four parts of Φ_1's antecedent are easy to leave in prose, and we do not. It ranges over *acting* nodes only; over the kinds in $\mathsf{eff}(n).\mathsf{required_kinds}$ rather than over every drawn edge; with k bound in the formula rather than in a sentence beneath it; and over *every* verification node v of kind k wired to n, not over one gate per kind. Each part matches what the executor (`_verification_satisfied`), the validator (`_wf3`) and the verifier compute.

The second part has a consequence worth stating. A graph whose root policy declares no required kinds, with an irreversible tool and a verification node joined to it by a verification-required-before edge, passes all eleven conditions, and Φ_1 asks nothing of the tool: the edge is inert, and the executor fires the tool with no verification event in the trace. `validate` keeps the graph well-formed and emits WF-3-inert naming the edge (`_wf3_inert_edges`).

The fourth part is what makes a spliced check bind. Put the provider's own static check on one branch and the consumer's spliced static check on another, both drawn verification-required-before to one irreversible tool call. Under a one-gate-per-kind reading the run could take the provider's branch, the tool would fire, and the consumer's check would never run: a provider would discharge a consumer's negotiated check by owning any check that shares its kind name. The executor and the external verifier require each gate separately (`test_phi1_every_gate.py`, which fails if either reverts). WF-3 stays per kind on purpose: as a static path property the per-gate form rejects graphs whose runs are fine, including our own spliced procurement graph, so it is a necessary condition checked in advance, and the per-gate obligation is enforced at run time and again at audit. A graph can therefore pass all eleven checks and still abort; Φ_1 holds on every run, which is what the theorem claims, and completion is not promised.

$$
\begin{aligned}
\Phi_1(C) = \forall n \in \mathsf{Acting}\ \forall i \le \mathsf{arrivals}_T(n)\ \forall k \in \mathsf{eff}(n).\mathsf{required_kinds}\ \forall v.\\
(v, n, \mathsf{verification\text{-}required\text{-}before}) \in E\\
\wedge\ v.\mathsf{check_kind} = k \Rightarrow \mathsf{verified}_T(v, n, i, k)\\
\Phi_2(C) = \forall (n, u, p) \in \mathsf{acts}(T)\ \big[\,\mathsf{permits}(n, u)\,\big]\\
\Phi_3(C) = \forall e = (n, u, p, \nu) \in \mathsf{acts}(T).\ \mathsf{approval\text{-}required}(G, n, p) \Rightarrow\\
\exists!\ \text{APPROVE}(a) \in \mathsf{bound}(T)\ \text{before}\ e.\\
\mathsf{consumes}(e, a) \wedge a \rightharpoonup (n, u, p)
\end{aligned}
$$

Three definitions carry the weight here, and each has to be as large as the system it describes.

Acting, once. $\mathsf{acts}(T)$ is the set of ACT events — TOOL-CALL, MEMORY-WRITE, ARTIFACT — and $\mathsf{permits}$ is the type's own admissibility condition (Def. 3). The executor and the external verifier gate all four acting types, so Φ_2 and Φ_3 quantify over all four; over TOOL-CALL events alone the theorem would be true of a smaller machine than the one released.

Arrivals, not nodes. $\mathsf{arrivals}_T(n)$ is the number of ENTER(n) events in T, and $\mathsf{verified}_T(v, n, i, k)$ is the occurrence-local judgment defined once, before STEP-CTRL (App. A), with all four arguments bound: a non-skipped VERIFY-PASS(v) with v of kind k, after the previous arrival at n and before the i-th. The antecedent of Φ_1 for one pair (n, i) is exactly the arrival obligation $\mathsf{gates}_T(n, i)$ that STEP-CTRL and STEP-FALLBACK-LAND carry as a premise; the invariant and the rules use one signature and one antecedent. Over the *set* of nodes entered, one pass would discharge the obligation for every later entry, and on a retry loop one check would license every irreversible call behind it. Φ_3 is checked per arrival for the same reason.

Two further definitions are carried by the events themselves rather than by prose.

One gating predicate. $\mathsf{approval\text{-}required}(G, n, p)$ holds iff *any* of three things is true: n has an incoming $\mathsf{requires\text{-}approval}$ edge from a $\mathsf{Human\text{-}approval}$ node; *or* $\mathsf{risk}(n) \succeq \tau_{\mathsf{hitl}}(\mathsf{eff}(n))$, where τ_{hitl} is High when the effective policy declares HITL triggers and $\mathsf{Critical}$ otherwise; *or* the effective guard $g_{\mathsf{eff}(n)}$ holds on the final arguments p, read together with the *guard inputs*: the values of the fields the guard names that the executor resolves from the store and records beside the arguments, which the verifier reads too. A guard that cannot be evaluated counts as holding. This happens when neither the arguments nor the recorded guard inputs carry a field it names. Reading it the other way would let a provider leave the field out of what it writes down and walk through the gate. One definition lives in `skillgraph/gating.py`, and the semantics, the executor, the verifier and the witness path all call it. Two definitions would be an attack: a verifier that read only the structural edge would accept a forged trace of a well-formed low-risk graph guarded by `amount >= 500`, with a call at `amount=900` and no approval event. `test_gate_agreement.py` holds that trace.

One consumption relation, scoped to the node. A gated TOOL-CALL event carries the decision nonce ν of the authorisation it spent, and

$$\mathsf{consumes}(e, a) \iff e.\nu = a.\nu_{\mathsf{dec}} \ \wedge \ e.\mathsf{node} = a.\mathsf{node}.$$

The node is part of the key on purpose, and the implementation agrees (`seen_bind` is a set of (node, ν) pairs). One authority decision may legitimately cover two different guarded nodes; what it may not do is cover the same node twice. Saying "a nonce is used once" globally would forbid the first and is not what is enforced. Two properties are wanted here, and they are not the same property. *Each gated call has exactly one matching bound approval* at its node: that is the $\exists!$ of Φ_3, and it makes the covering approval a function of the event. *No bound approval is consumed by two calls*: that is injectivity of the same function, and Φ_3 as displayed does not state it. Take two distinct gated events e_1, e_2 at one node ρ with equal arguments, so equal digests; a covering approval for e_1 is also bound to e_2's arguments, and nothing in the $\exists!$ stops it from being e_2's witness as well. What stops it is the spend state. STEP-ACT-GATED sets $\mathcal{A}[a_0.\mathsf{spent} := \top]$ when it writes e_1, so a_0 is no longer pending for e_2; STEP-APPROVE's freshness premise forbids issuing the pair $(\rho, \nu_{\mathsf{dec}})$ again; so e_2 needs a grant with a different decision nonce and its covering approval is a different event (the STEP-ACT-GATED case below spells this out). The executor's `Executor._authorize` is that transition: it skips every grant with `consumed` set, marks the grant it spends `consumed`, and stamps the event with that grant's nonce. Read together, the two properties make the $\exists!$ a bijection between gated calls and bound approvals *at that node*, and both are checkable by reading the trace: the verifier resolves each gated event's nonce to one bound approval and rejects a second event carrying an already-seen (node, ν) pair (`authorisation-nonce-reused`, the `seen_bind` set). The relation is carried by the event syntax, so an external verifier never has to guess which grant a call spent.

Φ_3 is a statement about T. It is deliberately *not* a statement about $\mathcal{A}$: the gated transition spends the grant, so requiring a live grant afterwards would require what the step just destroyed. $\mathcal{A}$ decides what may happen next; T records what did. An outside verifier has only T, and Φ_3 is what it can check.

Let $\Phi := \Phi_1 \wedge \Phi_2 \wedge \Phi_3$.

Base case. C_0 has $T = \text{ENTER}(\mathsf{Start})$ and $\mathcal{A} = \emptyset$. Start is not acting, so Φ_1 asks nothing of that arrival, and there are no ACT events, so all three conjuncts hold.

Preservation. Fix $C \to C'$ by rule R, assume $\Phi(C)$, and write T' for the trace of C'. Every rule appends to T and never removes, so any conjunct that quantifies over T needs checking only for the events R adds. We take the rules in turn.

$R =$ STEP-CTRL, appending $\text{ENTER}(\rho')$ for some $\rho' \in \mathsf{next}_\sigma(\rho)$. The case does not use that first premise, so it is the same whether the move follows a control-flow, calls-subskill, alternative or guard edge or a Decision branch. Φ_2 and Φ_3 quantify over ACT events and none is added, so both carry over unchanged. For Φ_1 the appended event makes this the i-th arrival at ρ' for some i, so the only new obligation is the pair (ρ', i); every other pair (n, j) keeps whatever witness it had, because events persist and $\mathsf{verified}_T(v, n, j, k)$ names a fixed interval of T that T' extends without disturbing. The new obligation is $\mathsf{gates}_{T'}(\rho', i)$: if ρ' is acting, $\mathsf{verified}_{T'}(v, \rho', i, k)$ for every $k \in \mathsf{eff}(\rho').\mathsf{required_kinds}$ and every v of kind k joined to ρ' by a verification-required-before edge; if ρ' is not acting, nothing. That is the rule's second premise, verbatim.

The interval matters here. Events persist, so a premise that asked only for $\text{VERIFY-PASS}(v) \in T$, with Φ_1 over the *set* of nodes entered, would let every later arrival at ρ' inherit the first arrival's witness, a check that ran before the previous arrival. Indexing by arrival and bounding the interval below by $\mathsf{prev}_T(\rho', i)$ rules that out; the executor and the verifier compute that interval.

$R =$ STEP-BIND (App. A). The rule updates σ, writes no event, and leaves $\mathcal{A}$ unchanged. All three conjuncts carry over. (Type safety of the assignment is Definition 1; it is orthogonal to Φ.)

$R =$ STEP-ACT-FREE, appending $e = \text{ACT}(\rho, u, p)$ with $\neg\, \mathsf{approval\text{-}required}(G, \rho, p)$. Φ_1: no ENTER event is added, so no new obligation. Φ_2: the second premise is $\mathsf{permits}(\rho, u)$, which is exactly the obligation the new event carries — for all four acting types, since permits is defined by cases on $\tau_V(\rho)$. Φ_3: for e the antecedent $\mathsf{approval\text{-}required}(G, \rho, p)$ is false, so e carries no obligation. For every earlier e' the obligation is unchanged: this rule writes no APPROVE event, so no earlier call gains a second covering authorisation, and it consumes no decision nonce, so no earlier call's uniqueness clause is violated.

$R =$ STEP-ACT-GATED, appending $\text{APPROVE}(a) \cdot e$ with $e = \text{ACT}(\rho, u, p, a.\nu_{\mathsf{dec}})$ and $\mathsf{approval\text{-}required}(G, \rho, p)$. Φ_1: no ENTER event is added. Φ_2: as above, from $\mathsf{permits}(\rho, u)$. Φ_3 for e, which is the case that matters. Existence: the premise $\text{CONFIRM}(a_0, G.\mathsf{id}, \rho, u, \mathsf{digest}(p), t) \Downarrow a$ returns an a with $d = \mathsf{digest}(p)$, the same contract and node, the tool URI u, a principal, a clock inside $[t_{\mathsf{iss}}, t_{\mathsf{exp}}]$, a_0's decision nonce, and a confirmation over exactly those four fields. So $a \rightharpoonup (\rho, u, p)$, and $\text{APPROVE}(a)$ is written into T' immediately before e. Uniqueness: $\mathsf{consumes}(e, b)$ requires $b.\nu_{\mathsf{dec}} = e.\nu$ *and* $b.\mathsf{node} = \rho$, and STEP-APPROVE's freshness premise forbids issuing the pair $(\rho, \nu_{\mathsf{dec}})$ twice. So at most one such b exists in T', and $\text{APPROVE}(a)$ is it. The argument is about the pair, not the nonce: the same decision may legitimately be issued for a *different* node. No double consumption *at this node*: a_0 was unspent in $\mathcal{A}$ before the step, so no earlier gated call at ρ consumed $(\rho, a_0.\nu_{\mathsf{dec}})$; the conclusion sets $\mathcal{A}[a_0.\mathsf{spent} := \top]$, every later STEP-ACT-GATED requires an unspent pending grant, and STEP-APPROVE's freshness premise forbids reissuing the same pair. This is the non-reuse half, and the case where it bites is a later gated event e' at ρ with the *same* arguments as e: the digests agree, so a is bound to e''s arguments as much as to e's, and only the spent flag and the freshness premise keep a from covering both. e' needs its own pending grant, its own nonce, and its own APPROVE event, so the map from gated events to bound approvals stays injective. Note what is *not* claimed: the decision nonce is not globally unique. The same authority decision may cover a second guarded node, which is why the key is the pair and not the nonce. Earlier calls are untouched: the appended $\text{APPROVE}(a)$ is bound to $\mathsf{digest}(p)$ and sits after them, so it cannot become a second covering event for any earlier e', whose covering approve must precede it.

$R =$ STEP-ACT-CONDITIONAL, appending $\text{APPROVE}(a) \cdot e$ where the gate fired on risk or the guard and no pending grant existed. Φ_1: no ENTER event. Φ_2: from $\mathsf{permits}(\rho, u)$, as above. Φ_3 for e: ASK returns an a already bound to $\mathsf{digest}(p)$ with this contract, node and URI, a principal, a clock inside its window, and the authority's confirmation over those four fields — so $a \rightharpoonup (\rho, u, p)$ — and $\text{APPROVE}(a)$ is written into T' immediately before e, which carries $a.\nu_{\mathsf{dec}}$. Uniqueness is the same argument as the pending case and rests on the same premise: $(\rho, \nu_{\mathsf{dec}})$ is fresh, so no other bound approval in T' consumes it at ρ. Earlier calls are untouched, because the appended approval is bound to *this* digest and sits after them.

The two gated rules are disjoint by their premises — this one requires $\neg$ structural-gate(G, ρ) and no pending grant — so no call is covered twice, and every call for which approval-required(G, ρ, p) holds is covered by one of them.

$R =$ STEP-VERIFY, appending VERIFY-PASS(ρ) or VERIFY-FAIL(ρ). Neither adds an ENTER or an ACT event, so no conjunct gains an obligation and all three carry over. The failing form also sets $\kappa = (\rho, \text{verify-fail})$; Φ does not mention κ, so this changes nothing here, and it is what licenses the next case. Note what VERIFY-PASS does *not* do under the occurrence-local reading: it does not discharge an obligation that already exists. Every verified$_T(v, n, i, k)$ names the fixed interval between two arrivals at n, and a pass appended now falls outside every such interval already closed. It is available to the *next* arrival, which is the intended meaning and was the bug in the old statement. The successful form is stationary (App. A): ρ is unchanged, so the case needs nothing about arrivals, and the next arrival, by STEP-CTRL, carries its own premise.

$R =$ STEP-FALLBACK-LAND, from a configuration with $\kappa = (\rho, c)$, appending FALLBACK-FIRE$(\rho, c) \cdot$ ENTER$(f) \cdot$ STEP(f), where ρ is the node the run is leaving and f is the landing target. The consumed failure (ρ, c) is cleared before f is dispatched; the rule dispatched at f may record a new one, so the post-state carries $\mathcal{A}'$ and κ' as that rule produces them (App. A), and $\kappa' = \bot$ only when it leaves nothing pending. The rule is enabled only by a pending failure — no other rule fires with $\kappa \neq \bot$, and this one cannot fire with $\kappa = \bot$ — so a traversal always follows a recorded failure and never precedes one. This case has real content, because the rule lands and *executes*: f is any executable node, not necessarily a Fallback.

Φ_1: the appended ENTER(f) makes this the i-th arrival at f, and the rule's fourth premise is gates$_{T'}(f, i)$ — the same arrival obligation STEP-CTRL carries, discharged the same way. When $\tau_V(f) =$ Fallback the obligation is vacuous, because a Fallback node is not acting.

Φ_2 and Φ_3: STEP(f) is whichever rule applies at f, and that rule has already been discharged above. If f is acting, the appended ACT event carries the obligations that STEP-ACT-FREE / -GATED / -CONDITIONAL establish; if not, no ACT event is added and both carry over. Nothing here rests on f's type.

Finite prefixes, and what is not claimed. Every configuration in $\mathcal{R}(G)$ is reached by finitely many steps from C_0, and each step appends finitely many events, so every reachable trace is a finite prefix and the induction above ranges over exactly those. The theorem says that Φ holds on every such prefix. It does not say that a run ends. WF-2 and the per-edge retry bound do not give that: they bound only how often a fallback-to edge is traversed, and STEP-APPROVE, STEP-FAIL and the successful STEP-VERIFY are stationary rules whose displayed premises put no bound on how many records one arrival may issue. Two bounds are stated and each is what it is: the per-edge counter retries$_T(\rho, f) <$ MAX-RETRIES-PER-EDGE bounds traversals of that edge and is re-counted by the verifier; fresh$_T(\rho)$ bounds acting events to one per arrival (App. A). Neither is a termination theorem, and Φ needs neither to be one. A global completion or termination claim would need an execution phase with bounded dispatch for every event-producing node, and would be proved then, separately from safety.

Composition is not a rule. $G \oplus_c G'$ is built during negotiation and re-sealed by the provider *before* C_0 (App. A), so no transition rewrites G mid-run and there is no case to discharge here. The induction runs over the composed graph from the start. A Subskill-call node is executed by STEP-ACT-FREE, STEP-ACT-GATED or STEP-ACT-CONDITIONAL with $u =$ `subskill:`$\langle$id$\rangle$, which is what the executor does, so delegation is covered by the three acting cases above rather than by a case of its own.

$R =$ STEP-APPROVE, appending one pending APPROVE(ρ, ν_{dec}) and adding a grant a_n with $d = \bot$ to $\mathcal{A}$ for each $n \in$ protects(ρ). This is the case that needs the bound(T) restriction. Since every $a_n.d = \bot$ and digest$(p) \neq \bot$ for every argument record, the appended event is not in bound(T'): a pending event cannot serve as the covering approve for any call. So the $\exists!$ in Φ_3 neither gains nor loses a witness, and Φ_1, Φ_2 are untouched. Without this restriction a pending approval would count as coverage, and a condition that asked only that *some* APPROVE event exist in T would let one approval at one node satisfy every later gated call.

$R =$ STEP-FAIL, appending r, which is empty or one DENY, GUARD-GATE or POLICY-VIOLATION record, and setting $\kappa = (\rho, c)$. None of these is an ENTER, an ACT or an APPROVE event, and ρ and $\mathcal{A}$ are unchanged, so no conjunct gains an obligation or loses a witness and all three carry over. The κ it sets is what licenses STEP-FALLBACK-LAND, whose case is above.

In every case $\Phi(C')$ holds. By induction Φ holds on all of $\mathcal{R}(G)$, which is Proposition 1. No case used a well-formedness condition, so the theorem holds for every procedure graph. □

What well-formedness adds. Φ_1 ranges over the verification nodes wired to n. If a kind in $\mathsf{eff}(n).\mathsf{required_kinds}$ has no such node, Φ_1 asks nothing for it: a run can reach n and act with Φ_1 intact and no check of that kind in its trace. WF-3 rules this out. Below, a *control path* is a path over the relation $E^{\rightarrow}$ of WF-1 together with the branch targets a Decision names; WF-1 and WF-3 both read this relation (`validator._successors`).

Lemma 3 (Well-formedness makes Φ_1 non-vacuous)**.** *Let G satisfy WF-1 and WF-3, let $C \in \mathcal{R}(G)$ have trace T, and let $e = \text{ACT}(n, u, p)$ be an event of T. Then T has an* ENTER(n) *before e; say the last one is the i-th. For every $k \in \mathsf{eff}(n).\mathsf{required_kinds}$ there is a node v with $\tau_V(v) =$ Verification, $v.\texttt{check_kind} = k$ and $(v, n,$ verification-required-before$) \in E$, and T holds a non-skipped* VERIFY-PASS(v) *after the $(i-1)$-th* ENTER(n) *and before the i-th.*

Proof. The event belongs to an arrival. Only STEP-CTRL and STEP-FALLBACK-LAND change the current node, and each appends the ENTER of the node it moves to (App. A); every other rule, STEP-FAIL and the failing forms included, leaves the current node where it is and appends no ENTER. C_0's trace is ENTER(Start) and its current node is Start. So, by induction over the rules, in every $C \in \mathcal{R}(G)$ the last ENTER event of T names the current node. The event e is written by one of the three acting rules, either on its own or as the STEP(f) that STEP-FALLBACK-LAND dispatches right after it appends ENTER(f). Either way the rule fires at the current node, requires it to be acting and leaves it where it is, so it fired at n with $\tau_V(n) \in$ Acting. Start is not acting, so the last ENTER event before e is an ENTER(n), the i-th for some $i \geq 1$.

A wired verifier of each required kind exists. WF-1 puts n on a control path from Start; take its prefix that ends at n. Since n is acting and $k \in \mathsf{eff}(n).\mathsf{required_kinds}$, WF-3 applies to that prefix and gives a v on it with $\tau_V(v) =$ Verification, $v.\texttt{check_kind} = k$ and $(v, n,$ verification-required-before$) \in E$.

That verifier passed in the interval. By Proposition 1, $\Phi_1(C)$ holds. Its instance at n, the arrival $i \leq \mathsf{arrivals}_T(n)$, the kind k and this v is $\mathsf{verified}_T(v, n, i, k)$: a non-skipped VERIFY-PASS(v) strictly between the $(i-1)$-th and the i-th ENTER(n), and so before e. □

The proof reads WF-3 only for the existence of a wired verifier of each required kind. Its per-path and exit clauses are not needed here, and we prove no progress result from them: a well-formed graph can still abort (the discussion of Φ_1 above). The released code is stricter than gates_T on this point. The executor refuses an arrival at an acting node with a required kind that no wired verifier performs (`no k verifier joined to n`), and `verify_trace` rejects such a trace (`missing-verify-gate`). Both refusals are tested (`test_required_kind_needs_a_joined_verifier.py`), and the mutation audit kills a mutant for each. So the released executor gives the lemma's conclusion on every graph, well-formed or not. On a well-formed graph that refusal never fires, and the model and the code agree.

What a caller gets inside a composed region. Lemma 1 keeps the composed graph well-formed, and Eq. (2) gives the policy it assigns an inlined node. With Proposition 1 and Lemma 3 they bound every mediated run of the composed graph by the caller's policy inside the region.

Theorem 1 is stated in §4; its proof follows.

Proof of Theorem 1. $\text{wf}(H)$ is a hypothesis; Lemma 1 gives it whenever its premises hold, and the corollary also covers a composite that is well-formed without them, as when a caller's verification is relayed onto an inlined node. Lemma 4, whose hypotheses are $\text{wf}(G)$, $\text{wf}(G')$, $\varsigma_G(c) = \bot$ and the fresh renaming the operator enforces, gives $\mathsf{eff}_H(n) = \pi_c \sqcap \mathsf{eff}_{G'}(n)$. Proposition 1 gives $\Phi(C)$. (a) Φ_2 gives $\mathsf{permits}(n, u)$ at $\mathsf{eff}_H(n)$, and permits is monotone in the policy (the consequences of

Lemma 2): what $\pi_c \sqcap \mathsf{eff}_{G'}(n)$ permits, π_c permits. (b) The gate reads $\mathsf{risk}(n) = \mathsf{eff}_H(n).\textsc{risk}$. The meet takes the higher risk, so $\mathsf{eff}_H(n).\textsc{risk} \succeq \pi_c.\textsc{risk}$, which gives the special case. τ_{hitl} is antitone under $\sqcap$, so $\tau_{\mathsf{hitl}}(\mathsf{eff}_H(n)) \preceq \tau_{\mathsf{hitl}}(\pi_c) \preceq \mathsf{risk}(n)$, and the risk clause of $\mathsf{approval\text{-}required}(H, n, p)$ holds. The guards meet disjunctively, $g_{\mathsf{eff}_H(n)} = g_{\pi_c} \vee g_{\mathsf{eff}_{G'}(n)}$, so if g_{π_c} holds on p or cannot be evaluated on it, the guard clause holds. Either way $\mathsf{approval\text{-}required}(H, n, p)$ holds, and Φ_3 gives the approval. (c) The meet unions required kinds, so $k \in \mathsf{eff}_H(n).\mathsf{required_kinds}$. H satisfies WF-1 and WF-3, and Lemma 3 applied to H gives a verifier of kind k wired to n. Φ_1 quantifies over every such verifier, so each passed after the previous arrival at n and before e. □

This is the composition result stated against the caller's policy, not against well-formedness of the result. Like Proposition 1, it is about runs the semantics mediates. It says nothing about which executor a provider ran, which is the trust-profile question (App. D).

What the theorem does not say. It is about the semantics of Appendix A and an executor that follows them. It says nothing about whether a verification predicate is meaningful, whether the tool does what its name suggests, or whether the provider ran this executor at all — that last one is the trust-profile question (App. D), not this one.

C Full Proofs of the Composition Lemmas

This appendix proves Proposition 4, states and proves the inlined effective-policy equation (Lemma 4), and gives the complete case enumeration, over all eleven well-formedness conditions WF-1–WF-11, for Lemma 1 (composition preserves well-formedness) and the selective-disclosure composition result stated and proved in this appendix. The definitions and premises it uses are the ones in `sections/_policy_defs.tex`, which both builds load.

Why Φ_3 binds to the action and its arguments. Φ_3 is stated over an APPROVE *bound to the guarded node and to that action's final arguments, and spent once*, rather than over the mere existence of an earlier APPROVE. The weaker reading is not a simplification, it is unsound for the property the mechanism is supposed to give. Our own executor first implemented it as `any(e.kind == "approve" for e in trace.events)`: one approval anywhere lifted every later value guard, on any node, over any arguments. An external audit produced a well-formed run that was approved while the store held $\mathsf{amount} = 100$ and then fired an irreversible tool with $\mathsf{amount} = 1000$, with the guard never evaluated. That run satisfies the temporal reading – an APPROVE event does precede the TOOL-CALL event – and violates the argument-level authorisation the paper claims. Tool arguments are resolved after the Human-approval node runs, so a grant issued there cannot yet name them; the executor therefore holds it unbound and requires the authority to confirm the concrete arguments before it can be spent. The guard is evaluated on the final arguments at *every* guarded call. `test_approval_binding.py` pins the honest path and five ways a grant can fail to authorise what actually runs: changed arguments, a grant naming another node, a replayed nonce, a refusing authority, and no authority at all.

Setting and notation. Throughout, G and G' are procedure graphs with $\mathrm{wf}(G)$ and $\mathrm{wf}(G')$; $c \in V(G)$ is a Subskill-call node with target G'; π_G and $\pi_{G'}$ are the two root policies; ς and $\widehat{\pi}$ are the scope map and frozen call-site environment of Def. 2; par is the parent relation over the ancestry relation $E^{\uparrow}$; and $H := G \oplus_c G'$. Every policy-valued object below is an element of Π and $\sqcap$ is the meet of Lemma 2 (§3.3 in the long build, App. A.2 in the conference build). The definitions, the premises and the equation of this appendix are the ones the conference method prints; there is no second copy.

The construction. H is a *partial* operation: the released operator returns H only when `validate()` accepts it. Premises (i)–(iv) of Lemma 1, printed again below, are sufficient for that and not necessary: a caller's verification relayed onto an inlined node can discharge WF-3 inside the region where premise (iv) fails, so the operator's domain is larger than the lemma's. Theorem 1 is stated on a well-formed H and covers both. Six steps:

1. **Rename.** Every node of G' is given a fresh identifier, and every id *reference inside a spec* is renamed with it. That is a closed list of typed places, not every string that hap-

pens to equal an id: Decision.branches targets; PortRef.node wherever a PortRef sits (bindings, tool parameters, verification targets, memory-write values); Fallback.handles; Human-approval.fallback_on_timeout; the disclosure coordinate (opaque node ids) of every policy of G' — its root, its overrides and its frozen values; and every region name of G' (its scope labels, the keys of $\widehat{\pi}_{G'}$ and of its nesting map, and the non-$\bot$ parents that map records). Every edge of G' is renamed too. A check kind, a store, a tool URI, a port name, a goal text or a fallback strategy is not an id and is left as written, whatever it says (an earlier operator renamed every equal string; a refuter's callee whose verifier was named after its check kind was then refused at WF-3, and a caller verifier whose kind equalled the call-site id was retyped). *Fresh* means: the renamed node ids and the renamed region names of G' lie outside $V(G) \setminus \{c\}$ and outside every region name of G (the same four sources), none of them is c, the renamed edge ids of G' and the edge ids step 4 and 5 generate lie outside the edge ids of G that survive the hop and outside each other, and the renaming is injective. This is a premise on the renaming, separate from $\mathrm{wf}(G)$ and $\mathrm{wf}(G')$, which speak about one graph each; it is what makes the unions of step 3 functions. Renaming only the ids the graph structure exposes leaves interior references dangling.

2. **Retype.** $\mathsf{Start}_{G'}$ and every $n \in \mathsf{End}(G')$ are retyped to Goal. This is the operation that makes them interior; without it H carries two Start nodes and two exit sets and fails WF-1 at once. Goal is a pass-through with no URI, no side-effect class and no policy of its own, so retyping adds no capability and leaves control-flow shape unchanged.
3. **Scope and transfer.** Every $n \in V(G')$ receives the scope $\varsigma_H(n) = c$ if $\varsigma_{G'}(n) = \bot$ and its renamed inner scope otherwise, and the override

$$\overline{\pi}_H(n) \;=\; \overline{\pi}_{G'}(n) \sqcap \pi_{G'} \sqcap \mathsf{eff}_G(c),$$

with $\mathsf{eff}_G(c)$ computed on G before step 6 deletes c. The frozen environment becomes

$$\widehat{\pi}_H \;=\; \widehat{\pi}_G \;\cup\; \{\, s' \mapsto \widehat{\pi}_{G'}(s) \sqcap \pi_{G'} \sqcap \mathsf{eff}_G(c) \;:\; s \in \mathsf{dom}(\widehat{\pi}_{G'}) \,\} \;\cup\; \{c \mapsto \mathsf{eff}_G(c)\},$$

where s' is s renamed into this namespace, and the root stays $\pi_H = \pi_G$. The nesting map of H is that of G, plus $c \mapsto \varsigma_G(c)$, plus one entry for *every* region name of G' (all four sources, the set step 1 renames and the operator reserves): its recorded parent renamed, or c when G' records none. An earlier operator carried the scope labels and the frozen keys only and dropped the rest, so it reserved a name the output never carried; the two sets are now one. A callee region now sits under the call site, so its boundary value carries the same three-way meet the node transfer applies one line above: the callee's value, the callee's root, and the caller's half. The root factor is needed. Carrying $\widehat{\pi}_{G'}$ across unchanged gives the callee's value alone, and $\widehat{\pi}_{G'}(s) \sqcap \mathsf{eff}_G(c)$ omits $\pi_{G'}$. On a freshly composed G' the two agree, since $\widehat{\pi}_{G'}(s)$ was a meet that contained $\pi_{G'}$; they differ once $\pi_{G'}$ is tightened after G''s own composition, as negotiation does when it seals a forbid into the root, and then the boundary value without the root would sit above the nodes it bounds, and WF-11 (c) at the nested region would not survive this composition (its case below uses the root factor). No effective policy changes: every node whose walk stops at that boundary is a callee node whose own override carries $\pi_{G'}$. This is the step that keeps a call-site obligation alive after the call site is deleted. An external audit of our released operator produced a well-formed G whose call site forbade a tool, carried a HITL trigger and a value guard, composed it with a well-formed G', and found all three absent at an inlined Tool-call — with $\mathrm{wf}(H)$ still holding. `compose.py` now freezes $\mathsf{eff}_G(c)$ at entry and records it under the region's scope, and `test_composition.py` pins every Policy field at the call site, above it, inside the region, at the caller's continuation past it, and across the canonical-JSON round trip that sealing performs. The callee's root policy, allowlist included, rides in the same override; because the walk of Def. 2 stops at a region boundary, that allowlist binds $V(G')$ and no sibling region.
4. **Rewire.** Every executable edge into c — $E^{\rightarrow}$ of WF-1 in full: control-flow, calls-subskill, fallback-to, alternative and guard — is redirected to the retyped $\mathsf{Start}_{G'}$ with its type, guard and specification unchanged; each such edge out of c is reissued from every retyped $\mathsf{End}_{G'}$, with its type, guard and specification unchanged, which also removes the need for a join when G' has several exits. (The released operator listed three of the four types by hand and dropped a guard edge with the call site, so a well-formed caller whose Decision guarded the call was refused at WF-5; the set is now derived from $E^{\rightarrow}$. It then left calls-subskill out as "the call itself". The call itself is not an edge of G — the callee is named by the call site's spec, and no released graph draws a calls-subskill edge inside one graph — and a caller that draws one out of c to its own node y has the executor advance along it and Def. 2 walk it; deleting it with c

widened $\mathsf{eff}_H(y)$ with every graph validating, and made H fail WF-1 when y was reachable only that way, both found by refuters on 2026-09-15. It reissued the out-edges without guard or specification, too, so a fallback-to out of a call site lost the failures it handles; both orders of a chain dropped the same fields, so Theorem 2 was unaffected.) No other edge of $E^{\uparrow}$ leaves c: a gate edge leaves a Verification or a Human-approval node (WF-1), and c is a Subskill-call; a branch target is named by a Decision, which c is not. An obligation edge — requires-approval or verification-required-before — into c is relayed onto every acting node of G', and onto nothing when G' has no acting node: it then has nothing to bind (WF-8 and Φ_3 quantify over acting nodes) and is deleted with c. An earlier operator put it on $\mathsf{Start}_{G'}$ in that case, and the two assembly orders of a chain then put it on different entries (Theorem 2, fixture F5).

5. **Bind.** SUBSKILL-CALL(c).binding and return_binding become data-flow edges. A data-flow edge into c lands on the retyped $\mathsf{Start}_{G'}$; one out of c, and every PortRef of G that reads an output port of c, is re-sourced at the first retyped $\mathsf{End}_{G'}$, which declares every output port of $\Sigma(G')$ (step 2 appends them there, the exit the return_binding edges use), and a port $\Sigma(G')$ does not name is refused. An earlier operator sent such a PortRef to $\mathsf{Start}_{G'}$, which declares no output, and deleted the edge with c: the output validated with a consumer reading a port that did not exist (a refuter's probe, 2026-09-15). Control references to c — a branch target, a Fallback.handles entry — go to $\mathsf{Start}_{G'}$.
6. **Delete.** c and every edge incident to it are removed.

So $V(H) = (V(G) \setminus \{c\}) \cup V(G')$, a disjoint union by freshness, with $V(G')$ renamed and its entry/exit retyped; all other edges of G and G' carry over unchanged, and the only control edges between the two parts enter the region at $\mathsf{Start}_{G'}$ and leave it from an $\mathsf{End}_{G'}$. The operator is implemented in `code/skillgraph/compose.py` and refinement-tested in `test_composition.py`: for every ordered pair of the five released task graphs, and for two-hop chains of them, the composed graph is checked against all eleven conditions. Freshness the operator decides before it writes. It prefixes every callee name with c__, which is deterministic and is not a collision check; `_preflight_names` then checks every name the hop would create — the prefixed node ids, the prefixed region names, the key c, the prefixed edge ids and the edge ids steps 4 and 5 generate (one formula per family, read by the preflight and the writer alike) — against every node id the caller keeps, every region name of the caller and every edge id that survives the hop, and the created sets against each other. A collision is a `CompositionError` naming the name and both sides, and nothing has been written. Edge ids joined the check on 2026-09-15: a callee whose edges were named `bind0`, `bind1` composed as far as the output check and was refused there with a bare "duplicate edge id", after every map was written, although both inputs were well-formed (`EdgeIdsArePreflighted` in `test_compose_domain_closure.py`). An existing region named c is a collision, never an unused key: the composition that made it deleted its call site, and its nodes still read its frozen value. Otherwise a live call named c beside such a region would have $c \mapsto \mathsf{eff}_G(c)$ written over the old value and a retained continuation widened from {`subskill:*`, `tool:allowed`} to $\star$ with every graph validating (`CheckERefusesALiveNodeNamedLikeARegion` and `ThePreflightRefusesAMappedInnerScopeCollision` in `test_wf11_materialization.py`). The operator carries the callee's entries of $\widehat{\pi}_{G'}$ into H, met with $\mathsf{eff}_G(c)$ and renamed. Without them, at a nested region boundary the walk finds no frozen entry and no surviving call site and falls back to the caller's root policy, so composing a callee restricted to one tool into an unrestricted caller returns an unrestricted node with $\mathrm{wf}(H)$ still holding; and without renaming two callees can collide on one region label. `test_scoped_policy_ancestry.py` pins both at depth two, and the mutation audit kills a mutant for each.

Proof of Proposition 4. (1) The two clauses of Def. 2 define a monotone operator on the finite powerset of the finite set $\{\pi\} \cup \overline{\pi}(V) \cup \widehat{\pi}(\mathcal{S})$, so its least fixed point exists. It is the set collected by a depth-first walk backwards from n over $E^{\uparrow}$ that, standing at a node m, enters a parent p iff $\varsigma(p) \in \{\varsigma(m), \bot\}$ and otherwise adds $\widehat{\pi}(\varsigma(p))$; a revisited node is skipped, so a cycle (possible only through fallback-to) changes nothing, and no acyclicity is assumed. `validator.effective_policy` is this walk with memoisation. (2) If $p \in \mathsf{par}(n)$ and $\varsigma(p) \in \{\varsigma(n), \bot\}$, clause (b) gives $Q_G(p) \subseteq Q_G(n)$, and a meet over a superset lies below the meet over the subset; transitivity extends this along any walk of such steps. (3) A walk that starts at n enters only nodes of region $\varsigma(n)$ or of the graph's own region, because every step either stays in the current node's region or moves to a $\bot$-node,

and from a $\bot$-node only $\bot$-nodes are entered. So no node of a region $s \notin \{\varsigma(n), \bot\}$ is visited, no override inside s is collected, and clause (b) contributes exactly $\widehat{\pi}(s)$ at each step that meets s. □

Effective policy after inlining.

Lemma 4 (Inlined effective policy)**.** *Let* $\mathrm{wf}(G)$ *and* $\mathrm{wf}(G')$ *— all eleven conditions,* **WF-11** *in the five-part form of App. A.2, so in particular (d) the current* **Start** *and every* **End** *of each graph is in its own region and (e) no live node id of either graph is one of its region names, and the structural rule of* **WF-1** *that a gate edge leaves a* **Verification** *or* **Human-approval** *node — let* $\varsigma_G(c) = \bot$*, i.e. the call site lies in the graph's own region, and let the renaming of step 1 be fresh. Then*

$$\mathit{eff}_H(n) = \begin{cases} \mathit{eff}_G(c) \sqcap \mathit{eff}_{G'}(n), & n \in V(G'), \\ \mathit{eff}_G(n), & n \in V(G) \setminus \{c\}. \end{cases} \tag{2}$$

Proof. Write $W_X(n)$ for the walk of Def. 2 from n in graph X. Freshness is used throughout in one form: the three key sets step 3 unions — the region names of G, the renamed region names of G', and $\{c\}$ — are pairwise disjoint, so $\widehat{\pi}_H$ and the nesting map are functions, $\widehat{\pi}_H$ agrees with $\widehat{\pi}_G$ on every region name of G, and $\widehat{\pi}_H(c) = \mathsf{eff}_G(c)$. For the key c this is already **WF-11** (e) for G: c is a live node id of G, so it is not a region name of G. Three facts about the construction, and one about its terminals.

(F2) At $\mathsf{Start}_{G'}$, which is in region c of H (its scope in G' is $\bot$ by (d) for G', and step 3 maps $\bot$ to c) and whose parents in H are parents of c in G (those on executable edges and branch targets; obligation edges go to the acting nodes, below), W_H continues into p iff $\varsigma_G(p) \in \{c, \bot\}$, i.e. iff $\varsigma_G(p) = \bot$, because $\varsigma_G(p)$ is $\bot$ or a region name of G and c is not one (by (e) for G); that is exactly when $W_G(c)$ continues into p (as $\varsigma_G(c) = \bot$); otherwise both add $\widehat{\pi}(\varsigma_G(p))$. The same holds at an inlined acting node for an obligation edge that the operator relays from a parent of c: its scope is c or a renamed region name of G', and neither is a region name of G. So everything $W_H(n)$ collects from nodes of G, for any $n \in V(G')$, is an element of $Q_G(c)$ or a frozen value in $Q_G(c)$, and lies above $\mathsf{eff}_G(c)$.

Terminals. The construction keeps the caller's own Start_G and every End_G, with their types and scopes: c is a **Subskill-call**, so it is neither; no step retypes or rescopes a node of G, and step 6 deletes only c. By (d) for G they stay in region $\bot$, so H satisfies (d): its **Start** and **End**s are G's. The callee's $\mathsf{Start}_{G'}$ and every $\mathsf{End}_{G'}$ are retyped to **Goal** (step 2) and, being in region $\bot$ of G' by (d) for G', receive scope c (step 3). Step 4 reissues every executable edge out of c from every retyped $\mathsf{End}_{G'}$, and every edge of $E^{\uparrow}$ out of c is executable: a gate edge cannot leave a **Subskill-call** (its source is typed, **WF-1**) and a branch target is named by a **Decision**. So each successor q of c in $E^{\uparrow}$ — every reconnected caller continuation — has in H a parent in region c where it had c in G; $\varsigma_H(q) = \varsigma_G(q)$ is $\bot$ or a region name of G, and c is neither (by (e) for G), so the walk from q stops at that parent and adds $\widehat{\pi}_H(c) = \mathsf{eff}_G(c)$. That is what (F3) needs. Without the source typing it is false: if the operator reissued only some of the seven edge types of $E^{\uparrow}$ out of c, a caller with a **requires-approval**, a **verification-required-before** or a **calls-subskill** edge from c to a retained node q would have $\mathsf{eff}_H(q)$ lose the call site's forbid, trigger and guard with every graph validating (`GateEdgesLeaveTypedSources` and `CallsSubskillOutOfTheCallSiteIsReissued` in `test_compose_domain_closure.py`). Without (d) for G' it is false: an $\mathsf{End}_{G'}$ authored inside an inner region u of G' keeps its renamed scope u', the walk from q stops there instead and adds $\widehat{\pi}_H(u') = \widehat{\pi}_{G'}(u) \sqcap \pi_{G'} \sqcap \mathsf{eff}_G(c)$, and a continuation whose allow-list was $\star$ read `tool:allowed` with every graph validating. `validate()` refuses that callee at (d), naming each terminal and its region, and its control, the same callee with only **Start** and **End** reset to region $\bot$, composes with the continuation unchanged (`CheckDRefusesScopedTerminals` in `test_wf11_materialization.py`). (F3) Every step of a walk in G that enters c continues (since $\varsigma_G(c) = \bot$), so any node q whose walk reaches c has $Q_G(c) \subseteq Q_G(q)$; in H the same walk reaches an $\mathsf{End}_{G'}$ instead, stops there by the terminals paragraph, and adds $\widehat{\pi}_H(c) = \mathsf{eff}_G(c) = \bigwedge Q_G(c)$.

(F1) Inside the region. Take $n \in V(G')$ with $\varsigma_{G'}(n) = s$ (s may be $\bot$). Every edge of G' is carried over, and ς_H maps $\bot \mapsto c$ and each inner region of G' to its renamed copy. Compare the two walks one parent at a time. At a parent p with $\varsigma_{G'}(p) = s$ both walks enter p. At a parent in another region $u \neq \bot$ of G' both walks stop ($u' \neq s'$ since the renaming is injective, and $u' \neq c$ by freshness):

$W_{G'}$ adds $\widehat{\pi}_{G'}(u)$ and W_H adds $\widehat{\pi}_H(u') = \widehat{\pi}_{G'}(u) \sqcap \pi_{G'} \sqcap \mathsf{eff}_G(c)$ (step 3). At a parent p with $\varsigma_{G'}(p) = \bot$ and $s \neq \bot$ the walks differ, and only here: $W_{G'}$ *enters* p and collects $Q_{G'}(p)$, whose meet is $\mathsf{eff}_{G'}(p)$; in H, $\varsigma_H(p) = c \neq s'$ (freshness again), so W_H *stops* and adds $\widehat{\pi}_H(c) = \mathsf{eff}_G(c)$. (For $s = \bot$ the case does not arise: p and n are both in region c of H and both walks enter p.) The two walks need not collect the same meet: on a well-formed inert chain of exactly this shape, identifying them would drop the callee's restriction.

Write $I(n)$ for the meet of what $W_{G'}(n)$ collects without entering a $\bot$-parent of the region — $\pi_{G'}$, the overrides of the region nodes it visits, and the frozen values of the regions it stops at — and $B(n)$ for the meet of $\mathsf{eff}_{G'}(p)$ over the $\bot$-parents p it enters from a region node ($B(n) = \top$ when $s = \bot$). Then $\mathsf{eff}_{G'}(n) = I(n) \sqcap B(n)$. In H, every override $W_H(n)$ collects is $\overline{\pi}_{G'}(m) \sqcap \pi_{G'} \sqcap \mathsf{eff}_G(c)$ (step 3) and every frozen value of G' it collects is the G' value met with the same two factors; the root π_G and everything collected from nodes of G lie above $\mathsf{eff}_G(c)$ by (F2). So $\mathsf{eff}_H(n) = I(n) \sqcap \pi_{G'} \sqcap \mathsf{eff}_G(c)$, and the tail the truncation drops is exactly $B(n)$.

The dropped tail is harmless when $I(n) \sqsubseteq B(n)$, and WF-11 for G' gives that with no appeal to how G' was built. By (b), $I(n) \sqsubseteq \overline{\pi}_{G'}(n) \sqsubseteq \widehat{\pi}_{G'}(s)$, and $I(n) \sqsubseteq \pi_{G'}$, so $I(n) \sqsubseteq \widehat{\pi}_{G'}(s) \sqcap \pi_{G'}$. Each $\bot$-parent p that $W_{G'}(n)$ enters from a region node m is a boundary parent of s with $\varsigma_{G'}(p) = \bot$, so by (c), $\widehat{\pi}_{G'}(s) \sqcap \pi_{G'} \sqsubseteq \kappa(p) = \mathsf{eff}_{G'}(p)$. Hence $I(n) \sqsubseteq B(n)$, and

$$\mathsf{eff}_H(n) \;=\; I(n) \sqcap \pi_{G'} \sqcap \mathsf{eff}_G(c) \;=\; I(n) \sqcap B(n) \sqcap \mathsf{eff}_G(c) \;=\; \mathsf{eff}_{G'}(n) \sqcap \mathsf{eff}_G(c),$$

because a meet does not move when elements above it are dropped, and $\pi_{G'}$ is already in $I(n)$. This is the first case of the lemma. Neither half of WF-11 is decoration. Without (c) the equation is false: the reviewer's graph has $\widehat{\pi}(d) = \top$ and $\mathsf{eff}(p) = \{\texttt{tool:allowed}\}$ at the boundary. Without (b) it is false too: widen the override of the region's entry to $\top$ under the same ancestor, and $\mathsf{eff}_{G'}$ still reads the restriction through p while $I(n)$ has climbed above $B(n)$. `validate()` refuses both (`test_wf11_materialization.py`).

Case $n \in V(G) \setminus \{c\}$. $W_H(n)$ and $W_G(n)$ agree on every step that does not touch c or the region: the nodes of G keep their overrides and scopes, and $\widehat{\pi}_H$ agrees with $\widehat{\pi}_G$ on every region name of G (freshness, and (e) for G for the key c — the step that probe B broke). Where $W_G(n)$ enters c, by (F3) it collects $Q_G(c) \subseteq Q_G(n)$, and $W_H(n)$ collects $\mathsf{eff}_G(c) = \bigwedge Q_G(c)$ instead. A meet is unchanged when a subset is replaced by its own meet, so $\bigwedge Q_H(n) = \bigwedge Q_G(n)$. □ □

The second case needs the scoped statement. Stated over *control-flow* ancestry it is false: inlining deletes c, so after composition $\mathsf{Start}_A \to \mathsf{Start}_B \to \mathsf{action}_B \to \mathsf{End}_B \to \mathsf{continuation}_A$, and action_B reaches $\mathsf{continuation}_A$ over ordinary control flow. What is preserved is the scoped contribution set: the continuation's walk stops at End_B and takes $\widehat{\pi}(c)$, never action_B's override. That is (F3), and it is also why the return boundary does not contradict Proposition 4(2): the step $\mathsf{End}_B \to \mathsf{continuation}_A$ is a boundary crossing, not an inclusion of contribution sets. The implementation always did this (`Node.compose_scope`, `SkillGraph.compose_scope_policy`, consulted in `validator.effective_policy`); it was the printed definition that was flat. Both fields are part of the sealed graph, so a consumer checking a commitment sees the same scoping the executor uses. The hypothesis $\varsigma_G(c) = \bot$ covers every one-hop composition whose call site lies in the caller's own region (every uncomposed caller, and the inside-out nesting $G \oplus_c (G' \oplus_{c'} G'')$); it is a premise of Lemma 1, and of the selective-disclosure result below and of the chain statement after Theorem 1 at every hop, not only of this equation, because their WF-3, WF-7 and WF-8 cases read it. The outside-in order places its second call site $c__c'$ inside region c, outside this hypothesis, and this lemma says nothing about that hop. Theorem 2 (App. C.2) covers it on the admitted pattern grammar, provided both orders compose: the two orders then give the same nodes, scopes and edge relation, the same frozen environment as values of Π, and the same eff at every node as a value of Π — not every serialised field of the implementation's policy objects — which `test_nested_delegation.py` checks on a hand-built three-graph chain and `check_composition_order.py` on random chains; sealed bytes are not claimed equal, and that the outside-in order composes at all is not claimed either — it is a hypothesis of the theorem, not a consequence of the inside-out order composing.

The four premises. We use premises (i)–(iv) of Lemma 1, printed here from `sections/_policy_defs.tex`, the same source as the lemma: Throughout, $\mathrm{wf}(G)$ and $\mathrm{wf}(G')$ mean all eleven conditions, WF-11 in its five-part form, and the call site lies in the

caller's own region, $\varsigma_G(c) = \bot$. That scope hypothesis is the hypothesis of Eq. (2); the WF-3, WF-7 and WF-8 cases of the proof read the equation, so it is a premise of the lemma itself and not only of the equation printed beside it. Write $\pi_n := \mathsf{eff}_G(c) \sqcap \mathsf{eff}_{G'}(n)$ for $n \in V(G')$, the combined policy the inlined node will carry.

(i) *(type compatibility)* the bindings on c are well typed: Δ holds for each binding edge, c's inputs against $\Sigma(\mathsf{Start}_{G'}).\mathsf{in}$ and $\Sigma(\mathsf{End}_{G'}).\mathsf{out}$ against c's outputs;

(ii) *(permission compatibility)* every acting node $n \in V(G')$ is permitted under π_n, that is permits (Def. 3) evaluated at π_n: a Tool-call or Subskill-call with URI u has $u \models \pi_n.\textsc{allowed}$ and $u \not\models \pi_n.\textsc{forbidden}$; a Memory-write to store s has $s \in \mathcal{L}(\pi_n.\textsc{writable})$; an Artifact has no condition. Because $\mathrm{wf}(G')$ already supplies the $\mathsf{eff}_{G'}(n)$ half and the meet intersects allow-lists and unions forbid-lists, this is the same three tests against $\mathsf{eff}_G(c)$ alone;

(iii) *(approval compatibility)* for every acting node $n \in V(G')$ with $\pi_n.\textsc{risk} \succeq \tau_{\mathsf{hitl}}(\pi_n)$ — the combined risk against the combined threshold — every control path from $\mathsf{Start}_{G'}$ to n in G' traverses, and leaves by an edge other than fallback-to, some h with $\tau_V(h) = \mathsf{Human\text{-}approval}$ and $(h, n, \mathsf{requires\text{-}approval}) \in E(G')$;

(iv) *(verification coverage and splice admissibility)* for every acting node $n \in V(G')$ and every $k \in \pi_n.\textsc{required-kinds}$, every control path from $\mathsf{Start}_{G'}$ to n in G' satisfies k in the sense of WF-3. For the kinds $\mathsf{eff}_{G'}(n)$ requires this is WF-3 for G'. For the kinds the consumer adds through $\mathsf{eff}_G(c)$ it is met by the Verification nodes that NEGOTIATE-VERIFICATION splices into G' before composition, and each spliced v satisfies: (a) $v.\texttt{check_kind} \in \{\mathsf{json\text{-}schema}, \mathsf{unit\text{-}test}, \mathsf{static}, \mathsf{llm\text{-}judge}\}$ and $v.\texttt{on_fail} \in \{\textsc{fallback}, \textsc{abort}\}$; (b) the insertion point has an incoming control-flow edge, or G' has a Start to hang the chain from; (c) every target joined to v by verification-required-before is a Tool-call or Artifact node, never a Fallback node (WF-10); (d) v dominates every such target.

Every premise is decided from G, c and G': $\varsigma_G(c)$ is a sealed field of G, $\mathsf{eff}_G(c)$ is computed on G, $\mathsf{eff}_{G'}(n)$ on G', and the path conditions on G'. None refers to $\mathrm{wf}(G \oplus_c G')$. The released operator decides (i) against the callee's signature, the Tool-call clause of (ii) directly at the call site (`_check_caller_admits_callee`), and the remaining clauses by validating H and raising if any condition fails. That makes the operator partial with a decidable domain, a domain that contains the premises' and can be larger (above); it is not the proof. The proof below uses the four premises as stated and nothing about the output.

C.1 LEMMA 1: FULL CASE ENUMERATION

WF-1 (unique entry/exit, reachability). H has the same unique Start as G ($c \neq \mathsf{Start}_G$ because Start nodes have type Start, not Subskill-call) and the same End set as G, *because* step 2 of the construction retypes $\mathsf{Start}_{G'}$ and every $\mathsf{End}_{G'}$ to Goal. They are interior by construction, not by assumption: carrying their types over unchanged would give H two entry nodes and two exit sets. Reachability: every $n \in V(G) \setminus \{c\}$ lies on a Start-to-End path in G; if that path passed through c, replace the single step through c by a path $\mathsf{Start}_{G'} \rightsquigarrow \mathsf{End}_{G'}$, which exists by $\mathrm{wf}(G')$ (WF-1 for G'). The two edges of that step are in $E^{\rightarrow}$, and step 4 rewires every type of $E^{\rightarrow}$, so the replacement path is in H; an operator that left calls-subskill out refused its own well-formed inputs here ("node y not reachable from Start", a refuter's probe of 2026-09-15). Every $n \in V(G')$ is on a $\mathsf{Start}_{G'} \rightsquigarrow \mathsf{End}_{G'}$ path by WF-1 for G'; prefixing the G-path from Start_G to c and suffixing the G-path from c to an End of G embeds it in an H-path. □

WF-2 (control acyclicity). Suppose H had a control cycle not passing through a fallback-to edge. Project the cycle onto $V(G)$ by contracting the inlined G' back to the single vertex c: edges internal to G' project to self-loops on c only if the cycle is internal to G', contradicting WF-2 for G'; otherwise the projected walk is a control cycle of G, contradicting WF-2 for G. □

WF-3 (verification coverage). The condition is path-wise and per kind, and so is this proof; no single dominating node is claimed, and none is needed. Two facts about paths in H are used throughout. *Contraction:* a control path P of H from Start to a caller node visits the region zero or more times, each visit entering at $\mathsf{Start}_{G'}$ and leaving from an $\mathsf{End}_{G'}$ by a reissued edge; replacing each visit by the single step through c gives a control path P/c of G through the same caller nodes,

and the edge by which P leaves a caller node v is the edge by which P/c leaves v, of the same type (step 4 preserves types). *Suffix:* a control path of H to $n \in V(G')$ enters the region only at $\mathsf{Start}_{G'}$, and its suffix after the last visit to $\mathsf{Start}_{G'}$ uses only edges of G', so it is a control path of G' from $\mathsf{Start}_{G'}$ to n.

Let $n \in V(H)$ with $\tau_V(n) \in$ Acting (all four types, Def. 3), let $k \in \mathsf{eff}(n).\mathsf{required_kinds}$ in H, and let P be any control path of H from Start to n. *Case* $n \in V(G) \setminus \{c\}$*:* $\mathsf{eff}_H(n) = \mathsf{eff}_G(n)$ by Eq. (2), so k is required in G; WF-3 for G applied to P/c gives a v on P/c with $\tau_V(v) =$ Verification, $v.\mathtt{check_kind} = k$, $(v, n,$ verification-required-before$) \in E(G) \subseteq E(H)$, left by a non-fallback-to edge; $v \neq c$ since c is a Subskill-call, so v lies on P and P leaves it by an edge of the same type. *Case* $n \in V(G')$*:* by Eq. (2), $\mathsf{eff}(n).\mathsf{required_kinds}$ in H is $\pi_n.\textsc{required-kinds}$, the union of the kinds required by $\mathsf{eff}_{G'}(n)$ and by $\mathsf{eff}_G(c)$; premise (iv) says every control path of G' from $\mathsf{Start}_{G'}$ to n satisfies each such k, and the suffix of P is one. Its verifier v, its edge (v, n) and its exit edge are edges of G' carried into H. □

WF-4 (fallback completeness). Side-effect classes are node-local and unchanged by inlining. For $n \in V(G) \setminus \{c\}$ the fallback edge or on_fail $=$ ABORT declaration of G is carried over; for $n \in V(G')$, that of G' is carried over. The substitution at c removes c itself, the only node whose WF-4 obligation could have changed, and c is not a side-effecting node (its type is Subskill-call). □

WF-5 (decision coverage). Decision nodes and their outgoing edge sets are carried over from G or G', with one change: a branch of a Decision of G that named c now names $\mathsf{Start}_{G'}$, in the branch map and in the redirected edge alike (step 4 redirects every executable type, guard included), so each declared label keeps a matching edge of an admitted type and the default-or-guard totality condition is unchanged. Inlining adds no Decision node and no guard edge. □

WF-6 (data-flow typing). Edges internal to G or G' keep their endpoints and port types, so Δ holds by $\mathrm{wf}(G)$, $\mathrm{wf}(G')$. The only new data-bearing edges are the binding edges realising $\textsc{Subskill-call}(c)$.binding (caller ports $\to \Sigma(\mathsf{Start}_{G'}).\mathsf{in}$) and return_binding ($\Sigma(\mathsf{End}_{G'}).\mathsf{out} \to$ caller ports); premise (i) states exactly that these are well-typed under $\sqsubseteq$. □

WF-7 (action permission). For $n \in V(G) \setminus \{c\}$ the effective policy is unchanged, so WF-7 holds by $\mathrm{wf}(G)$. For an acting $n \in V(G')$, $\mathsf{eff}_H(n) = \pi_n$ by Eq. (2), and WF-7 at n is $\mathsf{permits}(n, \mathsf{uri}(n))$ evaluated at π_n, one case per acting type. *Tool-call and Subskill-call* with URI u: needs $u \models \pi_n.\textsc{allowed}$ and $u \not\models \pi_n.\textsc{forbidden}$, which is premise (ii) verbatim. *Memory-write* to store s: needs $s \in \mathcal{L}(\pi_n.\textsc{writable})$, premise (ii) verbatim. *Artifact*: permits is $\top$. No case is inferred from tool compatibility, and no case reads the validator's verdict on H. □

WF-8 (approval gating). For $n \in V(G) \setminus \{c\}$: $\mathsf{eff}_H(n) = \mathsf{eff}_G(n)$, so the gating condition is unchanged, and the contraction argument of WF-3 carries the approval node h that WF-8 for G supplies on P/c onto P, with the same exit edge ($h \neq c$). For an acting $n \in V(G')$: the gating condition in H is $\pi_n.\textsc{risk} \succeq \tau_{\mathsf{hitl}}(\pi_n)$ by Eq. (2) — the combined risk (the higher of the two) against the combined threshold (High if either policy declares triggers) — which is precisely the antecedent of premise (iii); (iii) supplies, on every control path of G' from $\mathsf{Start}_{G'}$ to n, an h with a requires-approval edge to n left by a non-fallback edge, and the suffix of any H-path to n is such a path. □

WF-9 (subskill type compatibility). All calls-subskill edges of G other than the one at c (which is consumed by the inlining) and all such edges of G' are carried over with unchanged bindings, so WF-9 holds by $\mathrm{wf}(G)$, $\mathrm{wf}(G')$. □

WF-10 (fallback insulation). Fallback nodes and verification-required-before edges are carried over from G and G', where WF-10 holds. The only new verification-required-before edges are the splices of premise (iv), which by clause (c) target non-Fallback nodes. □

WF-11 (regions composed and current, terminals root-scoped, names disjoint). Five parts, from WF-11 for G and G', the freshness of the renaming, and the values step 3 writes; nothing about how G or G' came to be built is used. Throughout, $\kappa_X(p)$ is as in App. A.2: $\mathsf{eff}_X(p)$ at a boundary

parent in $\bot$ or in a region enclosing the one checked, the frozen value of p's region otherwise (what the walk takes at the first kind of parent, and stronger than what it takes at the second).

(a) Records. Every region label in H is one of G's, one of G''s renamed into this namespace, or the new label c. Step 3 records a frozen value for each — $\widehat{\pi}_G$ kept, $\widehat{\pi}_{G'}(s) \sqcap \pi_{G'} \sqcap \mathsf{eff}_G(c)$ at s', $\mathsf{eff}_G(c)$ at c — and a nesting: c sits in $\varsigma_G(c) = \bot$, a carried region sits in its renamed parent or in c, a region of G keeps its parent. Every chain resolves to $\bot$, through (a) for G and G' and the new link $c \mapsto \bot$.

(b) Region c binds inside. For $n \in V(G')$ with $\varsigma_H(n) = c$, $\overline{\pi}_H(n) = \overline{\pi}_{G'}(n) \sqcap \pi_{G'} \sqcap \mathsf{eff}_G(c) \sqsubseteq \mathsf{eff}_G(c) = \widehat{\pi}_H(c)$. *(c) Region c has a faithful boundary.* A boundary parent p of region c is either a node of a carried region — nested in c, so excluded from (c) — or a parent of c in G: step 4 moves c's incoming control edges to $\mathsf{Start}_{G'}$, and the operator relays an obligation edge into c onto the inlined acting nodes. No region of G encloses c (its parent is $\bot$), so either $\varsigma_G(p) = \bot$ and $\kappa_H(p) = \mathsf{eff}_H(p) = \mathsf{eff}_G(p)$ (second case of Lemma 4) with $Q_G(p) \subseteq Q_G(c)$, or $\varsigma_G(p) = u \neq \bot$ and $\kappa_H(p) = \widehat{\pi}_H(u) = \widehat{\pi}_G(u) \in Q_G(c)$. Either way $\widehat{\pi}_H(c) \sqcap \pi_H = \mathsf{eff}_G(c) \sqcap \pi_G = \mathsf{eff}_G(c) \sqsubseteq \kappa_H(p)$.

(b) A carried region s' binds inside. For n with $\varsigma_H(n) = s'$, $\overline{\pi}_H(n) = \overline{\pi}_{G'}(n) \sqcap \pi_{G'} \sqcap \mathsf{eff}_G(c) \sqsubseteq \widehat{\pi}_{G'}(s) \sqcap \pi_{G'} \sqcap \mathsf{eff}_G(c) = \widehat{\pi}_H(s')$, by (b) for G' at s. *(c) A carried region s' has a faithful boundary.* Let p be a boundary parent of s' in H with $\varsigma_H(p)$ not nested in s'. If $p \in V(G)$, it reaches the region by a relayed obligation edge, $\kappa_H(p)$ is $\mathsf{eff}_G(p)$ or a frozen value of G, both above $\mathsf{eff}_G(c)$ as in the previous paragraph, and $\widehat{\pi}_H(s') \sqcap \pi_G \sqsubseteq \mathsf{eff}_G(c)$. Otherwise $p \in V(G')$ is a boundary parent of s in G', and $\kappa_H(p) \sqsupseteq \kappa_{G'}(p) \sqcap \pi_{G'} \sqcap \mathsf{eff}_G(c)$: when $\varsigma_{G'}(p) = \bot$, p is now in region c, which encloses s', and $\kappa_H(p) = \mathsf{eff}_H(p) = \mathsf{eff}_G(c) \sqcap \mathsf{eff}_{G'}(p) = \mathsf{eff}_G(c) \sqcap \kappa_{G'}(p)$ (first case of Lemma 4; $\mathsf{eff}_{G'}(p) \sqsubseteq \pi_{G'}$); when $\varsigma_{G'}(p) = u$ encloses s, u' encloses s' and the same equation holds; when u is neither, $\kappa_H(p) = \widehat{\pi}_H(u') = \widehat{\pi}_{G'}(u) \sqcap \pi_{G'} \sqcap \mathsf{eff}_G(c) = \kappa_{G'}(p) \sqcap \pi_{G'} \sqcap \mathsf{eff}_G(c)$. Now (c) for G' at s gives $\widehat{\pi}_{G'}(s) \sqcap \pi_{G'} \sqsubseteq \kappa_{G'}(p)$, hence

$$\widehat{\pi}_H(s') \sqcap \pi_H \;\sqsubseteq\; \widehat{\pi}_{G'}(s) \sqcap \pi_{G'} \sqcap \mathsf{eff}_G(c) \;\sqsubseteq\; \kappa_{G'}(p) \sqcap \pi_{G'} \sqcap \mathsf{eff}_G(c) \;\sqsubseteq\; \kappa_H(p).$$

The factor $\pi_{G'}$ in $\widehat{\pi}_H(s')$ is what makes the first step true when $\pi_{G'}$ was tightened after s was composed: (c) for G' holds through its own $\sqcap \pi_{G'}$, and H inherits that only if the boundary value carries the root. *Regions of G.* Their nodes, overrides, frozen values, boundary parents and root are unchanged from G (freshness: no entry of $\widehat{\pi}_G$ or of the nesting map is overwritten), except that where c was a boundary parent (an edge out of c into a region of G), the reissued edge makes an $\mathsf{End}_{G'}$ in region c the boundary parent instead, with $\kappa_H = \widehat{\pi}_H(c) = \mathsf{eff}_G(c) = \kappa_G(c)$ — so (b) and (c) for G carry over. *(d) Terminals.* $\mathsf{Start}_H = \mathsf{Start}_G$ and $\mathsf{End}(H) = \mathsf{End}(G)$, kept with their scopes, which are $\bot$ by (d) for G; the callee's terminals are Goals in H (step 2), which (d) does not read. This is the terminals paragraph of Lemma 4. *(e) Names.* The live node ids of H are $V(G) \setminus \{c\}$ and the renamed $V(G')$; the region names of H are among those of G, those of G' renamed, and c (the scopes step 3 writes are c or renamed scopes of G'; the parent it records for c is $\varsigma_G(c) = \bot$; a carried region's parent is a renamed region name of G' or c). Six pairs. $V(G) \setminus \{c\}$ against the region names of G: (e) for G. The renamed $V(G')$ against the renamed region names of G': (e) for G', the renaming being injective. $V(G) \setminus \{c\}$ against c: distinct ids. The remaining three — $V(G) \setminus \{c\}$ against the renamed region names of G', the renamed $V(G')$ against the region names of G, and the renamed $V(G')$ against c — are freshness, which the construction assumes and the operator decides before it writes: `_preflight_names` checks these three, the rest of freshness (the renamed node ids of G' against $V(G) \setminus \{c\}$, and the renamed region names of G' and c against the region names of G), and, although (e) for G' already gives it, the renamed node ids against the renamed region names. □

All eleven conditions hold, so $\mathrm{wf}(H)$. ■

Four instances, checked by hand against the definitions. *(1) Nested composition* $H = G \oplus_c (G' \oplus_{c'} G'')$, both call sites in their graph's own region. Lemma 4 at the inner hop gives, for $n \in V(G'')$, $\mathsf{eff}_{G' \oplus_{c'} G''}(n) = \mathsf{eff}_{G'}(c') \sqcap \mathsf{eff}_{G''}(n)$, and at the outer hop $\mathsf{eff}_H(n) = \mathsf{eff}_G(c) \sqcap \mathsf{eff}_{G'}(c') \sqcap \mathsf{eff}_{G''}(n)$. For the G'-node q that follows c', the walk stops at $\mathsf{End}_{G''}$ and takes $\widehat{\pi}(c') = \mathsf{eff}_{G'}(c')$: $\mathsf{eff}_H(q) = \mathsf{eff}_G(c) \sqcap \mathsf{eff}_{G'}(q)$, with no override of G'' in it. *(2) Two sequential callees* c_1, c_2 in G, where callee 1's root forbids tool t and callee 2 calls t at a node n with t admitted by $\mathsf{eff}_G(c_2)$. Inside region c_1 every node has t in its forbid set, so callee 1 may not call t (WF-7 for the callee). At n the walk stops at End_{c_1} and takes $\widehat{\pi}(c_1) = \mathsf{eff}_G(c_1)$, the *caller's* policy at c_1, in which callee 1's forbid does not appear; so $t \notin \mathsf{eff}_H(n).\textsc{forbidden}$ and WF-7 holds at n exactly when premise (ii) holds at c_2. Callee 1's local restriction ended at its return. *(3) A restricted memory write.* The call site

has WRITABLE $= \{\texttt{scratch}\}$ and the callee, with WRITABLE $= \{\star\}$, writes to `prod` at n. Then $\mathcal{L}(\pi_n.\text{WRITABLE}) = \{\texttt{scratch}\}$, $\texttt{prod} \notin$ it, premise (ii) fails at n, and the composition is refused; the same callee writing to `scratch` is admitted. This is the case the tool-only precondition could not see, and the one that once produced an unrestricted composed policy, an executed write and a verified audit. *(4) An opaque output at a concrete type* is instance (4) of the next subsection.

C.2 Composing a chain in either order

Instance (1) above is one chain, checked by hand. The claim the paper leans on is general: a consumer hires a provider, the provider subcontracts, and what the consumer is guaranteed must not depend on the order somebody assembled the graph in. Otherwise the induction along a chain is not available.

What is compared, and which equality. Two composed graphs are compared as the data that Def. 2 and the eleven conditions read: node identifiers, types, specifications, ports, scopes and overrides; the root; the frozen environment and the nesting map; and the edge relation. Policies are compared *as values of* Π: its seven coordinates, not every field of the implementation's `Policy` object, which also serialises a memory TTL and a disclosure list. Two pattern sets are equal when their canonical forms are — drop subsumed entries, deduplicate, sort; equal iff the languages are (App. A.2) — and two guards when their canonical disjunctions are. The checker below decides this equality by name (`pi_projection`, the seven coordinates in canonical form, applied by `pi_differences` to eff at every node and to the root, every override and every frozen value, each reported by where it differs) and, beside it, a stronger one: the implementation's `validator._policy_key` compares the two auxiliary fields as well, and the full serialisation is compared field by field. The projection applies to eff, the overrides, the frozen values and the root alike. A difference the stronger comparison reports is a diagnostic; the theorem's claim is the projected equality. It is language equality on the pattern coordinates, not equality of lists as written: $\{\star\}$ and $\{\star, \texttt{tool:a}\}$ are the same value. Sealed bytes are compared too, but only *reported*; the theorem does not claim them equal, and they are not.

Proposition 5 (Either order is sound)**.** *Let G, G', G'' be well-formed, let $c \in V(G)$ target G' with $\varsigma_G(c) = \bot$ and $c' \in V(G')$ target G'', and let both orders compose. Then in either order, for every $z \in V(G'')$,*

$$\mathit{eff}_H(c __ c' __ z) \;\sqsubseteq\; \mathit{eff}_G(c) \sqcap \mathit{eff}_{G'}(c') \sqcap \pi_{G''} \sqcap \overline{\pi}_{G''}(z),$$

and for every $y \in V(G') \setminus \{c'\}$, $\mathit{eff}_H(c __ y) \sqsubseteq \mathit{eff}_G(c) \sqcap \pi_{G'} \sqcap \overline{\pi}_{G'}(y)$. So neither order admits a tool, a store or an unguarded call that the contributing policies jointly refuse, and Φ holds whichever order was used.

Proof. Clause (a) of Def. 2 puts a node's own override in its contribution set, so $\mathsf{eff}_H(n) \sqsubseteq \overline{\pi}_H(n)$. Step 3 of the construction writes the override of an inlined node as its old override met with the callee's root and the caller's policy at the call site. Inside out, z's override is $\overline{\pi}_{G''}(z) \sqcap \pi_{G''} \sqcap \mathsf{eff}_{G'}(c')$ after the inner hop, met with $\pi_{G'} \sqcap \mathsf{eff}_G(c)$ at the outer. Outside in it is $\overline{\pi}_{G''}(z) \sqcap \pi_{G''} \sqcap \mathsf{eff}_{G \oplus_c G'}(\overline{c})$, and $\mathsf{eff}_{G \oplus_c G'}(\overline{c}) = \mathsf{eff}_G(c) \sqcap \mathsf{eff}_{G'}(c')$ by Lemma 4 at the first hop, whose premise $\varsigma_G(c) = \bot$ holds. Both are below the stated bound because $\sqcap$ is a lower bound (Lemma 2); the G' case is one hop. Nothing here uses that the meet is *exact*: the bound holds for any lower bound of the contributions. That is why this proposition is stated on its own although it now follows from Theorem 2: it is what remained true while the allow-list meet was an approximation (below), and what would remain true of a coordinate whose meet is only a lower bound. What the orders could then differ on is how much *more* they restrict. □ □

Theorem 2 (Order independence on the admitted grammar, as canonical policies)**.** *Let G, G', G'' be well-formed — all eleven conditions, WF-11 in the five-part form of App. A.2; the three may carry regions from earlier compositions — with every pattern list in the admitted grammar (`validate()` rejects any other shape as a structural error). Let $c \in V(G)$ be a Subskill-call targeting G' with $\varsigma_G(c) = \bot$, let $c' \in V(G')$ target G'' with $\varsigma_{G'}(c') = \bot$, and let premises (i)–(iv) of Lemma 1 hold at both hops of the inside-out order, so that H_{in} is well-formed by the lemma applied twice. Suppose the outside-in order also composes, i.e. the operator refuses at neither of its hops. Write $\overline{c} = c __ c'$ and*

$$H_{\text{in}} = G \oplus_c (G' \oplus_{c'} G''), \qquad H_{\text{out}} = (G \oplus_c G') \oplus_{\overline{c}} G''.$$

Then (a) $V(H_{\text{in}}) = V(H_{\text{out}})$, and each node has the same type, specification, ports, disclosure witness and scope in both; (b) at every node the override, at every region label the frozen value, and

at the root the policy are equal as elements of Π *— equal canonical forms — and the nesting map is the same; (c) the edge relation* $\{(\mathsf{src}, \mathsf{dst}, \mathsf{type}, \mathsf{guard}, \mathsf{spec})\}$ *is the same. Hence* $\mathsf{eff}_{H_{\mathrm{in}}}(n) = \mathsf{eff}_{H_{\mathrm{out}}}(n)$ *in* Π *for every* n*;* $\mathrm{wf}(H_{\mathrm{in}}) \Leftrightarrow \mathrm{wf}(H_{\mathrm{out}})$*, so* H_{out} *is well-formed although its second hop is at* $\overline{c}$*, which lies in region* c *and is outside the caller-root premise of Lemma 4; and* Φ *holds on the same trajectories.* Not claimed: *equality of sealed bytes, or of every serialised field. Equality of policy values in (b) is after projection to the seven coordinates of* Π*, together with the node, scope and edge-relation components of (a) and (c); it does not assert equality of every serialised implementation field. The implementation's* `Policy` *carries two fields outside* Π*, a memory TTL and a disclosure list. Generated edge identifiers differ between the orders, and an auxiliary disclosure reference to a deleted call site is not currently guaranteed to agree: with* `opaque_nodes` $= \{c'\}$ *on* G'*'s root, the inside-out order deletes* c' *before the outer hop renames, and the reference stays* c'*, while the outside-in order renames it to* $\overline{c}$ *first (fixture F7 of the checker, an accepted input; both outputs validate;* Π *agrees at every node). The checker reports those additional differences separately from the projected equality. The* wf *equivalence above is over the model's eleven conditions, which read* Π*; the implementation's* **WF-11** *also compares the two auxiliary fields (a TTL bound and an opaque-set inclusion). On F7 and every input checked, their difference never changed a verdict — each order's inclusion test compares a list with the same renaming of itself — but that is observed, not proved here.*

Proof. The meet on canonical forms. Two admitted patterns are nested or disjoint: $\star$ contains every pattern; a literal lies inside $p\star$ iff it begins with p and otherwise misses it; $p\star$ and $q\star$ are nested iff one prefix begins with the other, and are disjoint otherwise, since a string beginning with both would make one prefix a prefix of the other. So for pattern sets A, B, with $\mathcal{I}$ keeping the narrower of each nested pair and nothing for a disjoint pair, $\mathcal{L}(\mathcal{I}(A,B)) = \bigcup_{p\in A,\, q\in B} \mathcal{L}(p) \cap \mathcal{L}(q) = \mathcal{L}(A) \cap \mathcal{L}(B)$: the intersection is exact. Intersection of languages is associative, commutative and idempotent; canonical forms are equal iff languages are; and $\mathcal{I}$ returns a canonical form. So on canonical forms $\mathcal{I}$ has all three properties, and $\mathcal{L}(A) \subseteq \mathcal{L}(B)$ gives $\mathcal{I}(A,B) = A$. The forbid coordinate is the canonical union, with the same properties; the remaining coordinates are finite sets under union, total orders under $\max$, and the guard under canonical disjunction. Hence $\sqcap$ computed on canonical representatives is associative, commutative, idempotent and absorbing, which is what the two bracketings below need. The released meet did not have this before the grammar was fixed; see the next paragraph.

Node identifiers. Inlining at a call site d prefixes the callee with $d__$. Inside out, $z \in V(G'')$ becomes $c'__z$ and then $c__c'__z$. Outside in, c' becomes $\overline{c}$, and inlining at $\overline{c}$ prefixes z with $\overline{c}__$, which is the same string. Nodes of G' become $c__y$ either way and nodes of G are untouched.

Types, specifications, ports. Each callee's Start and End are retyped to Goal exactly once, at the hop that inlines that callee, in both orders, with the same goal text (it names the callee graph, and $G' \oplus_{c'} G''$ keeps G''s identifier); the exit's ports are extended by the same signature at that hop. Identifier references inside specifications are renamed by the same maps as the identifiers. Disclosure witnesses are carried verbatim.

Overrides. Let $z \in V(G'')$. Inside out, the inner hop gives $\overline{\pi}_{G''}(z) \sqcap \pi_{G''} \sqcap \mathsf{eff}_{G'}(c')$ and the outer hop meets in $\pi_{G'\oplus G''} \sqcap \mathsf{eff}_G(c)$; the root of a composed graph is the caller's, so $\pi_{G'\oplus G''} = \pi_{G'}$ and the total is

$$\overline{\pi}_{G''}(z) \sqcap \pi_{G''} \sqcap \mathsf{eff}_{G'}(c') \sqcap \pi_{G'} \sqcap \mathsf{eff}_G(c).$$

Outside in, the first hop does not touch z and the second gives $\overline{\pi}_{G''}(z) \sqcap \pi_{G''} \sqcap \mathsf{eff}_{G\oplus_c G'}(\overline{c})$, and $\mathsf{eff}_{G\oplus_c G'}(\overline{c}) = \mathsf{eff}_G(c) \sqcap \mathsf{eff}_{G'}(c')$ by Lemma 4 at the first hop (its hypotheses hold: both graphs are well-formed and $\varsigma_G(c) = \bot$). The two differ only by the factor $\pi_{G'}$, and $\pi_{G'} \in Q_{G'}(c')$ by clause (a) of Def. 2, so $\mathsf{eff}_{G'}(c') \sqsubseteq \pi_{G'}$ and that factor is absorbed. The two are equal in Π, so their canonical forms coincide; the implementation stores every meet in canonical form on the pattern coordinates. For $y \in V(G') \setminus \{c'\}$ both orders give $\overline{\pi}_{G'}(y) \sqcap \pi_{G'} \sqcap \mathsf{eff}_G(c)$, at the outer hop and the first hop respectively; nodes of $G \setminus \{c\}$ are untouched in both, and so is the root, π_G.

Scopes. For z in region $\bot$ of G'': inside out it takes c' at the inner hop and is renamed to $\overline{c}$ at the outer; outside in it takes the call-site label $\overline{c}$ directly. Nodes of G' in region $\bot$ take c in both. Nodes inside a region their graph already carried: below, under pre-existing regions.

Frozen environment and nesting. Inside out, the inner hop records $c' \mapsto \mathsf{eff}_{G'}(c')$ and $c' \mapsto \bot$ in the nesting; the outer hop transfers the value as $\overline{c} \mapsto \mathsf{eff}_{G'}(c') \sqcap \pi_{G'} \sqcap \mathsf{eff}_G(c)$ (step 3), re-parents $\overline{c}$ under c, and records $c \mapsto \mathsf{eff}_G(c)$, $c \mapsto \bot$. The factor $\pi_{G'}$ is absorbed as above, so the transferred value is $\mathsf{eff}_{G'}(c') \sqcap \mathsf{eff}_G(c)$. Outside in, the first hop records $c \mapsto \mathsf{eff}_G(c)$, $c \mapsto \bot$, and the second

$\overline{c} \mapsto \mathsf{eff}_{G \oplus_c G'}(\overline{c}) = \mathsf{eff}_G(c) \sqcap \mathsf{eff}_{G'}(c')$ with parent $\varsigma(\overline{c}) = c$. The same map in Π, and the same nesting, on the entries the two hops create.

Pre-existing regions. The paragraphs above read every node as $\bot$-scoped in its own graph and list only the entries the hops create. WF-11 admits bases that were composed before, so let t be a region of G'', s a region of G' and r a region of G, each with its frozen value, its recorded parent and its nodes. Inside out, the inner hop renames t to $c'__t$ and writes $\widehat{\pi}_{G''}(t) \sqcap \pi_{G''} \sqcap \mathsf{eff}_{G'}(c')$ for it (step 3); the outer hop renames it again, to $\overline{c}__t$, and meets in $\pi_{G'} \sqcap \mathsf{eff}_G(c)$. Outside in, the first hop does not touch G'', and the second renames t to $\overline{c}__t$ and writes $\widehat{\pi}_{G''}(t) \sqcap \pi_{G''} \sqcap \mathsf{eff}_{G \oplus_c G'}(\overline{c})$, with $\mathsf{eff}_{G \oplus_c G'}(\overline{c}) = \mathsf{eff}_G(c) \sqcap \mathsf{eff}_{G'}(c')$ as before. The two differ by the factor $\pi_{G'}$ alone, and $\mathsf{eff}_{G'}(c') \sqsubseteq \pi_{G'}$ absorbs it, exactly as for the overrides. So in both orders

$$\widehat{\pi}_H(\overline{c}__t) \;=\; \widehat{\pi}_{G''}(t) \sqcap \pi_{G''} \sqcap \mathsf{eff}_{G'}(c') \sqcap \mathsf{eff}_G(c).$$

Its parent is renamed by the same maps as its label: a recorded parent t_2 that the hop renames becomes $c'__t_2$ and then $\overline{c}__t_2$ inside out, and $\overline{c}__t_2$ at once outside in (every non-$\bot$ parent is a region name, so every hop renames it); the parent $\bot$ becomes the call-site label, c' and then $\overline{c}$ inside out (the outer hop renames c' like every region name of $G' \oplus_{c'} G''$), and $\overline{c}$ at once outside in. A region name of G'' that no scope label and no frozen entry names — a nesting record alone — is carried by the same two renamings with the same parent rule, so both orders record it under $\overline{c}__$ with parent $\overline{c}$. A node z of G'' with $\varsigma_{G''}(z) = t$ takes the scope $\overline{c}__t$ by the same two renamings, and its override is the expression of the override paragraph, since step 3 writes every node of the callee alike. For a region s of G': inside out, the inner hop leaves s, its nodes and its records alone — it copies G' and adds entries under keys that are new by freshness — and the outer hop renames it to $c__s$ with value $\widehat{\pi}_{G'}(s) \sqcap \pi_{G'} \sqcap \mathsf{eff}_G(c)$, parent $c__s_2$ or c, scope $c__s$ on its nodes and override $\overline{\pi}_{G'}(y) \sqcap \pi_{G'} \sqcap \mathsf{eff}_G(c)$ on each; outside in, the first hop writes exactly these and the second hop, at $\overline{c}$, leaves them alone. A region r of G and its nodes are untouched by every hop in both orders. So (a) holds at every node and (b) on the whole frozen environment and nesting map, not only on the entries the hops create. The factor $\pi_{G'}$ is not idle: when the middle root was tightened after G''s own composition it is what keeps WF-11 (c) at $c__s$ (the closure case of Lemma 1), and it sits in both orders' value there; at $\overline{c}__t$ it is absorbed, so the orders agree there too. An earlier version of this proof stopped at the created entries; had these cases not closed, the theorem would have been narrowed to uncomposed bases.

Overrides on old nodes. A node inside a pre-existing region is written by step 3 like any other node of its callee, so the override paragraph applies to it unchanged: the same expression in both orders, the extra $\pi_{G'}$ absorbed.

Edges. Inlining redirects each executable edge into the call site to the callee's entry, keeping the edge otherwise as it was; reissues each executable edge out of the call site from every callee exit, with the same type, guard and specification; re-sources each data-flow edge out of the call site at the first exit; relays each obligation edge into the call site onto the inlined acting nodes, and onto nothing when there is none; and deletes the call site. Both orders perform the two inlinings on the same incident sets — an edge of G' at c' is rewired at the inner hop and then carried, or carried and then rewired at $\overline{c}$ — so the resulting relation, with each edge's type, guard and specification, is the same. The relay needs the "onto nothing" clause: the acting nodes of $G' \oplus_{c'} G''$ are those of $G' \setminus \{c'\}$ and of G'', so an obligation into c reaches the same nodes inside out (one relay) and outside in (a relay onto G''s acting nodes, c' among them, then the edge at c' relayed onto G'''s); with G'' inert the outside-in second relay lands on nothing, as the inside-out relay never reached G''. An earlier operator relayed onto the callee's *entry* in that case, and the two orders then differed by one inert edge on different entries (fixture F5 of the checker). The *identifiers* of the rewired edges are generated from the prefix and the old edge id, and those strings differ; see the last paragraph.

Conclusion. Definition 2 reads only nodes, overrides, scopes, the root, the frozen environment and the edge relation, and meets policy *values*; so $Q_{H_{\mathrm{in}}}(n)$ and $Q_{H_{\mathrm{out}}}(n)$ are the same set of values and eff agrees in Π. Every wf condition reads those and the node types, and Φ is a property of traces of the semantics, which reads the same. □ □

What replaced the comparability hypothesis, and why there was one. The proof meets the same policies in two bracketings, so it needs $\sqcap$ to be associative. On the allow-list coordinate the released meet was *not*: it decided containment by matching one pattern *string* against the other as though it were a URI, kept the narrower of a nested pair, and *dropped* a pair where neither subsumed the other.

That is sound (Prop. 5) and exact only when no such pair occurs. Three well-formed graphs with overrides `tool:a*`, `tool:*b` and `tool:ab` gave

$$\mathcal{I}(\mathcal{I}(A,B),C) = \{\texttt{subskill:*}\}, \qquad \mathcal{I}(A,\mathcal{I}(B,C)) = \{\texttt{tool:ab}, \texttt{subskill:*}\},$$

so the two assembly orders disagreed at the sub-contracted node, and `tool:ab` — which all three policies admit — was refused by one order and allowed by the other. Swapping the overrides reversed the direction, so neither order was uniformly the stricter, and the operator's *domain* moved too: a Tool-call on `tool:ab` inside the innermost graph made one order compose and the other refuse at WF-7. This was found by an adversarial reading of the proof, not by our randomised checks, whose generator used literal tool names and so never built such a pair. The theorem then carried a hypothesis that every contributing allow-list is pairwise *comparable*, checked list by list on the released artifacts. That hypothesis and its checker had the wrong shape: the triple above satisfies them list by list and fails across the three lists that contribute to one composition, and the checker's matcher was not the runtime's. We did not repair this with a cross-list checker. The repair is the grammar: `tool:*b` is not a pattern (App. A.2), so `validate()` rejects the graph that carries it with a structural error naming the list, both orders refuse before any meet is taken, and no order is chosen for it. On the grammar every pair of patterns from any two lists is nested or disjoint, so the theorem needs no hypothesis about the lists beyond well-formedness, and the intersection is exact rather than a safe approximation. What we ship is inventoried against the grammar rather than assumed to be in it (`scripts/check_comparable_allowlists.py`, which keeps its historical name): 103 pattern lists in 14 of 879 candidate files — 68 root policies, 32 node overrides, 0 frozen maps (no composed graph is serialised in the release, so that source is exercised only at run time, where `validate()` reads it), 3 other — holding 56 distinct patterns, 110 literal entries, 40 universal, 0 prefix-star, and 0 outside the grammar. The triple is pinned as fixture F2 of the checker and as a test; its positive control, `tool:ab*` in place of `tool:*b`, composes in both orders with $\{\texttt{tool:ab}, \texttt{subskill:*}\}$ at the leaf.

Why the theorem is stated in canonical form. The old meet was not commutative as a function on lists: $\{\star\} \sqcap \{\star, \texttt{tool:a}\}$ returned $\{\star\}$ and the reversed arguments returned $\{\star, \texttt{tool:a}\}$ — the same language, two representations — so a theorem promising identical overrides and frozen values promised more than the algebra of lists could give. Clause (b) is therefore equality in Π, decided as equality of canonical forms; the meet now returns the canonical form, $\{\star\}$ in both argument orders; and fixture F4 of the checker pins this example and composes a chain whose call-site override is *written* redundantly.

Where the algebra is actually used, and the one place it failed. The override step above needs $\sqcap$ to be commutative, associative, idempotent and absorbing. Every component of a policy is a set operation — intersect the allow-lists, union the forbids, take the higher risk, union the triggers, intersect the stores — and has all four. The *value guard* is the exception: it is a string, and the meet built it as `(a) or (b)`. Logical $\vee$ is idempotent, commutative and associative; that string is none of the three. So without canonicalisation the theorem is false on the guard component: a two-hop chain with guards x and y gives the leaf gate `((x) or ((y) or (x))) or ((y) or (x))` one way and a longer, differently-parenthesised equivalent the other. A randomised check whose generator never sets a guard tests only the components where $\sqcap$ is a set operation, so ours sets guards. The guard meet is canonical — flatten, de-duplicate, sort — which restores the three properties and, incidentally, stops the gate condition growing by a duplicated copy of itself at every hop. Only the shape the meet itself builds is split apart, so a guard an author wrote is kept verbatim; every guard in the released evidence is a single atom, unchanged.

What is *not* order-independent, and what the checker compares. The generated identifiers of the rewired edges differ between the two orders (`c__c'__out0_x1` against `c__c'__out0_c__x1`). Edge identifiers are serialised, so the two orders produce different sealed bytes and therefore different *commitments*, even though they are the same graph: two parties who assemble the same chain in different orders will not agree on a digest. No semantics reads an edge identifier — the retry bound is keyed on $(\mathsf{src}, \mathsf{dst})$ — so wf, eff and Φ are unaffected. We do not canonicalise edge identifiers, because that would change the digest of every sealed graph already frozen in this paper's evidence and would buy no guarantee. `scripts/check_composition_order.py` compares the theorem's endpoints by name — eff at every node, and the root, every override and every frozen value, each after projection to the seven coordinates of Π — and, beside that, a comparison

stronger than the theorem: the full implementation policy at every node (TTL and disclosure list included) and every other serialised field except the edge identifier, reported by its path when it differs. The stronger comparison can report auxiliary implementation differences the theorem does not cover — edge multiplicity, a disclosure reference — and keeps doing so; a line from it is a diagnostic, and the projected check says whether the theorem's equality holds. It runs seven decisive fixtures first — a comparator self-test that perturbs each of the 221 serialised leaves of a composed graph one at a time and requires the comparator to name it, and requires the Π projection to *report* a tightened frozen value at the region and at the nodes whose walk crosses it (without this, a projection that never reports would pass every fixture), the `tool:*b` triple (refused at the grammar in both orders), its positive control, the redundant-universal example, F5, an obligation edge into a call site whose innermost callee is inert (the shape on which a relay could put the edge on a different entry in each order, step 4), F6, one relation duplicated under a fresh edge id, which the comparator must report at `edges.len`, and F7, the accepted deleted-call-site disclosure triple of the theorem's last sentence, on which the full comparison must report `opaque` at the frozen value of $\bar{c}$ and at every inlined override and nothing else, the projection to Π must report no node, and whose twin F7′ with an empty disclosure list must show no difference (the count of fixtures here is checked against the checker's list by `check_source_sync.py`) — and only then the randomised diagnostic, on the held-out seeds 100–119 that every earlier run used. Because those seeds were reused after each amendment, the run is *post-amendment validation*, not a fresh confirmation (`experiments/nesting-order/protocol.md`, Amendments 4 and 5). It attempted 4,000 triples: 997 were not built by the generator, 0 were refused by both orders, 0 by exactly one order, and 3,003 composed both ways, with 0 disagreements (0 at eff, 0 at another serialised field); 2,007 of the triples carried an obligation edge into a call site and 750 had an inert innermost callee, two shapes the generator could not draw before F5, so its earlier zero could not have seen the relay disagree. Edge identifiers were identical in 0 of the 3,003 and sealed bytes in 0. So "differ only in edge identifiers" is what the comparison shows on these chains, and it is reported as an observation rather than deduced from the theorem, which does not claim it. The generator draws no disclosure list, so this zero says nothing about F7's shape: the sweep is kept with that domain label, and a `--verify` run with the current checker re-derives every macro and matches the stored file byte for byte. A Π disagreement is always an eff disagreement of the stronger comparison (the key compares the seven coordinates and two more), so 0 at eff covers the theorem's projected equality on these chains.

C.3 Selective-disclosure composition

This subsection states and proves a selective-disclosure result: composition still preserves well-formedness when some of the callee's nodes are hidden behind facades. Let $G'^{\flat}$ replace a set $O \subseteq V(G')$ of internal nodes with opaque facades; each facade ν carries witness set $W(\nu)$ with $\{\mathsf{WF\text{-}3}, \mathsf{WF\text{-}7}, \mathsf{WF\text{-}8}\} \subseteq W(\nu)$ (witness completeness, premise (v) of the lemma) and an exposed signature $\Sigma(\nu)$ with the hidden node's port names. The hypotheses of the base lemma carry over unchanged: $\mathrm{wf}(G)$ and $\mathrm{wf}(G'^{\flat})$ in all eleven conditions, premises (i)–(iv) on the exposed nodes, and $\varsigma_G(c) = \bot$, since the exposed-node cases below are the base lemma's cases and read Eq. (2) through that scope hypothesis.

How a facade is typed, and the boundary premise. The released operator (`a2a_adapter.with_selective_disclosure`) keeps the facade's port names and re-types every port to `Opaque`, tagged by port. The subtyping rules of Appendix A.1 are directional: rule (3) gives $\tau \sqsubseteq$ `Opaque` for every τ, and `Opaque` $\sqsubseteq \tau$ holds only for $\tau =$ `Opaque` with the same tag. So the two kinds of incident edge must be argued separately, and the outgoing kind needs a premise: *(v) boundary typing:* every data-flow edge (ν, w) leaving a facade satisfies Δ on $G'^{\flat}$, i.e. $\Sigma(\nu).\mathsf{out} \sqsubseteq \Sigma(w).\mathsf{in}$. Under the released retyping this holds iff w's port is the matching `Opaque`; a facade that instead retained its interface types would satisfy (vi) for free, since the edge would then be the one G' already typed. Either way (v) is decided on $G'^{\flat}$ alone.

The case enumeration of the previous subsection applies verbatim to all conditions whose hypotheses mention only *exposed* structure: WF-1, WF-2, WF-4, WF-5, WF-9, WF-10 are checked on the facade graph exactly as before (facades are single nodes with in/out edges preserved, so reachability and acyclicity project as in WF-1/WF-2; a facade conservatively inherits the strictest side-effect class of its hidden interior, capped at irreversible, so WF-4 requires the facade to carry a fallback-to edge

or on_fail = ABORT, which selective disclosure preserves by retaining the facade's incident edges). WF-6 is argued in two directions. *Incoming* edge (u, ν): the destination port is `Opaque`, so Δ holds by rule (3) for every source type. *Outgoing* edge (ν, w): rule (3) does not apply, because the opaque type is the source; Δ holds by premise (vi) and by nothing else. Trust qualifications about the hidden interior say nothing about this edge: they concern what the facade did, not what type its output has.

The three conditions whose hypotheses reach *inside* the hidden region—WF-3 (are required verifications covered inside ν?), WF-7 (do hidden actions respect the effective policy?), WF-8 (are hidden high-risk nodes gated?)—cannot be checked by the consumer on $G'^{\flat}$ and are exactly the conditions delegated to the witness set $W(\nu)$.

Implementation, and why the witness is not optional. $W(\nu)$ is implemented as `DisclosureWitness` (`code/skillgraph/schema.py`): the union of tool URIs the hidden region may call (WF-7), the verification kinds it vouches for (WF-3), whether its high-risk nodes are approval-gated (WF-8), and its strictest risk level and side-effect class. The consumer's trust check reads it, and it is serialised with the sealed graph so it survives the commit.

This is a soundness requirement, not bookkeeping. An earlier implementation replaced a hidden node with a bare Goal carrying no witness. Hiding a node then removed it from the consumer's forbidden-tool check entirely, so a provider could evade the consumer's whitelist simply by declaring the node opaque—defeating Φ_2 across the disclosure boundary. The regression is covered by `test_e2e_negotiation.py`. We note one gap that remains: the witness *records* the conservative side-effect class, but WF-4 and WF-8 are not yet enforced *through* a facade by the validator, so a consumer must currently treat `approval_gated` as an assertion by the provider rather than a checked fact—which is what trust profile T1 already means.

Witness completeness makes the delegation total: every condition the consumer cannot check is asserted by the witness. Soundness of the assertion is a trust assumption graded by profile: at T1 the witness is believed; at T2 the consumer additionally replays the trace checker, which verifies every VERIFY-PASS, TOOL-CALL, and APPROVE event that fired *including inside the facade if the trace is forwarded*. T3 execution binding does *not* discharge this assumption—it binds gated steps to an execution, not the contents of an opaque interior—so discharging it entirely needs hardware attestation of that interior, which we do not implement. Hence $\mathrm{wf}(G \oplus_c G'^{\flat})$ holds relative to the stated trust profile, which is the qualified conclusion of the selective-disclosure result. ■

Instance (4): an opaque output consumed at a concrete type. A facade ν hides a node whose output port `x` had type `str`, and a data-flow edge carries `x` to a port of type `str` at w. After retyping, the edge is `Opaque(out/x)` $\rightarrow$ `str`. Rule (3) does not apply, because the opaque type is the source, not the destination. No other rule matches an `Opaque` source: (1) and (2) need a base type or a reference on both sides, and (4) and (5) need an `Array` or a `Struct` on both sides. So Δ fails, premise (vi) is violated, and the facade graph is rejected at WF-6. Retyping w's port to the same tagged `Opaque`, or retaining `str` on the facade's port, restores Δ. The earlier proof concluded Δ for this edge from $\tau \sqsubseteq$ `Opaque`, which is the incoming direction, not this one.

Remark (why the qualification is necessary). Without witness completeness, a malicious provider could hide a forbidden tool call behind a facade and the consumer's static check would pass; the composed graph would be "well-formed" only in the consumer's view. The lemma therefore *cannot* be strengthened to an unconditional wf claim at T1, and the paper does not claim it: end-to-end audit of hidden interiors requires T2 with trace forwarding, or hardware attestation of the interior, which is beyond the T3 we implement (§3.2, App. D).

C.4 LIMITS OF CONTROL-PLANE ENFORCEMENT

Both statements below are one-line observations, not theorems. We state them so that the experiments that show them have a precise target.

Remark 1 (A value-dependent gate needs a trusted value). Let φ be a non-trivial gate predicate: there exist values v_0, v_1 with $\varphi(v_0) = \bot$ and $\varphi(v_1) = \top$. Let the executor decide from a *reported* value $\hat{v}$ chosen by the constrained party. Then no control-plane mechanism attains both zero unsafe-pass and zero needless-gate.

Table 3: **What each profile proves.** A provider can obtain a valid co-signature and then skip the tool, so a co-signed run is not execution-bound. The two T3 rows are separate statuses, and the verifier returns the one the trace can support. No status shows that the trace contains every gated call: the chain has no anchored tail, so a truncated trace still verifies (App. D). Nothing here survives a provider and witness who *collude*.

Status	Proves	Does *not* prove
verified (T1/T2)	the trace is internally consistent and every gated call was authorised for the arguments the trace records	that it describes anything that happened: it is the provider's own report, and a world effect its approval never covered can pass
witnessed (T3)	an outside party was told this step, these arguments, this order, *before* it was claimed to run	that the step ran — the provider can get the signature and then do nothing
execution_bound (T3)	the step ran on those arguments: the witness ran the tool and signed its result	which code ran; and it needs a tool the witness can reach, not one inside the provider

Proof. Fix any mechanism E and consider its behaviour when it reads $\hat{v} = v_0$. There are two cases and both fail. If E passes ungated, then a party whose true value is v_1—a case the policy says must be gated—reports $\hat{v} = v_0$ and passes: unsafe-pass > 0. If instead E gates, then a party whose true value really is v_0 is gated although the policy does not require it: needless-gate > 0. E was arbitrary. □ □

The converse is immediate: when the value is observation-derived, so $\hat{v} = v$, evaluating φ on it attains both zero. Data-plane integrity is therefore not a convenience for value-dependent gating—it is necessary and sufficient. This is the formal content of the guard row of Table 1, and the trusted-guard study in the technical report measures both sides: the claimed-value guard leaks exactly where the model misreports, and the observed-value guard does not leak at all.

Remark 2 (Semantic steering is invisible to the control plane). Suppose two intents $i_0 \neq i_1$ are realisable by the same permitted action—the same tool with arguments the effective policy admits. Then no control-plane mechanism distinguishes them, and any attack that stays inside the permitted action set is undetectable by such a mechanism.

Proof. A control-plane decision is a function of (G, T, a) where a is the proposed action. By hypothesis i_0 and i_1 induce the same a and the same admissibility verdict, so the function returns the same decision on both. A mechanism that returns the same decision cannot separate them. □ □

D T3: BINDING THE TRACE TO AN EXECUTION

One rule, stated first. A signature proves whoever signed it saw something. What matters is *what* they saw.

- A **co-signature** is signed over a *declaration*: the provider says "I am about to call this tool with these arguments", an independent witness signs that sentence, and the provider then makes the call itself. This proves *prior disclosure* — an outside party was told the step before it was claimed to run. It does **not** prove the step ran. Such a run verifies as witnessed.
- A **receipt** is signed over a *result*: the witness holds the tool, calls it, and signs what came back. A provider that skipped the call has no result the witness produced, so it cannot make one up. This proves *witness-observed invocation*, and only this verifies as execution_bound.

Everything below follows from that distinction. Neither form survives a provider and witness who *collude*, and neither says anything about *which code* ran — a TEE quote would, and we do not have one.

How it is built. The bytes a witness signs are canonical JSON. A co-signature covers seven fields: contract id, node, tool URI, digest of the final arguments, step index, the hash of the previous statement, and the run id. A receipt covers those plus a digest of the result. Each field earns its place. Contract and node say *which* step. The argument digest puts the arguments inside the signed bytes, so they cannot be rewritten afterwards. Index and previous hash chain the statements *within* a run, so a step cannot be dropped from the middle or reordered. The result digest is the one that turns a declaration into evidence of an execution.

The run id stops a receipt from being replayed. The verifier restarts the index-and-previous chain at $(0, \varnothing)$ for every trace it reads, so without a run id two runs of the same sealed graph on the same arguments produce *byte-identical* receipts. The provider writes its own witness client, and nothing obliges that client to contact the witness. A provider could keep one genuine receipt and re-serve it, and every later run would verify as execution_bound, certifying a call nobody made. The run id inside the signed bytes closes this, provided the executor gives each trace its own identifier: an executor that builds its trace with `run_id=run_id or ""` gives every run the same empty name. `test_t3_receipt_replay.py` fails if either the signed run id or the per-trace identifier is removed.

What the chain still does not do. It has no anchored tail. Dropping a step from the middle breaks the links, but truncating the *end* leaves a shorter chain that is still contiguous and still verifies, so a provider can hide the last gated call the witness really mediated. Fixing it needs the witness to sign a closing statement over the session length, which our service does not issue. Until it does, T3 says: every gated call you can see was mediated, and it does not tell you that you can see all of them.

The verifier prefers a receipt wherever the trace carries one, and never reports a co-signature as though it were a receipt.

Evidence. The tests assert the boundary in both directions, which is the part that matters. `test_t3_execution_binding.py` builds the hand-written forgery a provider acting alone could write — correct nodes, correct order, valid hash chain, a bound authorisation with every field right and a matching confirmation — and asserts that it **passes** T2 and **fails** T3. If the first of those ever starts failing, T2 has become stronger and this appendix needs rewriting; the test says so. Further cases cover a forged witness record signed by the wrong key, arguments rewritten after a real signature, an unwitnessed run, a deleted receipt, a rewritten result, and a witness asked to sign for a tool it does not hold. One asserts that adding T3 does not change what T2 accepts.

Why the tool has to be reachable by the witness. A receipt exists only if the witness can call the tool, not only the provider. That is true of an external endpoint — a payment API, a ticketing system, anything the provider calls rather than owns — and false of a tool that only exists inside the provider's process. For those, `witnessed` is the most that can be claimed, and the verifier says so rather than rounding up.

What it costs. Over 100 episodes per arm, interleaved (`code/scripts/t3_remote_cost.py`). Each line gives the *total* time per gated step, then the *increment* over the line above it, as the median of the per-episode paired difference:

- co-signature, witness in the provider's process: 5.44 ms per gated step in total.
- co-signature, witness in its own process holding its own key: 7.83 ms in total; the increment over the in-process witness is 2.30 ms (IQR $[1.20, 3.36]$). 100 of 100 runs verify as `witnessed`.
- **mediated**, the witness runs the tool and signs the result: 9.43 ms in total; the increment over the out-of-process co-signature is 1.87 ms (IQR $[0.54, 3.01]$). 100 of 100 verify as `execution_bound`, and the witness performed 100 of the 100 side effects — the provider's own copy of the tool raises if it is ever called, so the test fails loudly if the executor routes around the witness.

So on loopback an execution-bound gated step costs 9.43 ms in total, against 5.44 ms with the witness inside the provider's process. Medians of paired differences do not add, so the two increments need not sum to the difference of the totals. Across a real network the round trips cost whatever that network costs. Host: Darwin x86_64, measured 2026-09-26 UTC. Every number here is generated from `code/results/t3_remote_cost.json`.

What none of this buys. A provider and witness who collude are indistinguishable from two honest parties, at every profile. And the witness process is still one we launch, on our own machine, in our own trust domain. The key is out of the provider's reach; the *organisation* is not. A deployment needs a witness run by someone with no stake in the provider's answer, and we do not have one.

Why a co-signature is not execution binding. A provider acting alone can ask the honest witness to sign, skip the tool, and write the trace with the real signature in it. Nothing is forged and the witness behaves correctly, yet a verifier that treats a co-signature as execution binding returns `execution_bound` on a run in which nothing was paid. A co-signature alone is a signing oracle, not a monitor. The split above — witnessed for a declaration, execution_bound for a witnessed invocation — closes this, and `test_t3_signing_oracle.py` keeps the attack as a regression test. App. N records the earlier claim.

E A2A PROTOCOL WIRE FORMAT

This appendix specifies each of the six PROCPACT negotiation messages as a JSON Schema (Draft 2020-12), specifies the signature scheme on COMMIT-EXECUTION, and shows a worked diff against a published Agent Card. All schemas are released in machine form at `code/skillgraph/protocol/schemas/*.schema.json`.

E.1 COMMON ENVELOPES

Every message carries:

```
{"sg_protocol_version": "0.1",
 "message_id":          "<uuid>",
 "in_reply_to":         "<uuid-or-null>",
 "from":                "<did:agent:...>",
 "to":                  "<did:agent:...>",
 "ts":                  "<rfc3339>",
 "nonce":               "<base64url(16-byte)>"}
```

Acceptance window for (nonce, ts) is fixed by the consumer in REQUEST-SKILL-GRAPH and re-stated in every reply.

E.2 MESSAGE SCHEMAS

```
{"type": "request-skill-graph",
 "skill_id":      "<URI>",
 "skill_version": "<semver-range>",
 "trust": {
   "required_verification_kinds": ["schema", "llm-judge"],
   "forbidden_tools":             ["tool:wire_transfer", "tool:purge_data"],
   "max_risk":                    "medium",
   "hitl_threshold":              {"price_gt": 1000},
   "acceptance_window_s":         300,
   "trust_profile":               "T2"}}
```

```
{"type": "offer-skill-graph",
 "graph":     { /* SkillGraph JSON; opaque facades allowed */ },
 "facade_witnesses": [
   {"facade_node_id": "n7_opaque",
    "wf_witnesses":   ["WF-1","WF-2","WF-3","WF-7"],
    "exposed_signature": {"in": [...], "out": [...]}}
 ]}
```

```
{"type": "negotiate-verification",
 "splice": [
   {"target_node": "n3_search",
    "verification_node": {"type":"Verification",
                          "spec":{"check_kind":"schema",
                                  "predicate":"tool:check_pci_compliance"}},
    "edge_type": "verification-required-before"}
 ]}
```

```
{"type":      "commit-execution",
 "payload":   {"sealed_graph": "<base64(canonical-json(graph))>",
               "policy_snapshot": { /* eff(root) of sealed_graph */ },
               "commit_window_s": 600},
 "signature": "<JWS-compact-detached>"}
```

Signature is JWS (RFC 7515) over the canonical JSON encoding of the `payload` (RFC 8785 JCS). The public key is resolved from the provider's W3C Decentralized Identifier in the envelope's `from` field (W3C, 2022). Algorithms: EdDSA (Ed25519, `EdDSA`) by default, ECDSA (`ES256`) for legacy.

```
{"type":  "execution-trace",
 "ref":   "<message_id of commit-execution>",
 "trace": [ /* trace events with hash chain */ ],
 "outcome": "ok" | "fallback" | "aborted",
 "artifacts": [ /* content-addressed artifact descriptors */ ]}
```

```
{"type":  "verify-trace",
 "result":     "verified" | "failed",
 "Phi_check":  {"Phi_1": true, "Phi_2": true, "Phi_3": true},
 "violations": [ /* node ids where the soundness invariant failed */ ]}
```

E.3 VERSION NEGOTIATION

The envelope carries `sg_protocol_version` (the version of the sender) and an optional `sg_compat_versions` array (a list of older versions the sender is willing to downgrade to). The negotiation rule for a consumer-provider pair is:

1. If `request.sg_protocol_version` matches `provider.sg_protocol_version` exactly, proceed.
2. Else, the following intersection is computed:

   $$R = \{\texttt{request.sg_protocol_version}\} \cup \texttt{request.sg_compat_versions}$$
   $$P = \{\texttt{provider.sg_protocol_version}\} \cup \texttt{provider.sg_compat_versions}$$

   and the negotiated set is $R \cap P$;

   if non-empty, the highest version in the intersection is selected and the provider responds with the chosen version in its envelope.
3. Else, the provider returns an error with the PROCPACT-specific `type: "version-mismatch"` reply (out of scope of the six core messages); the consumer aborts or retries with a different provider.

This is the same shape as TLS/HTTP version negotiation: monotone in the acceptable set, no silent downgrade, no fallthrough to a default.

E.4 WORKED AGENT CARD

The card below is what `code/skillgraph/a2a_extension.py` actually emits, and it is the card the official-SDK conformance run parses (App. N). The A2A schema has no top-level `agent_url` and no extension keys inside a `skills` entry. An endpoint belongs in `supportedInterfaces`, each entry carrying its own binding and protocol version, and extension data belongs in that extension's `params`. A parser we did not write checks this shape; our own tests alone could pin a wrong one and still pass.

```
{
  "name": "did:example:travel-bot",
  "description": "SkillGraph-capable provider",
  "version": "1.0.0",
  "supportedInterfaces": [
    {"url": "https://travel.example.com",
     "protocolBinding": "HTTP+JSON",
     "protocolVersion": "1.0"}
  ],
  "capabilities": {
    "streaming": false,
    "pushNotifications": false,
    "extensions": [
      {"uri": "https://skillgraph.dev/ext/procedure-contract/v1",
       "description": "Typed procedure contracts: the caller may require named
                       verification kinds and human-approval gates, statically
                       checkable before execution and auditable afterwards.",
       "required": false,
       "params": {"verifyingKey": "9f86d081884c7d65..."}}
    ]
  },
  "defaultInputModes":  ["text/plain"],
  "defaultOutputModes": ["text/plain"],
  "skills": [
    {"id": "skill:travel.book_flight",
     "name": "Book a flight",
     "description": "Search, book, and email a confirmation.",
     "tags": ["skillgraph"]}
  ]
}
```

A consumer that does not know the extension reads the card normally and ignores the entry in `capabilities.extensions`; it then gets the prose `description` and none of the guarantees.

Three versions, three jobs. These are easy to confuse, so we name them separately. The *A2A protocol version* is `1.0`; it appears in each `supportedInterfaces` entry and in the `A2A-Version` header on every request, and an absent header means `0.3`, not "current", so our server rejects it rather than serving a silent downgrade. The *extension version* is the `/v1` at the end of the extension URI; changing the payload shape means a new URI, not a new field. The *contract schema version* is the `version` inside a procedure graph document, which says nothing about the protocol carrying it.

E.5 AGENT CARD ROUND-TRIP

A round-trip check: an A2A Agent Card carrying the four PROCPACT extension keys (§E.3) round-trips correctly through both (a) a base Agent Card schema with forward-compat `additionalProperties: true` (modelling a consumer that does not understand the extension) and (b) the PROCPACT-extended schema (modelling a consumer that does). The supplementary script `code/scripts/agent_card_roundtrip.py` runs five checks and writes `code/results/agent_card_roundtrip.json`:

```
$ python3 code/scripts/agent_card_roundtrip.py
{
  "all_checks_pass": true,
  "n_checks": 5,
  "sg_extension_keys": ["id", "name", "description",
                        "skill_graph_uri", "verification_policy",
                        "trust_profile", "composability_constraints"]
}
```

The five checks: base validates against base schema; SG-extended validates against base schema under `additionalProperties: true`; SG-extended validates against the SG schema; serialise-parse-serialise is byte identity in canonical JSON; a base-only consumer that strips the SG keys reconstructs the original base Agent Card. *This is not reference conformance.* The schema it validates against is one we transcribed from the published A2A specification, so the check is self-authored on both sides: if we mis-transcribed the schema, our card satisfies our mistake. It has one virtue – it runs offline with no SDK – and we keep it for that reason only. "Reference conformance" in this paper means exactly one thing: the run against Google's official `a2a-sdk`, whose protobuf types we

did not write (App. N). That run found a mis-shaped Agent Card this check had passed for months, which is the whole argument for keeping the two labels apart.

E.6 LOCALHOST CROSS-AGENT RUN

The supplementary archive includes `code/scripts/run_cross_agent_demo.py`, a stdlib-only (∼430 LOC) script that boots a Consumer and a Provider on localhost, exchanges all six PROCPACT messages over HTTP, validates each message against the JSON Schemas of §E.2, signs the COMMIT-EXECUTION payload with a fresh key, replays the trace through the soundness verifier of §3.3, and emits the full log to `code/results/cross_agent_demo.jsonl`. Running the script on a 2024-vintage MacBook with Python 3.9.4 yields:

```
$ python3 code/scripts/run_cross_agent_demo.py
{
  "total_wallclock_ms": 37.9,
  "n_messages": 6,
  "all_schemas_valid": true,
  "signature_verified": true,
  "verdict": "verified",
  "sig_decomposition_ms": {
    "ed25519_available": true,
    "sign_ms_total":   6.3,
    "verify_ms_total": 0.9
  },
  "log": [
    {"step":"REQUEST",   "schema_errors":[], "latency_s":0.013,
     "chosen_version":"0.1", "version_negotiation_ok":true},
    {"step":"NEGOTIATE", "schema_errors":[], "latency_s":0.003,
     "signature_verified":true},
    {"step":"EXECUTE",   "schema_errors":[], "latency_s":0.003, "trace_len":5},
    {"step":"VERIFY",    "schema_errors":[],
     "hash_chain_ok":true,
     "Phi_check":{"Phi_1":true,"Phi_2":true,"Phi_3":true}}
  ]
}
```

The version-negotiation step intersects the consumer's `sg_compat_versions` list with the provider's, picks the highest common version, and proceeds; the `chosen_version: "0.1"` field in the log records the negotiated value. The Ed25519 cost is decomposed: single `sign` at 6.3 ms (key resolution + signature), single `verify` at 0.9 ms; the rest of the 37.9 ms is JSON Schema validation, three HTTP round-trips on localhost, and the trace hash-chain recomputation. The signature scheme is RFC-7515 JWS compact serialisation with `alg=EdDSA` (Ed25519), implemented via the `cryptography.hazmat.primitives.asymmetric.ed25519` module of the `cryptography` package. The provider's signing key is generated fresh per launch; in deployment the key is resolved from the provider's W3C DID. Should `cryptography` be unavailable, the demo falls back to HMAC-SHA256 with header `alg=HS256-canonical` and announces the substitution at startup; the wire format and the trust-profile semantics are otherwise identical. The retained on-disk artefact is the JSONL log; the verifier's `Phi_check` confirms the three soundness conjuncts of Proposition 1 on the actually-executed trace.

F EXPERIMENTAL SETUP

Models and routes. All model runs use hosted endpoints through OpenRouter at temperature 0, on paid routes, with the route and the token budget recorded in every result file. The delivery control uses twelve models from eleven labs: `claude-haiku-4.5`, `deepseek-v4-flash`, `gemma-4-31b`, `glm-5.2`, `gpt-4o-mini`, `gpt-5.6-luna`, `kimi-k2.5`, `llama-3.3-70b`, `mimo-v2.5`, `minimax-m3`, `nemotron-3-nano` and `qwen3.7-flash`. The procurement study and the decision-interface experiment use three of them, `claude-haiku-4.5`, `gpt-4o-mini` and `gpt-5.6-luna`, with a budget of 400 output tokens per call.

The delivery control. Two task families. Research summarisation with claim verification has a catalogue with two forbidden shortcut tools and a procedure that requires `verify_citations` before the final artifact; procurement requires an approval before the irreversible order. The provider's own policy calls the caller's step optional, and only what the caller sends varies. A run is scored

only if it reaches the terminal action, and a skip is a terminal action with no passing instance of the required step before it.

Statistics. Temperature 0 does not make a hosted API deterministic, so the analysis rests on distinct cases rather than repeated calls, and every count is reported per model as well as pooled, because pooling can hide opposite effects by model (App. J). The procurement comparisons are pre-registered as counts against fixed thresholds, not as significance tests: with zero violations in both enforcing arms, no test could separate them.

G IS THE CONFLICT CONSTRUCTED OR OBSERVED?

The paper's premise is a caller whose requirement the provider's procedure does not already satisfy. We wrote both sides, so the premise could be an artefact of our authorship. This appendix replaces the assumption with a measurement.

Design. For each of five domains we elicit, from parties that never see each other's output: a *provider* procedure ("list the steps your service actually runs; you are paid per completed request and compete on turnaround and cost") and a set of *caller* requirements ("list what your company requires of an outside provider"). Neither prompt contains the words verification, approval, checking, safety, compliance or risk—otherwise the conflict would again be built in, which is the criticism rather than an answer to it. Procedures come from three model families, requirements from three different families, and pairs are formed at random across families so a requirement is never judged against a procedure by the same author. A seventh model, told nothing about this paper, classifies each requirement as naming an action or not, and judges whether the paired procedure already covers it.

The instrument is validated before the result is read. A judge told not to give the supplier the benefit of the doubt answers NOT-SATISFIED on nearly every pair, uptime and insurance clauses included, so its conflict rate measures the prompt. The released judge has no such instruction, and a positive control gates it: paraphrase one of a procedure's *own* steps into a requirement and check the judge recognises it. The judge scores $10/10$ on that control. A judge that fails it reports a conflict rate of 1.0 whatever the data, so the control gates the number rather than decorating it.

Result. Of 33 judgments on requirements that name an action, 31 are not covered. De-duplicating to one vote per distinct requirement—the correct unit, since a requirement is paired with several procedures—gives **18** of **19**, 0.95, Wilson $[0.75, 0.99]$; a bootstrap resampling whole domains gives $[0.84, 1.00]$. Both locked predictions hold (procedural rate $0.55 \geq 0.50$; conflict rate $0.95 \geq 0.30$). Of the 12 requirements judged against more than one procedure, 11 drew the same verdict every time.

What this settles, and what it does not. When the two sides are written independently and neither is primed, the caller's procedural requirements are overwhelmingly not already met. The premise is observed. Two limits are worth stating plainly. First, 45% of what callers write down is not procedural at all—service levels, insurance, uptime, payment-data compliance, liability—and PROCPACT addresses none of it; 0.55 is the honest scope of a procedure contract. Second, the parties are models. They are independent of each other and unprimed, but they are not human procurement officers, and a human elicitation remains the stronger test.

H THE PROCUREMENT STUDY

Arms and matched privileges. All arms act on one world, one tool catalogue, one approval authority and one scorer. The free-form arms run one agent loop: the model sees the catalogue and the request and replies with a tool call or a final report. *Rules in the prompt* adds Northwind's four rules to the system prompt. *Caller's monitor* is the same agent with a stateful monitor between its calls and the world. PROCPACT is the executor driving the sealed graph of Fig. 1(a). Table 4 lists what is matched. Two privileges are not matched, and both favour the monitor. The free-form models read Northwind's rules in prose, and the graph's model never does, because the contract is meant to replace the prompt, not add to it. And the monitor's model may ask the approver at any step, while

the graph asks only where the contract gates; since the approver grants only requests within scope, it can stop an out-of-scope order for the monitor that the graph never shows it. We report those stops separately.

Table 4: **Matched privileges of the monitor and PROCPACT.**

privilege	caller's monitor	PROCPACT
mediates every tool call	wraps the world's tools	the executor drives the tools
can veto any call	yes	yes
sees the arguments	yes	yes
carries state across steps	the automaton	the trace
approval authority	the same	the same
published catalogue	the same	the same
consults the approver	whenever its model asks	at the gate only
reads the caller's rules in prose	yes	no

The monitor was written to win. It is 138 lines of consumer-side code, an automaton over the ordered stream of attempted calls. It remembers which approval was granted, for which vendor, item and amount, and whether it was spent; it treats a delegated purchase as a purchase; and it applies the threshold to the largest amount it can establish (the request, the quote, and the agent's declaration), so an agent cannot get under $500 by declaring less than it spends. Only the approval tool grants approval; a request or a tool observation that claims otherwise is ignored.

Cases. 26 hand-written cases in two halves. The graph, the trust constraint and the monitor were authored against the first 13; the other 13 were written before the monitor existed and add three classes: a price known only mid-procedure, a delegated purchase, and an intermediary that tells the agent the rule does not apply. The halves are balanced so that always refusing scores 0.50 on task success. Three cases ask to cancel or modify an existing order, which rule R4 forbids.

The oracle. The oracle reads the world's event log and the authority's ledger, and never a mechanism's verdict. For each episode and rule it returns safe, violation or unknown; a rule whose evidence the log does not carry is unknown, not safe. R1 is a purchase of $500 or more with no granted approval bound to it before the purchase; R2 an unapproved vendor; R3 a purchase with no budget check before it, or an invalid code; R4 a cancellation or modification.

Two runs. A historical panel (624 episodes, with an unconstrained arm as the floor) found two defects that the results in the main text depend on. The provider's approval was granted over an empty argument map and the world filled in the order afterwards, so the graph arm made 33 purchases with no approval bound to them, each of which the trace verifier passed; and the world did not log the budget check, so R3 was unknown on many episodes. We bound the approval to the materialised order, logged the check, and re-ran the three enforcing arms on the same cases and models as a pre-registered regression check (468 episodes; 3 predictions held and the condition for comparing the arms was met; 0 unknown verdicts). Table 5 gives the full result, and the main text reports it. The historical episodes are kept and not re-scored under the fix.

Approval requests by stage. The graph arm's 162 requests are 84 first stages of the caller's approval (asked before the order's arguments exist; 42 denied), 42 confirmations of that approval at the final order, and 36 requests at the client report. The report is gated because the order node's guard `amount >= 500` is met into the report node's effective policy, and the executor evaluates it on the order's amount, so the report of every order of $500 or more is gated and no smaller order's report is (`test_report_gate_is_the_orders_guard.py`). The monitor's 85 requests are all at the order.

What the artifacts carry. Table 6 asks which view, before execution, shows a provider's change to its skill: a tool catalogue, the catalogue plus the provider's edge list, or the sealed graph. Table 7 gives the monitor the provider's edge list as a signed manifest and compares it with the graph on six scripted cases; the pre-registered prediction was parity, and parity is what it found. So the representation, not

Table 5: **The repaired system on real models** (pre-registered in `experiments/repaired-replication/`). The same 26 cases and three models as the historical panel, two repetitions, three arms, after two repairs: the provider's approval is bound to the materialised order, and the world logs the budget check. Scored by the independent event oracle, which reads the world's log and the authority's ledger and never a mechanism's verdict. *R1*: a purchase at or above $500 with no granted approval bound to it (historical panel: 33 on the graph side). *Correct*: the right outcome with no violation. *Unasked*: orders placed on the 18 episodes per arm whose request needs a tool the contract forbids; each breaks no rule. The two mediated arms agree because neither violates, so this is a regression check of the repairs, not a test that could separate them.

arm	episodes	violations	R1	unknown	rate	approvals/ep	correct	unasked
rules in the prompt	156	12	4	0	0.077	0.51	144	0
caller's monitor	156	0	0	0	0.000	0.54	156	0
ProcPact	156	0	0	0	0.000	1.04	147	9

enforcement, is what distinguishes the contract, and even the representation can be matched by an equally informed manifest.

Client messages are not an outcome here. Each arm ends with one model call that writes the client a message. The graph arm's message step was told less about the run than the free-form arms' (whether an order was placed, and a flag that misreported the approver), so its messages cannot be compared with theirs, and we draw no conclusion from them.

I The caller's approval: six fixtures

A negotiated approval clause should mean more than "some approval preceded the purchase". It should mean that the approval the caller required was obtained from the caller's own authority, for the effect that ran. Each fixture below negotiates a provider graph against a trust constraint that requires approval before every irreversible step and names Northwind's principal, then runs the executor with a scripted authority and checks the trace with the verifier (`test_caller_obligation_fixtures.py`).

F1 *The caller requires approval and the provider has no gate.* Negotiation splices an approval reserved for Northwind before the order. Northwind is asked twice, before the arguments exist and at them, and the order fires only after the second; a denial stops it.

F2 *The provider approves, the caller denies.* The provider's own approval node, reserved for its manager, does not discharge Northwind's clause, so negotiation splices Northwind's beside it. Northwind's denial stops the order, and a decision by the provider's manager at Northwind's node is refused, at either stage.

F3 *Both require approval.* The order fires only after each principal has confirmed at the final arguments; one refused confirmation stops it; a trace missing either confirmation fails verification.

F4 *An approval before the arguments is not enough.* An authority that grants the first stage and refuses every confirmation never lets the order fire.

F5 *Changed arguments.* A confirmation over a digest other than the final arguments' is refused.

F6 *One authorisation, one effect.* A second call on the same grant fails verification, and so does a trace that records Northwind's approvals under another principal.

Before these fixtures existed, F2 and F3 failed. Negotiation counted any approval node as the caller's, and nothing checked who answered one: the schema had a field for the required role that no code read. And when two approval nodes gated one action, the executor confirmed the first grant and fired, so the second principal never saw the final arguments. The fixes are the ones stated in §3.3 and the approval row of Table 1: a caller that names its approver gets an approval node reserved for that principal, and an action needs a confirmed grant from every approval node that gates it. Seven mutants of the mutation audit disable these checks one at a time, and each is killed. No recorded episode of the procurement study changes, because its graph has one approval node and one authority.

Table 6: **The provider changes its skill; which pre-execution view shows it?** A check mark means the view's projection of the new skill differs from its projection of the old one — nothing more. Three views, each from one provider-signed artifact: the tool catalogue (uri to class); the catalogue plus the provider's edge list, the information-matched control; and the sealed graph, which also carries per-node policy. Deterministic, so no interval. Last column: does the caller's approval edge cover every irreversible step in the sealed result, or does negotiation refuse to seal. Frozen cases and the dated post-hoc amendment are in `experiments/provider-change/protocol.md`; the version-1 output is preserved unchanged.

change	catalogue	+edges	sealed graph	covered
tool-level				
T1: new irreversible tool, on the existing path	✓	✓	✓	✓
T2: new irreversible tool, as an alternative branch	✓	✓	✓	✓
T3: anticipated tool as a branch (control)	✓	✓	✓	✓
T4: declared class raised pure → irreversible	✓	✓	✓	aborts
T5: declared class lowered irreversible → pure (a lie)	✓	✓	✓	–
T6: a tool is removed from the skill	✓	✓	✓	aborts
mixed: a deleted step is also a tool				
P1: the budget-check step (a tool) is deleted	✓	✓	✓	✓
procedure-only (same catalogue)				
P2: the budget check is moved after the order	–	✓	✓	✓
P3: an edge reaches the order skipping the budget check	–	✓	✓	✓
P4: a second branch reaches the order from triage	–	✓	✓	aborts
P5: the order step is split into two nodes, same tool	–	✓	✓	✓
P6: the fallback edge out of the order is removed	–	✓	✓	aborts
policy-only (exploratory)				
P7: policy only: approval guard 500 → 50 (exploratory)	–	–	✓	✓
controls				
O1: same artifacts; the tool implementation is swapped	–	–	–	✓
O2: same artifacts; the provider ignores its graph at run time	–	–	–	✓
O3: null control: nothing changes	–	–	–	✓
O4: harmless rename of one node id (exploratory)	–	✓	✓	✓

The main theorem on the running example. A second deterministic test, `test_running_example_delegation.py`, runs Theorem 1 end to end on the example of Fig. 1. Vendorly's order step is a Subskill-call to a third agent's skill (Fulfilco). Northwind negotiates its trust constraint into Vendorly's graph, Vendorly inlines Fulfilco's skill at the call site and signs the result, and the executor runs it. Five checks pass: a Fulfilco skill that calls `cancel_order` cannot be composed in; inside the inlined region the effective policy forbids Northwind's tools and requires its check; Fulfilco's order runs only after Northwind's check passed and Northwind's own approver approved the order's final arguments, and the trace verifies against the signed composed graph; an approval from Vendorly's manager does not let the order run; and a tool Fulfilco's own policy forbids is still allowed in Vendorly's step after the call returns.

Table 7: **Against an equally informed signed workflow (P11).** Six frozen cases, two arms fed from one provider input: (A) the negotiated graph run by the released executor, (B) a signed manifest exported from the same graph, read by the App. H monitor. One authority, one deterministic world, one independent scorer over the world's events and the authority's ledger. Per arm: outcome; effects; approval requests (bound to the transaction in parentheses); rule violations; a refusal names where it was decided. The repair-required indicator on W3–W5 (1 when the obligation fails to hold with the frozen consumer artifacts, else 0) sums to 0 in arm A and 0 in arm B; it is not an edit count. W6 is the shared negative control. Pre-registered in `experiments/signed-workflow/protocol.md`; locked prediction: parity.

	(A) negotiated graph				(B) signed manifest + monitor			
case	outcome	eff.	appr.	viol.	outcome	eff.	appr.	viol.
W1: valid purchase over threshold	completes	1	3 (1)	0	completes	1	1 (1)	0
W2: authority denies	refuses (authority)	0	1 (0)	0	refuses (authority)	0	1 (1)	0
W3: check moved after the order	refuses (negotiation)	0	0 (0)	0	refuses (manifest)	0	0 (0)	0
W4: new irreversible branch	completes	1	3 (1)	0	completes	1	1 (1)	0
W5: broker delegates downstream	completes	1	3 (1)	0	completes	1	1 (1)	0
W6: effect around the boundary (control)	effect, unapproved	1	0 (0)	2	effect, unapproved	1	0 (0)	2

J THE DECISION-INTERFACE EXPERIMENT

Protocol. Locked on 2026-09-23 before any episode, in `experiments/decision-interface/protocol.md`; the locked text and the digests of the ten files the run depends on are kept beside it. The predictions, with their outcomes on the data:

I1 The released graph places an unasked order on at least 6 of the 36 out-of-scope episodes (the failure reproduces). held.
I2 Propose-then-veto places at most half as many, pooled, and no more for any model. failed.
I3 Showing the terms alone does not reach that bound. held.
I4 Both together place at most two more than propose-then-veto alone. held.
I5 No cell loses more than three in-scope completions against the released graph. held.
I6 No graph cell violates a rule, and fewer than 10% of verdicts are unknown. failed.
I7 The monitor places at most two unasked orders. held.

The locked decision rule: the paper may say that the menu-only interface causes the unasked orders, and that a refusal with a reason removes them, only if I1 and I2 both hold; if I1 holds and I2 fails, the cause is unresolved. The verdict is *unresolved*.

The cases. Out of scope (gold: refuse): a cancellation phrased four ways, a later delivery date, a different delivery address, a different variant, fewer units, more units, a different vendor, a line item removed, and a hold. Every item field names the goods, and the amount is the existing order's value, under $500, so an unwanted order is a real purchase and no approval gate interferes. In scope (gold: place): twelve new purchases, including a reorder, a replacement, and six that name the same order or goods as an out-of-scope case. The vendors are approved and the codes valid.

Pilot. Eight episodes on `gpt-4o-mini` in the two propose-then-veto cells, not pooled, were all scorable. On two out-of-scope cases the model did not propose a tool; it took the purchase branch from the menu, and a real order was placed. We recorded this in the protocol before the full run.

The violations (I6). All 5 are of R2 on `L09`, where the request asks to move an existing order to "Acme". The model wrote that name into the order's vendor argument; it is not on Northwind's list, which names "Acme Supply". The trust constraint of the procurement study carries R1, R3 and R4

Table 8: **Why a compliant run does the wrong task.** Three pre-registered runs on the 24 new cases (12 that need a forbidden tool, 12 new purchases), three models, one run each. *Unasked orders*: orders placed on the requests that need a forbidden tool (12 per model). *In scope*: new purchases completed correctly with no violation. †No model proposed a tool in any veto cell, so the refusal path was never exercised. The only rule violations in any cell are R2 on `L09`, a name mismatch (App. J).

	unasked orders				
cell	haiku	mini	luna	all	in scope
first run: the decision-interface 2×2					
menu only (released)	3	7	6	16/36	36/36
propose-then-veto†	5	11	6	22/36	36/36
menu, contract's terms shown	3	12	1	16/36	36/36
propose-then-veto†, terms shown	5	12	1	18/36	36/36
follow-up run					
menu only (released, rerun)	3	8	6	17/36	36/36
menu, caller's rules in prose	1	3	0	4/36	29/36
propose-then-veto, permissive prompt†	5	9	5	19/36	36/36
caller's monitor (first run)	0	0	0	0/36	35/36
rule-blind run					
caller's monitor (rerun)	0	0	0	0/36	35/36
caller's monitor, model not told the rules	1	4	2	7/36	34/36

and has no field for R2, which it enforces only through the approver, and the approver is asked only on orders of $500 or more. The executor fills a declared binding only for an argument the model left out, so the model's vendor replaced the bound one. A vendor allow-list in the trust constraint, or bindings the model cannot override, would close this; neither exists in the released code, and we leave both to future work rather than change the contract after the run.

Run. 360 episodes (0 harness errors), scored by the oracle of App. H: 5 violations in the four graph cells and 0 unknown verdicts in all five. Approval requests per episode: 0.67, 0.67, 0.67 and 0.67 for the four cells, 0.39 for the monitor. The monitor's one in-scope loss (`L21`, `gpt-5.6-luna`) placed the right order and then wrote an empty client message; it is a reporting loss, not a wrong action. Table 8 gives every cell of the three earlier runs on these cases.

What the first run could not test. The first run's design had three flaws, recorded as dated corrections in the protocol. The executor's fixed system prompt tells the model to "select exactly one of the allowed actions" and "not invent tools or nodes", so the invitation to propose a tool contradicted it, and no model ever did. The pilot's stop criterion could not fail. And "terms shown" prints two tool URIs, not Northwind's rule, while the monitor's model reads all four rules in prose, so the task-failure contrast was confounded with that privilege.

The follow-up. Locked on 2026-09-23 before any episode, in `experiments/decision-followup/protocol.md`, one run on the same 24 cases and three models (216 episodes, 0 harness errors): the released graph rerun (A2); the released graph with Northwind's four rules in prose at the decision, the same text the free-form arms read, and nothing else (E); and propose-then-veto under a decision prompt that permits a proposal (B2). The predictions, with their outcomes:

G1 A2 places unasked orders on 12 to 20 of 36 (the first run's released cell placed 16). held.

G2 E places at most 4. held.

G3 In B2, at least 6 out-of-scope episodes contain a refused proposal (a manipulation check that can fail). failed.

G4 If G3 holds, B2 places at most half as many unasked orders as A2. not evaluated (G3 failed).

G5 E and B2 each lose at most three in-scope completions against A2. failed.

G6 No violation in any cell other than R2 on `L09`. held.

The locked decision rules: the paper may say the graph's model orders because it never reads the caller's rule only if G1, G2 and G5 hold (not supported); and that a refusal with a reason helps only if G3 and G4 hold (feedback: untested). In B2, 0 of 36 out-of-scope episodes proposed a tool.

Terms shown at the decision. In the informed-triage run, the graph's model was shown the contract's own terms at its decision. Out of 14 out-of-scope requests per model, unasked orders went from 4 to 0 (haiku), 10 to 14 (mini) and 4 to 0 (luna): information about the contract does not by itself stop substitution.

The simulated approver. Northwind's approver is a deterministic function of the case: it grants an order if and only if the vendor is approved, the code is valid, and the request is within scope, as a human approver who knows the request would. It does not read what it is shown, so in this harness a vendor the model substitutes would pass it at any amount; a human approver shown the order would catch one at $500 or more. It was asked on none of the graph's out-of-scope episodes, all below $500, and on three of the monitor's, where it stopped an order the monitor's model had proposed.

The rule-blind monitor. Locked on 2026-09-23 before any episode, in `experiments/rule-blind-monitor/protocol.md`, one run of two monitor cells on the same cases and models (144 episodes, 0 harness errors): the monitor as released, and the same monitor with its agent not told Northwind's rules. The predictions: the rule-blind monitor places at least 8 unasked orders (failed); the released monitor at most 2 (held); the rule-blind monitor breaks no rule (held) and completes at least 30 of 36 in scope (held). The locked decision rule gives the verdict *the agent loop*. The rule-blind agent called a forbidden tool on 36 of 36 out-of-scope requests and was refused each time; the released agent did so on 2, and the approver stopped 4 of its attempted orders.

A rule the contract carried only in part. On one case, "Cedar is late on PO-4675. Move that order over to Acme.", the graph's model wrote `Acme` into the order's vendor and the order was placed (5 of 288 graph episodes in the first run). The oracle counts these as R2 violations because Northwind's list names "Acme Supply", so they are a name mismatch. What they expose is real, and a deterministic run shows it with a vendor Northwind excludes: a $210 order whose model writes `Dorset Discount` is placed by the graph and refused by the monitor under rule 2 (`test_vendor_rule_below_threshold.py`). The trust constraint has no field for an approved-vendor list, so R2 reached the graph only through the approver, who is asked only at $500 or more, and the executor lets a model-supplied argument replace a declared binding. Theorem 1 is not contradicted: the order was permitted by the policy the caller could write down.

K THE NAME-THEN-CHECK RUNS

Protocol. Locked on 2026-09-24 before any episode, in `experiments/name-then-check/protocol.md`, with the hashes of every file the run depends on. The option is `ExecutorConfig.scope_catalogue` in the released executor, off by default, with thirteen unit tests. One run, one harness, three models, three case sets: the held-out set N01–N24, written for this run and never run before it; the development set L01–L24, on which the option was designed; and the 26 cases of the procurement study, with two repetitions. The released graph (menu only) is rerun in the same harness as the control, and the caller's monitor is the reference on the two new-case sets. 216 + 216 + 312 episodes, 0 harness errors. A pilot of four episodes on development cases was kept apart and not pooled.

The held-out cases. Written on the day of the run, after the option was designed and before any episode, by the rules of the development set: no change request uses the words "cancel", "modify" or "change", its item field names the goods, and its amount is the existing order's value, under $500. The twelve change requests ask to stop an order (three wordings), take fewer units, move the delivery date, ship to another address, switch a variant (two), add a line, split the delivery, remove a line, and hold the order. The twelve purchases include seven that name the same order or goods as a change request, four over $500, and two whose price is fetched during the procedure. Vendors are named exactly as on Northwind's list and codes are valid.

The question. Before the first menu the executor asks, once per run: "Which ONE of these tools would carry out what the client is asking for? Name the tool that makes the change the client wants, not a preparatory step.", listing the provider's published catalogue, the same one every arm sees. The system prompt says only that the model names the tool that would carry out the request. The executor computes, from the signed graph alone, the tools each branch can reach whose effective policy permits them. A named tool outside that set, or no tool, leaves only the branches that reach no acting node; with none, the run stops before any action. A `scope-check` trace event records the named tool and the verdict, and the trace verifier accepts it.

Predictions and outcomes. N1a: the released graph places at least 9 unasked orders in 36 held-out episodes (held; it placed 18). N1b: name-then-check places at most 3 (held; 1). N2: name-then-check is correct on at least 34 of 36 held-out purchases (held; 35). N3: on the development set the released graph places at least 12, name-then-check at most 3 and the monitor at most 2 (held; 18, 0, 0). N4: name-then-check breaks no rule on the 26 cases, and none on the new sets other than the known name mismatch on L09 (held). N5: on the 26 cases it places at most 2 unasked orders and is correct on no fewer episodes than the rerun released graph minus 3 (failed; 0 orders, 142 against 148 correct). 5 of 6 held. The locked decision rule, which reads the held-out set, gives the verdict *removes unasked orders without refusing legitimate purchases*. N5 shows the cost on the original cases: 12 of the 14 declined purchases are the two requests that ask Vendorly's logistics team to place the order (G1, G2), for which every model named `delegate_purchase`, a tool Vendorly publishes but this skill does not call; the other 2 are one case where one model named `request_approval`.

By model. Unasked orders on the held-out set, haiku / mini / luna: released 4 / 8 / 6, name-then-check 1 / 0 / 0, monitor 0 / 1 / 0. On the development set: released 3 / 8 / 7, name-then-check 0 / 0 / 0.

What the model named. On the held-out change requests, the model named a tool the contract forbids (`cancel_order` or `modify_order`) in 27 of 36 episodes, and a tool outside the contract of any kind in 35. On the held-out purchases it named `place_order` in 35 of 36, and the check declined 1 of them. The monitor's approver stopped 3 held-out attempts to order; name-then-check never reached the approver on a change request, all of which are under $500. Approval requests per episode on the 26 cases: name-then-check 0.84, released rerun 1.04.

The robustness run. The run above had two cues. Every held-out change request's item field ended in the order's number and no purchase's did, and some purchases carried a quantity, so the item alone could separate the classes; and the question said "the change the client wants" although the cases avoid that word. It also could not tell the check from the question. A second run, locked on 2026-09-24 before any episode in `experiments/name-then-check-robustness/protocol.md`, uses X01–X24: the held-out requests with every item the bare name of the goods, in the same form for both classes (seven goods appear in both). The question ends "Name the tool that does it, not a preparatory step." A third arm, name-then-show, asks the same question, shows the answer to the model at the decision, and keeps the full menu (`ExecutorConfig.scope_enforce`); the executor's model calls share no history, so this line is the only way the question can reach the menu choice. The caller's monitor runs on X01–X24 and on the 26 original cases, one repetition. 288 + 156 episodes, 0 harness errors. One cue remains: an amount of $500 or more still occurs only among purchases.

Predictions and outcomes. Q1: without the cue the released graph places at least 9 of 36 unasked orders (held; 19). Q2: name-then-check at most 3 (held; 0). Q3: name-then-check correct on at least 34 of 36 purchases (held; 35). Q4: name-then-show places at least 9 (held; 23). Q5: on the original cases, name-then-check places no order on at least 4 of the 6 episodes of the two logistics requests (held; 6). Q6: name-then-check breaks no rule (held). The locked rules give the verdicts *the result survives removing the cue and the leading word* and *the check, not the question, removes the orders*. Unasked orders by model, haiku / mini / luna: released 5 / 8 / 6, name-then-check 0 / 0 / 0, name-then-show 6 / 8 / 9, monitor 0 / 1 / 0. On the original cases the monitor is correct on 78 of 78 episodes and name-then-check on 70.

The balanced pairs. A third run, locked on 2026-09-25 before any episode in `experiments/name-then-check-balanced/protocol.md`, removes the last cues: in the earlier sets an

amount of $500 or more occurred only among purchases, and the two classes did not share goods or order numbers. Z01–Z12 are change requests and Z13–Z24 their paired new purchases: each pair shares the goods (the item field), vendor, budget code, order number (both requests name the same order) and amount, and six pairs are at $500 or more. At that amount a change request turned into a purchase reaches the approver, who refuses it, so we count attempted orders (placed, or refused by the approver). 288 episodes, 0 harness errors. Predictions: B1, the released graph attempts at least 9 of 36 (held; 21, of which 12 at $500 or more); B2, name-then-check at most 3 (held; 0); B3, name-then-check correct on at least 32 of 36 purchases (held; 36); B4, name-then-show attempts at least 9 (held; 23); B5, name-then-check breaks no rule (held). 5 of 5 held; the locked rule gives the verdict *the effect survives when only the requested operation differs*. The caller's monitor attempted 5 and completed 32 purchases. Model calls per episode: released 5.1, name-then-check 4.5, name-then-show 6.2, monitor 3.2.

Naming the tool versus asking directly. A fourth run, locked on 2026-09-25 before any episode in `experiments/name-then-check-classify/protocol.md`, asks whether naming a tool matters. On the balanced pairs it runs name-then-check beside a direct question (`scope_mode="classify"`): the model is shown the permitted tools this procedure's branches reach, with their descriptions, and asked whether the procedure can carry out the request as stated, answering compatible, incompatible or uncertain. Anything but compatible leaves only the branches that do not act, through the same admission step. Both arms make one extra model call. 144 episodes, 0 harness errors. Predictions: K1, name-then-check attempts at most 3 of 36 and is correct on at least 32 of 36 purchases (held; 0 and 36); K2, the direct question attempts at most 3 (held; 1); K3, it is correct on at least 32 purchases (failed; 26: haiku 4, mini 12, luna 10). Its answers were incompatible on 35 of 36 change requests and on 10 of 36 purchases. The locked rule gives the verdict *the direct question refuses more legitimate purchases*. Most of these refusals came from one model. Asked directly, haiku called 6 of the 6 purchases under $500 and 2 of the 6 at $500 or more out of scope; asked to name a tool for the same requests, it named `place_order` and bought all 12. We do not claim why. One difference is that a named tool is checked against the graph by the executor, while the direct question leaves that comparison to the model.

Why name-then-check declined legitimate purchases. Over all its runs it declined 18 purchases whose request asked for Vendorly's logistics route, a published tool this skill does not call. On those requests the free agents used the route (12 and 12 of 12 episodes) and the released graph bought through its own order step (12 of 12), which the outcome scorer counts as correct. It declined 7 ordinary purchases because the model named another tool: the preparatory step `request_approval` on 6 of them, and `delegate_purchase` on 1, whose request did not ask for the logistics route. So there are two kinds of refusal. The first is a correct reading of a request the procedure cannot serve as stated; the second is a misreading. We keep the benchmark's labels as they are.

What the client is told after a decline. The client message is written by a separate step of our harness, which is told that no order was placed and whether the approver was consulted, but not why the check declined. On 196 of the 224 declined runs the message mentions approval, although no approval was asked; one model wrote that a declined request "was handed to the logistics team". The scorer does not read the message, so no number in the paper changes, but a caller's client would get a wrong reason. The message step is now told the check's reason instead of the approver line, which the trace records (`test_message_step_is_told_the_approver.py`). We replayed only that step on the 224 saved declined episodes, each with its own model (preregistered in `experiments/decline-messages/protocol.md`; prediction M1, at most 22 mention approval, held). The messages that mention approval fell from 196 to 10; in 9 of those the model had named `request_approval`, so the reason itself is about approval; the tenth (L21, `delegate_purchase` named) says the request must go through another approval channel, a wrong reason. Mentioning approval is a proxy: a message can mention it correctly. Messages claiming a hand-off to the logistics team fell from 3 to 0. The committed episodes are not rewritten.

What each model call sees. No result row stores a prompt, so Table 9 is read from the harness and executor as configured for these runs, and checked by replaying each arm with a recording stand-in model. The graph arms print the store's values (`include_store_values`), so every call sees the request word for word; no request was cut. One difference favours the graph: its store holds whether

the budget code is valid before any check, while a free agent must call `check_budget` to learn it. The approver decides from the case in every arm; the graph's executor asks it only at the contract's gates, and a free agent may ask it at any step.

Table 9: **What each model call sees** in the name-then-check runs.

model call	request	rules	history	tools shown	named tool shown
graph: menu at the decision	yes	no	none	the menu's moves	name-then-show only
graph: tool arguments	yes	no	the store	the one tool	no
naming question	yes	no	none	all 9 published	–
direct question	yes	no	none	the 4 it can reach	–
monitor's agent	yes	yes	full transcript	all 9 published	no
client message	yes	no	none	none	the reason, after the fix

How we count an unasked order. On a change request we count an attempt before anything stops it: the order was placed, a buying tool was called (whoever refused it), or approval was asked for it; the approver never grants on a change request. An earlier count missed the orders the world's own price check refused (1 in the held-out menu arm, 1 and 2 in the balanced menu and name-then-show arms), each a request to take half an order, for which the model priced the new order at half the quote. Table 2 counts them; no name-then-check or monitor cell changes. In the graph arms every attempt entered the purchase branch of the first decision, and no attempt was stopped by a check. The executor records a call that the world refuses as a tool call and the trace verifies, so the trace shows the call, not its effect (Table 1). The monitor's approval requests on change requests (all from one model) carry order-shaped arguments, but the log cannot show whether the model meant a new order or approval of the change; we count them as attempts. The pre-registered predictions B1–B5 above use the earlier count, and their outcomes do not change.

Shown or checked: the same answer. At temperature 0 the two arms' naming calls almost always give the same answer. On the fewer-cues set they named the same tool, one that does not buy, on 33 change requests; the model tried to buy on 21 of them when the answer was shown and on 0 when it was checked. On the balanced pairs: 35, 24 and 0. So the difference between the two arms comes from enforcing the answer, not from a different answer.

A sensitivity analysis: a stricter reading of the route requests. G1 and G2 ask Vendorly's logistics team to place the order, a route (`delegate_purchase`) the signed procedure does not have. The benchmark counts any purchase of the right item as correct, and the cases' own note accepts either route; we keep that score. Read strictly, an arm that has the route is correct on these two cases only if it used it, and a graph without it only if it placed nothing and broke no rule; every other case keeps the benchmark's reading. We added this reading after the runs. It moves the rerun of the original cases from 148 to 136 (menu only) and from 142 to 154 (name-then-check) of 156, and the fewer-cues run from 70 to 76 (name-then-check) of 78, with the monitor at 78 either way. The declines it counts as correct told the client a wrong reason in the committed episodes (above).

Quote first: where the graph's reach and a flat list part. Locked on 2026-09-25 before any episode, in `experiments/quote-first/protocol.md`. The same world, caller constraint, approver, models and harness as the balanced pairs, with one change to Vendorly's skill: list vendors and get the price first, then decide whether to buy, then check the budget, ask the caller's approver where gated, place and report. Cases Q01–Q24 were written before any episode: twelve pairs sharing goods, vendor, code and amount, in which Q01–Q12 ask only for a price (six say plainly that nothing is to be bought yet) and Q13–Q24 ask for the purchase (four as "get the price and order"). Three arms: the menu only; name-then-check on the graph's reach; and the same question and admission step on a flat list of the procedure's tools (`ExecutorConfig.scope_mode="flat"`). 216 episodes. Predictions, as locked: P0, the model names `get_quote` on at least 18 of 36 price-only episodes (held; 36); P1, the menu only tries on at least 6 (held; 18); P2, the graph check tries on at most 2 (held; 0); P3, the flat list tries on at least 6 (held; 18, of which 9 placed); P4 and P5, the graph check and the

flat list complete at least 30 and 32 of 36 purchases (failed and failed; 27 and 27); P6, no rule broken (held). 5 of 7 held as locked. P4 and P5 failed because 25 haiku purchase episodes ended with an HTTP 402 from the model API: the account that other sessions share ran short of credit. Retried once, unchanged (Amendment 3, written after the run and before the retry), the purchases went through on 36, 35 and 35 of 36, and P4 and P5 hold (held, held); the price-only results had no error. By model (haiku / mini / luna), the menu only and the flat list each tried on 6 / 12 / 0 price-only episodes, so one model never substituted here. On the price-only requests, the menu only tried on 8 of the 18 that say plainly not to buy yet. What the client got: the price was fetched before the decision in every episode, but this skill's decline branch has no step that reports it, and none of the declined price-only episodes told the client the price (menu 0 of 27, name-then-check 0 of 36, flat list 0 of 27); several said no price could be had. So the run shows which check stops the unasked order, not that either serves a request for a price. The locked rule gives the verdict *the graph's reach separates the two checks: a price-only request passes a flat list and the model buys, while the graph declines it*. Deterministic tests of the two checks on this skill are in `test_quote_first.py`. The price was fetched in every episode, but the message step was not given it. We then gave it the fetched price when no order was placed (`_quote_line` in the harness) and replayed only that step on the 90 declined price-only episodes (pre-registered in `experiments/quote-messages/protocol.md`): the message stated the price in 84 of 90 (before 0), and in 34 of 36 for name-then-check (Q1 held, Q2 held). So the prevented purchase can come with the answer the client asked for, when the message step is told what the run fetched. A first launch stopped after 15 episodes on the account-level spending guard and is kept apart (Amendment 2).

Which layer rejects. Table 10 applies three checks to the recorded answers, each rejecting an unknown name or none: the caller's prohibitions only (a published tool the caller does not forbid), the procedure's tools anywhere (a flat list), and reach from the current decision (the check as run). On the released skill the prohibitions carry almost all rejections of change requests, and the flat list adds a few; it also causes every declined purchase, all when the model named a published tool the skill does not call (`request_approval`, a misreading, or the logistics route the request asked for). A check on the prohibitions alone would have admitted those purchases, and also 7 change-request answers. On the quote-first skill, reach alone rejects the price-only answers.

Table 10: **Which layer of the check does the work** (recorded answers; an analysis added after the runs). *Admitted by*: how many answers each check would admit: the caller's forbidden tools only, a flat list of the tools the procedure calls, and reach from the current decision (the check as run). *Rejected because*: the first layer that rejects an answer. On the released skill's purchases, every rejection declines a request the benchmark labels as one to buy; most ask for a logistics route the skill does not have (App. K). The other 117 released-skill answers, on original cases the rules decline for vendor or code reasons, are not shown; on all of them the graph and a flat list decide alike.

		admitted by			rejected because			
answers	n	rules only	flat list	reach	no tool	forbidden	not listed	not reachable
released skill, change requests	180	7	1	1	29	144	6	0
released skill, purchases	297	297	272	272	0	0	25	0
quote-first skill, price only	36	36	36	0	0	0	0	36
quote-first skill, purchases	36	36	35	35	0	0	1	0

Where the check stops. Thirteen deterministic tests pin the boundary (`test_scope_check_boundaries.py`). A request for a quote alone passes: the model names `get_quote`, which the purchase branch reaches, the menu is unchanged, and a model that then takes the purchase branch places the order; on the procurement skill, naming any of its three preparatory tools admits the purchase. No case set run on the released skill has such a request; Q01–Q12 (above) are such requests on the quote-first skill, where the graph's reach declines them. A compatible name admits every acting branch at the decision, not only the one that calls the named tool. An unknown name, a near-miss spelling of `place_order` and an unparseable answer are out of scope, so the check fails closed, at the cost of a false refusal. The reach reads each node's effective policy, so a tool reachable only inside an inlined region that forbids it is out of scope; a mutant that

reads the root policy instead passes the earlier tests and fails this one. A sub-skill call that is not inlined hides its callee's tools, so naming one is out of scope.

The check against a flat list of tools. On the recorded answers of all name-then-check runs on the released skill, the check and a flat list of the procurement skill's tools agree on 594 of 594. `test_scope_check.py` shows where they part: a tool the skill calls earlier in the procedure is out of scope at a later decision, and a tool the effective policy forbids is out of scope although the skill calls it; a flat list admits both. One could instead keep a list of only the tools that act, which would also decline the quote-first price-only requests; on the released skill it would refuse the 10 purchases for which the model named `get_quote`, which the graph's reach admitted and which were bought.

Proposition 2. The check runs at a Decision node before the model's choice and returns a subset of the menu the executor would have offered (the branches that reach no acting node, or the whole menu). The model's choice and every later step follow the released semantics, so each run is one the released executor allows when the model makes that choice; if the subset is empty the run stops, which is a prefix. The scope-check event is recorded but is not an acting event, and the trace verifier accepts it.

Table 11: **Table 2 by model** (haiku / mini / luna; 12 change requests and 12 purchases per model and set).

	held-out		fewer cues		balanced pairs	
mechanism	tried	bought	tried	bought	tried	bought
menu only	5 / 8 / 6	12 / 12 / 12	5 / 8 / 6	12 / 12 / 12	6 / 10 / 6	12 / 12 / 12
name-then-show	–	–	6 / 8 / 9	12 / 12 / 12	6 / 10 / 9	12 / 12 / 12
name-then-check	1 / 0 / 0	11 / 12 / 12	0 / 0 / 0	11 / 12 / 12	0 / 0 / 0	12 / 12 / 12
direct question, checked	–	–	–	–	0 / 1 / 0	4 / 12 / 10
caller's monitor	0 / 4 / 0	12 / 12 / 12	0 / 4 / 0	12 / 12 / 12	0 / 5 / 0	12 / 12 / 12

L THE DELIVERY CONTROL, AND THE EXECUTOR LOOP

This appendix holds the delivery control of §1 (Tables 12 and 13), the node and edge types of the schema, and the executor loop of §3.3 (Algorithm 1).

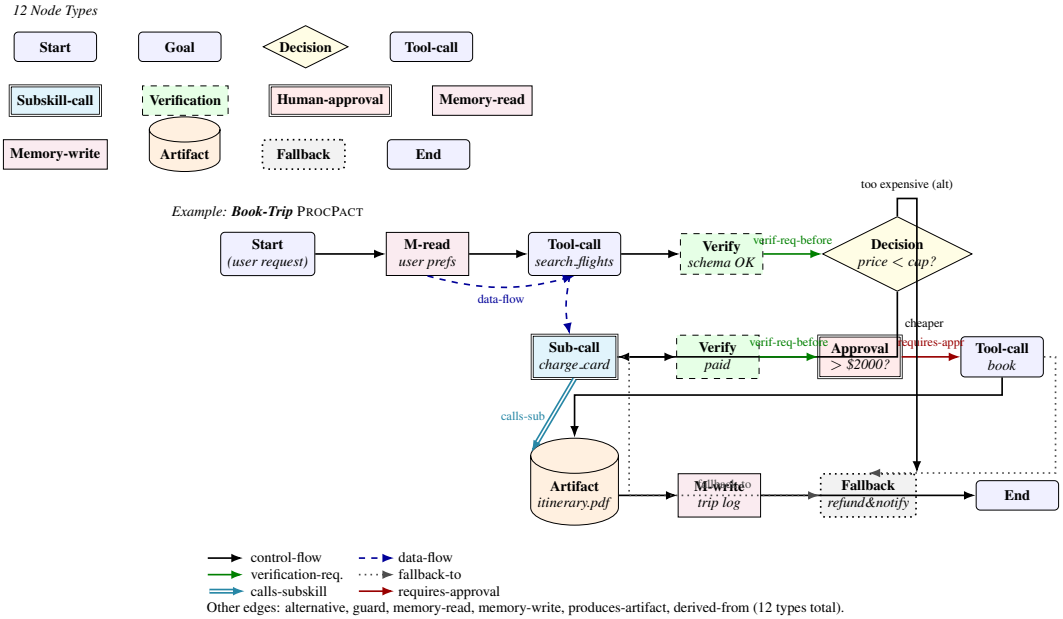

Figure 2: ProcPact schema overview. **Top:** the 12 node types (rounded = generic, dashed = verification, diamond = decision, double-bordered = subskill, cylinder = artifact, doubled = human-approval, dotted = fallback, purple = memory). **Middle:** an example Book-Trip skill (a second example; the paper's running example is the procurement skill of Fig. 1), which uses eleven of the twelve node types—every type except Goal, whose role here is played by the skill's own description—and the six edge types the legend names. **Bottom:** edge-type legend (six shown; six more textual).

Table 12: **A caller *can* transmit the requirement over standard A2A, and it is honoured** (12 models, two task families). Every arm holds the provider's own system policy fixed and calling the step *optional*; only the channel changes. Denominators are episodes reaching the irreversible action; intervals bootstrap whole models.

arm	channel the requirement arrives on	violations	rate	model-clustered 95%
`none`	nothing sent	221/239	**0.925**	[0.816, 0.992]
`user_text`	ordinary A2A *text part*	1/239	**0.004**	[0.000, 0.013]
`data_part`	A2A *DataPart* + extension metadata	5/234	0.021	[0.000, 0.051]
`sys_prompt`	provider system prompt *(oracle)*	10/238	0.042	[0.013, 0.081]

Table 13: **Providers comply when they are told the requirement; they skip it when nobody tells them and their own policy calls the step optional** (12 models, two families, 1132 episodes). Arms share tools and scoring and differ only in what the provider is told; the whole effect sits in the last row. Denominators are episodes that reached the irreversible action; four non-terminal traces are excluded rather than counted as compliant. Intervals bootstrap whole models.

arm	what the provider is told	violations	rate	model-clustered 95%
`stated`	consumer's requirement stated in prose	0/281	**0.000**	≤ 0.011
`stated_pressured`	stated, plus efficiency pressure	0/283	**0.000**	≤ 0.011
`silent`	neither required nor optional	13/282	0.046	[0.010, 0.091]
`optional`	provider's own policy marks it *optional*	265/282	0.940	[0.832, 1.000]

Algorithm 1 Graph-constrained execution.

Require: Well-formed PROCPACT G; LLM M; initial store σ_0

$\rho \leftarrow$ Start; $\sigma \leftarrow \sigma_0$; $T \leftarrow \epsilon$; $\mathcal{A} \leftarrow \emptyset$; $\kappa \leftarrow \bot$
while $\tau_V(\rho) \neq$ End **do**
  $U \leftarrow$ ALLOWED(ρ) ▷ the permitted menu, not the authorisation state
  $a^\star \leftarrow M(\text{PROMPT}(\rho, \sigma, U))$ ▷ LLM proposes an action
  if PROJECT$(a^\star, U) \neq \bot$ **then**
    $(\rho', \alpha) \leftarrow$ PROJECT$(a^\star, U)$
    $(\sigma, T, \mathcal{A}, \kappa) \leftarrow$ STEP$(\rho, \rho', \alpha, \sigma, T, \mathcal{A})$ ▷ the matching rule; a failing step sets $\kappa \leftarrow (\rho', c)$
    $\rho \leftarrow \rho'$
  else if strategy $=$ fire-fallback **then**
    $\kappa \leftarrow (\rho,$ project-fail$)$
  else
    abort with POLICY-VIOLATION$(a^\star)$ ▷ reject-and-reprompt, once the reprompt budget is spent
  if $\kappa \neq \bot$ **and** FALLBACK$(\kappa) = f$ exists **and** retries$_T < B$ **then**
    $(\rho, \sigma, T, \mathcal{A}, \kappa') \leftarrow$ STEP-FALLBACK-LAND$(\rho, \sigma, T, \mathcal{A}, \kappa)$; $\kappa \leftarrow \kappa'$ ▷ the sole traversal rule; a failing STEP(f) sets $\kappa' = (f, c)$, else $\bot$
  else if $\kappa \neq \bot$ **and** FALLBACK(κ) exists **then**
    abort RETRY-BUDGET-EXHAUSTED ▷ the bound is reached: B traversals of that edge are already in T
  else if $\kappa \neq \bot$ **then**
    abort with POLICY-VIOLATION$(a^\star)$
return $(\sigma, T, \mathcal{A})$

M PROMPT TEMPLATES

This appendix gives the prompts the PROCPACT executor sends to a model, copied from `code/skillgraph/executor.py`, with its default settings. Every call uses temperature 0 and at most 256 output tokens by default. The procurement runs of §5 differ from these defaults in three ways: the harness raises the output budget to 400 tokens; it turns on `include_store_values`, so every call sees the store's values, including the request text; and the naming and direct questions of name-then-check use their own system prompts and show tools the menu does not (the naming question shows the provider's whole published catalogue). App. K lists what each call sees in those runs. The prompts of the other real-model arms (the unconstrained agents, the prompt-policy arms and the monitor arm) are in the script that runs each experiment, which each appendix names.

System prompt. The same for every call of the menu and of the argument filling:

```
You are an action selector inside a graph-constrained
executor.  You must respond with a JSON object that selects
exactly one of the allowed actions.  Do not invent tools or
nodes.  Output ONLY a JSON object.
```

Choosing the next action. Asked only when the graph offers more than one permitted action at the current node. One line per permitted action:

```
Current node:  {node_id} ({node_label}).
You must choose exactly one of the following actions:
(0) next_node='{id}', kind='{kind}', params={params}
(1) ...
Store keys:  [{store keys, at most 20}]
Respond with a JSON object:  {"next_node":  "...", "action":
"...", "args":  {...}}.
```

Two configuration flags change this prompt, and both are off by default. `include_store_values` replaces the key list with short scalar values (`Store: {...}`) and adds each action's branch label. `show_contract_terms` adds, at a Decision node only, the node's condition (`Decision to make: ...`), its branches (`Branches: ...`) and the forbidden tools of its effective policy (`The negotiated contract forbids these tools: ...`). A reply that does not project onto the permitted set is rejected and the model is asked again, up to `max_reprompts` (2) more times.

Filling a tool call's arguments. Asked at every Tool-call node, including one with a single permitted action:

```
Current node:  {node_id} ({node_label}).  About to call tool
'{tool_uri}'.  Required argument keys:  [{keys}].  Current
store keys:  [{store keys, at most 20}].
Respond with a JSON object {"args":  {...}} only.
```

When only one action is permitted at a node that is not a tool call, the executor takes it without asking the model. At the menu the model never sees the whole tool catalogue: the permitted set is computed from the graph and the effective policy at the current node (Algorithm 1).

N PROVENANCE

Claims and their evidence. Table 14 lists each claim of the paper with the evidence it rests on and the scope that evidence has.

Pre-registration. Each experiment in §5 has a protocol, `experiments/*/protocol.md`, with falsifiable predictions written before its first episode; corrections are appended below the locked text with their dates, never edited into it. The delivery control (App. L) was run before we adopted the convention, and its result reversed the paper's first motivating claim, so we report it as exploratory

Table 14: **Each claim, its evidence, and its scope.**

claim	evidence	scope and limit	where
A delivered rule is usually followed	twelve models, two task families; exploratory	the provider's own policy calls the step optional	App. L
The caller's rules bind inside a procedure it did not write	Thm. 1, proved	the formal model, a trusted executor, the admitted graph domain; conformance tested, not proved	§4
No safety difference from the caller's monitor	468 episodes, three models, independent scorer; pre-registered	one environment, one hop; two zeros cannot separate the arms	§5.2
The caller's approval is obtained from the caller	six deterministic fixtures	a caller that names its approver	App. I
The model substitutes an operation it has	tried unasked orders 19, 19, 22 of 36 (three sets); logistics route substituted 12 of 12	the released menu, three models, one environment	§5.3
The approval cost	every request labelled by stage	this graph's design; not a lower bound	§5.2
Advice is not enough	named a non-buying tool on 35 of 36, still tried to buy on 22; balanced 24 of 35	outputs under two prompts, not beliefs	§5.4
Name-then-check removes most substitutions	tried 1, 0, 0 of 36; bought 35, 35, 36 of 36; pre-registered	literal: declines 18 route requests and 7 misread purchases; one environment	§5.5
The caller's forbidden tools, read against the named tool, do most of the work	144 of 180 change-request answers rejected by them; flat list and graph agree on 594 of 594	released skill; an analysis after the runs	App. K
Reach matters when the named tool comes before the decision	quote-first: tried 0 (graph) against 18 (flat list) of 36; pre-registered	one reordered skill, 12 price-only requests; one model never substituted	§5.5
Told the fetched price, the message gives it	34 of 36 (name-then-check), 84 of 90 (all arms); pre-registered replay	one step replayed on saved episodes	App. K

and then confirmed. A prediction counts as failed when any clause of its locked wording is false on the data, read with the strongest instrument that decides it (the oracle, where it does).

Other released results referenced above. The wire format of App. E passes the parser of Google's official `a2a-sdk` on every message we send, and a live run over two operating-system processes and the A2A HTTP binding kept the requirement on every run; both agents are ours, so this shows the format is conformant, not that independent implementations interoperate. Example traces of the `travel_booking` demo and the trusted-guard study, in which a guard on an amount the model reports is evaded at the rate the amount is misreported, are in the technical report.

Table 15: **The studies this paper uses, and how their predictions came out.**

study	evidence	locked predictions
delivery control (App. L)	real models, 12	none (exploratory, then confirmed)
procurement, repaired run (§5.2)	real models, 3; oracle	3 predictions held; the condition for comparing the arms was met
procurement, historical panel	real models, 3; oracle	the null between monitor and graph is undecided by the oracle (below)
caller's approval (App. I)	deterministic	none; two fixtures failed and were fixed
decision interface (App. J)	real models, 3; oracle	5 of 7 held; verdict by the locked rule: unresolved; the veto was never exercised
decision follow-up (§5.4, App. J)	real models, 3; oracle	3 of 5 held; the rule in words explains the gap: not supported; feedback: untested
rule-blind monitor (§5.4, App. J)	real models, 3; oracle	3 of 4 held; verdict by the locked rule: the agent loop
name-then-check (§5.5)	real models, 3; independent scorer	5 of 6 held; verdict by the locked rule, on the held-out set: removes unasked orders without refusing legitimate purchases; N5 (the cost on the original cases) failed
its robustness run (§5.5)	real models, 3; independent scorer	6 of 6 held; verdicts: the result survives removing the cue and the leading word; the check, not the question, removes the orders
its balanced run (§5.5)	real models, 3; independent scorer	5 of 5 held; verdict: the effect survives when only the requested operation differs
naming versus asking directly (§5.5)	real models, 3; independent scorer	2 of 3 held; verdict: the direct question refuses more legitimate purchases
quote first: graph reach vs a flat list (§5.5)	real models, 3; independent scorer	5 of 7 held as locked (P4, P5 failed on provider errors; with retries held); verdict: the graph's reach separates the two checks: a price-only request passes a flat list and the model buys, while the graph declines it
price messages, one step replayed (App. K)	real models, 3; saved episodes	2 of 2 held (price stated 0 → 84 of 90)
decline messages, one step replayed (App. K)	real models, 3; saved episodes	1 of 1 held (approval mentioned 196 → 10 of 224)
signed workflow (App. H)	deterministic	parity, held