COARSE-GRAINED DYNAMICS OF MULTI-AGENT COLLECTIVE AUTORESEARCH

Anonymous authors
Paper under double-blind review

ABSTRACT

Multi-agent collective AutoResearch generates research findings through shared knowledge and collaborative exploration. As the number of agents increases, collective behavior arising from complex interactions becomes increasingly important, making system-level performance and scalability difficult to predict and motivating an effective dynamical description. We model the effective dynamics of collective behavior from a statistical physics perspective using coarse-graining methods and formalize a multi-agent collective AutoResearch system as a stochastic growth process in which agents are coupled through a growing knowledge hypergraph. We describe each agent by the phase-type distribution of its action duration and outcome, specify local hypergraph interaction rules, and average over the agents to obtain the coarse-grained dynamics. This formulation can model and predict collective behavior using measurements on single agents, mapping different research agents onto their phase-type distributions and different frameworks onto interaction rules and correction terms. We evaluate the model's predictive ability in formal theorem-proving tasks: from the phase-type distribution of a single agent, it predicts how the completion time of many agents sharing a record changes with their number, in close agreement with experiments. This provides a dynamical framework for understanding and predicting the scalability of collective AutoResearch systems.

1 INTRODUCTION

AutoResearch systems now carry out the research cycle with little human input (Lu et al., 2026a; Gottweis et al., 2026), often with many agents building on one shared record of verified results (Bianchi et al., 2026; Zhang et al., 2026). Whether more agents help is not known in advance. They can speed up the growth of the record, but they also duplicate work and race on the same questions. At matched compute, adding agents can lower performance (Kim et al., 2026). In practice the answer is found by running the collective at each size of interest, at a cost that grows with population size, and the resulting conclusions do not automatically transfer across agents, tasks or sharing rules. Scaling laws for AI scientists have been conjectured (Zhang et al., 2025); we ask whether they can be derived and predicted from measurements on single agents.

Existing scaling studies do not directly address this prediction problem for a growing body of interdependent research results. On fixed tasks, voting over repeated samples raises accuracy, though not always monotonically (Li et al., 2024; Chen et al., 2024), collaboration networks improve with size (Qian et al., 2025), and the gain of collaboration must be separated from that of a larger budget (Fan et al., 2026). Statistical physics derives population laws of language-model agents from single-agent behaviour (Flint et al., 2026; De Marzo et al., 2026; Tanaka, 2026; El et al., 2026), but for conventions and opinions on a fixed set of states. In research, as in human science (Rzhetsky et al., 2015; Chu & Evans, 2021), agents interact through results: a result can supply a premise another agent needs, end ongoing attempts on the same question, and enlarge the record they must search. Predicting collective growth therefore requires a model of how individual action times and outcomes couple through the evolving record.

To address these issues, we first model collective AutoResearch as agents posing and researching questions on a shared, growing knowledge hypergraph (Figure 1). We identify the agents' local actions and how their interactions through the hypergraph generate collective knowledge growth

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.

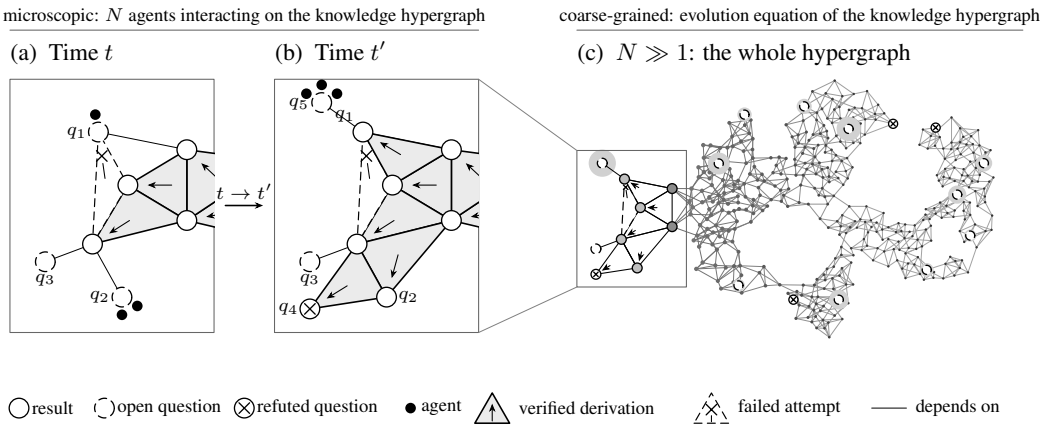

Figure 1: The knowledge hypergraph and its field limit. (a) A patch of the hypergraph at time t with N agents on it: arrows point to derived results, lines join questions to the results they were posed on, dots are agents at the questions they research. An attempt at q_1 has failed; nobody works on q_3. (b) The same patch at $t' > t$: q_1 and q_2 derived, q_4 posed and refuted, q_5 posed, agents moved on. (c) The whole hypergraph for $N \gg 1$, with (b) boxed; far from it only nodes and edges are drawn. Coarse-graining: averaging over the agents replaces them by a few collective variables on the hypergraph, drawn as shades and halos, which set the hazard at which each question is established (Section 4.2).

and competition. We then coarse-grain this microscopic description: each agent is modeled by the phase-type distribution of its action duration and outcome. Combining these distributions with the interaction rules, we derive collective dynamics and conditional laws for knowledge growth, population scaling and efficiency (Sections 3–5). Practically, we fit the phase-type distributions from single-agent runs and use them to compute collective growth curves and predict completion times at different population sizes (Section 6).

Our experiments follow two steps. We first measure the phase-type distribution of single agents: for Claude Code, Codex and the NewtonBench agent, on lemmas of the Natural Number Game in Lean (Buzzard & Eugster, 2023), hidden physical laws in NewtonBench (Zheng et al., 2025) and implications between equational laws from the Equational Theories Project(ETP) (Bolan et al., 2025), a phase-type distribution fits the duration and outcome of their actions. We then fit the single-agent performance distribution on ETP tasks, and insert the distribution into the coarse-grained dynamics to predict the multi-agent completion time of many such agents sharing a record changes with their number. Nothing is fitted to runs with more than one agent, and the prediction agrees closely with the measured completion times at every population size tested.

Contributions. In this paper, our contributions are: (i) A model of collective AutoResearch linking local agent actions to coarse-grained collective dynamics. (ii) Conditional laws for knowledge growth, population scaling and efficiency. (iii) Empirical tests of the agent description and collective dynamics, identifying the reduction's range of validity.

2 Background and related work

AutoResearch. We use this term for AI agents that pose questions, attempt and check answers, and record the results with limited human intervention. Such systems run chemistry experiments, search for mathematical constructions, write papers and generate hypotheses later tested experimentally (Boiko et al., 2023; Romera-Paredes et al., 2024; Lu et al., 2026a; Gottweis et al., 2026; Ghareeb et al., 2026), and many agents now share verifiers, public leaderboards and records of failed attempts (Bianchi et al., 2026; Wang, 2026). These works evaluate what one system produces; we ask how the results of many agents accumulate, and how this depends on their number.

Collective behavior and coarse-graining. A system of many interacting units shows collective behavior when the whole has properties that no unit has and that do not follow from adding up the

units (Anderson, 1972; Castellano et al., 2009; te Vrugt et al., 2022); for language-model agents it has been modelled for opinions, conventions, team size and interaction structure (El et al., 2026; Tanaka, 2026; Bertalanič & Fortuna, 2026; Lu et al., 2026b). Coarse-graining describes such a system by a few macroscopic variables, as in aggregate models of knowledge shared by humans and language models (Nettasinghe & Zhao, 2026). Their equations close only approximately: the discarded details still act on the retained variables and must be represented through them (Givon et al., 2004).

3 A MICROSCOPIC MODEL OF COLLECTIVE AUTORESEARCH

We model N agents coupled to one another on a growing knowledge hypergraph. The agents pose questions, which adds nodes to the hypergraph, and research them, which establishes nodes through verified derivations; the hypergraph sets what each agent can do, and the agents make it grow. The structure of knowledge depends on the philosophical setting, which is discussed together with the modelling assumptions in Appendix A. We first define the knowledge hypergraph and admissible actions, then describe individual behaviour using phase-type distributions and couple these descriptions through shared knowledge. This provides the microscopic process used in Section 4 to derive collective dynamics.

3.1 KNOWLEDGE HYPERGRAPH AND AGENT ACTIONS

The fixed hidden directed hypergraph $\mathcal{H}^\star = (V^\star, E^\star)$ represents possible statements and valid derivations using hyperedges with a single conclusion (Gallo et al., 1993). A node q is a statement, and a hyperedge (P, q) means that the finite premise set $P \subseteq V^\star$ jointly suffices for $q \in V^\star$. This structure specifies validity independently of discovery. Agents cannot inspect it, and we do not require reconstructing it from their logs. Their public knowledge hypergraph is

$$\mathcal{H}_t = (V_t, E_t, \ell_t), \qquad \ell_t : V_t \longrightarrow \{\text{true}, \text{false}, \text{undetermined}\}, \tag{1}$$

where $V_t \subseteq V^\star$ contains public statements and $E_t \subseteq E^\star$ their published, verified derivations. Labels express what the knowledge hypergraph establishes: new statements are undetermined; true statements are the initial true seeds and conclusions obtained by repeatedly applying recorded edges whose premises are all true; false statements require checked refutations.

Agent i accesses the public knowledge hypergraph and its private hypergraph $h_i(t)$ of unpublished findings; labels in its view are computed from their combined contents. Its knowledge base $K_i(t)$ holds at most C statements available for immediate use, including undetermined hypotheses. We write $z_i(t)$ for its working state: the action it is performing, the question it is working on together with the statements it has loaded for it, and whatever internal state its own procedure keeps. Together these variables define the system state

$$\mathcal{S}_t = (\mathcal{H}_t, (h_i(t), K_i(t), z_i(t))_{i=1}^N). \tag{2}$$

At the model level, an agent has two actions: *posing a question* and *researching a question*. Conjecturing and proving in Dong & Ma (2025) provide one concrete example of this distinction. Posing proposes a new statement for investigation, without establishing it. Research selects an undetermined target q and seeks a checked derivation or refutation. Loading knowledge, reasoning, using tools, submitting a candidate and giving up are steps within this action. When an action ends, the agent chooses the next one, posing with probability α and research otherwise; α is constant in the homogeneous baseline. A derivation from usable knowledge is possible only if

$$g_i(q, K_i) = \mathbf{1}\{\exists (P, q) \in E^\star : P \subseteq K_i\} = 1. \tag{3}$$

This condition is hidden from the agent and does not ensure that it finds a valid proof. Intermediate statements to be recorded or shared enter through the posing action. Section 3.2 specifies how these actions update the hypergraphs and knowledge bases.

3.2 INTERACTING DYNAMICS

An agent interacts with the hypergraph in two ways. It *reads*: to start research it selects an open question in its view, and during research it loads statements from its view into K_i, evicting a held

statement when K_i is full. It *writes*: when an action ends, it adds its output to its private hypergraph h_i. Posing can produce a new undetermined statement. Research produces an accepted derivation (P, q) with $P \subseteq K_i$, a failure case listing the statements loaded for q, or nothing if the agent gives up; an accepted derivation also adds q to K_i. Each new output is published with probability σ, or σ_{f} for a failure case, and enters $\mathcal{H}_t$ together with the private derivations it relies on. Labels are then recomputed in the public knowledge hypergraph and in every view, and research whose target has become true or false is cancelled. In this baseline, agents are coupled only through publication; extensions such as cooperation among the researchers of one question add other couplings.

For each type of action, we describe its duration and terminal outcome by an outcome-marked phase-type (PH) law. For an action $r \in \{\mathrm{pose}, \mathrm{res}\}$ of agent i we write

$$(T, J) \sim F^{\mathrm{PH}}_{i,r}(\cdot \mid I_i), \tag{4}$$

where T is the duration, J the terminal outcome and I_i the inputs listed in assumption A3. For posing, J is a new question or no output; for research, J is a result, a rejected submission or giving up. The PH representation is only a convenient stochastic description of this law: it assigns no meaning to the agent's internal reasoning steps, and different agents can have entirely different phases. It allows durations that are not exponential while keeping the whole system a continuous-time Markov process, and it can approximate any joint law of a positive duration and finitely many outcomes (Bladt & Nielsen, 2017). In this representation, the working state z_i consists of the action, its target and the current phase ϕ_i of the chain; the phase retains whatever the rates depend on, such as the statements loaded so far. Loads occur on some phase transitions.

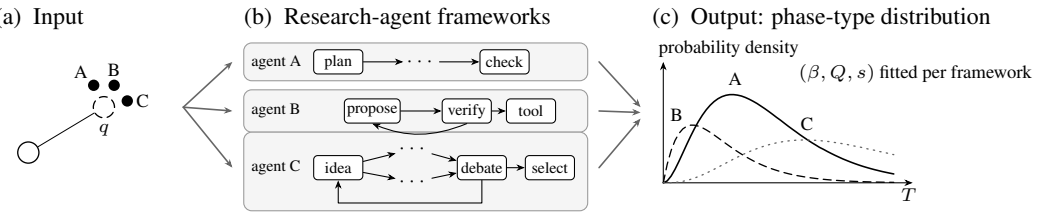

Figure 2: In the coarse-grained model, an agent is described by a phase-type distribution regardless of its designed pipeline. (a) Three agents A, B and C research the same open question q. (b) Examples of agents following their own designed procedure. (c) For each framework the model fits a phase-type distribution F^{PH} with parameters (β, Q, s); the curves show the probability density of the time T at which the action returns its output.

The PH representation specifies the rates at which events occur, while the read and write rules specify how each event changes the state. We combine the two through the generator, written as in stochastic graph rewriting (Behr et al., 2016; 2020): a sum over events of the rate of the event times the change it causes. Given the hidden hypergraph, let $\mathcal{E}(S)$ be the set of events possible in state S, let $a_e(S)$ be the rate of event e, and let S^e be the state just after it, which may be random. For every function f of the state,

$$(\mathcal{L}f)(S) = \sum_{e \in \mathcal{E}(S)} a_e(S)\, \mathbb{E}\big[f(S^e) - f(S)\big]. \tag{5}$$

Each event belongs to one agent and is of one of two kinds: a phase transition within an action, which may load a statement into K_i, or the end of an action, which carries out the write and its consequences described above. The agent, and every agent whose research was cancelled, then starts a new action at once. The full rules, the phase-type details and these extensions are given in Appendix B. Since $\frac{\mathrm{d}}{\mathrm{d}t}\mathbb{E}[f(\mathcal{S}_t)] = \mathbb{E}[(\mathcal{L}f)(\mathcal{S}_t)]$, Section 4 derives the collective balance equations from this generator.

4 From agent behaviour to collective dynamics

We derive collective balances and construct an effective process on the knowledge hypergraph from the action laws of Section 3. We follow a standard route to collective variables in statistical physics: start from the quantities of interest, write their rates of change with the generator, add the variables

these rates depend on, and, when this chain does not end, close it by separating time scales or by assuming independence, as in the projection-operator method (Zwanzig, 1961; Mori, 1965). Section 4.1 takes the exact steps, ending with the elimination of the phases, which leaves a memory kernel, and Section 4.2 closes the balances into a process on the knowledge hypergraph. We work in the *public setting*: all outputs are public, only true premises are used, and refutation and pooling are excluded.

4.1 Collective Variables

Let M and M_{open} count public established and open statements,

$$M(t) = \#\{q \in V_t : \ell_t(q) = \mathrm{true}\}, \qquad M_{\mathrm{open}}(t) = \#\{q \in V_t : \ell_t(q) = \mathrm{undetermined}\}, \quad (6)$$

and M_{ready} the open questions with an established sufficient premise set, the only ones that can be established. Only results increase M. A researcher of q in phase ϕ of its PH law (β, Q, s) (Figure 2) ends with a result at the rate $s_{q\phi}$, the entry ϕ of s; with $X_{q\phi}$ such researchers, results appear on q at the exact rate $\lambda_q = \sum_\phi X_{q\phi} s_{q\phi}$, and $\lambda = \sum_q \lambda_q$. The expected growth is therefore exact:

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathbb{E}[M] = \mathbb{E}[\lambda]. \quad (7)$$

The growth balance is not closed: the success flux depends on the researchers' phase occupancies. Eliminating these phases introduces memory (Appendix C.1). When the entries change slowly, only two functionals of the action law remain. For fixed inputs and independent exponential cancellation at rate κ, the action law gives the success probability per entry and the mean occupation time:

$$L_{\mathrm{s}}(\kappa) = \mathbb{E}\big[e^{-\kappa T}\mathbf{1}_{\{J=\mathrm{result}\}}\big] = \beta(\kappa\mathbf{I} - Q)^{-1}s, \qquad \tau_{\mathrm{R}}(\kappa) = \beta(\kappa\mathbf{I} - Q)^{-1}\mathbf{1}, \quad (8)$$

Here T and J are the duration and outcome of Eq. (4), and τ_{R} is the mean time until completion or cancellation. The factor $e^{-\kappa T}$ is the probability that an action is not cancelled before it ends at time T. At $\kappa > 0$, L_{s} depends on the joint law of duration and outcome. Section 4.2 combines these quantities with entry rates to construct the effective dynamics.

4.2 The coarse-grained dynamics

We treat the coarse-grained dynamics at two orders: a baseline and its first-order corrections.

Zeroth order: the baseline. In the baseline, all agents are alike, choose open questions uniformly and keep no memory between actions. Other researchers cancel an action at a rate κ_q that does not depend on its phase. Appendix C.1 lists all the assumptions.

Let Z record each statement as unrevealed, open or established. Its rates follow in four steps. (1) An agent starts a new action as soon as the last one ends. It poses with probability α and otherwise researches an open question. One such cycle lasts on average $\bar\tau = \alpha\tau_{\mathrm{pose}} + (1-\alpha)M_{\mathrm{open}}^{-1}\sum_{q\ \mathrm{open}}\tau_q$, where τ_{pose} is the mean time of a posing action and τ_q the mean time of research on q. If q is ready, $\tau_q = \tau_{\mathrm{R}}(\kappa_q)$ of Eq. (8). If q is not ready, no one can establish it, so $\tau_q = \tau_{\mathrm{R}}(0)$ of its own law, which has no exit to a result. (2) Each agent starts $(1-\alpha)/\bar\tau$ research actions per unit time, spread evenly over the open questions. So q receives entries at the rate $\lambda_q^{\mathrm{in}} = N(1-\alpha)/(\bar\tau\, M_{\mathrm{open}})$. (3) Each entry into q succeeds with probability $L_{\mathrm{s},q}(\kappa_q)$ of Eq. (8), averaged over the holdings at entry. This probability depends on the whole law, not only on its mean. So the research law enters only through τ_{R} in step (1) and L_{s} in step (3). (4) If the phases relax faster than Z changes, we take the resulting mean rates as the rates of a Markov jump process on Z. This is an effective model, not an exact reduction of the agent process (Appendix C.1). An open ready question q is established, $Z \to Z^{+q}$, or an unrevealed statement u is posed, $Z \to Z^{+u}$, at the rates

$$\begin{aligned} W(Z \to Z^{+q}) &= \lambda_q^{\mathrm{in}} L_{\mathrm{s},q} = \frac{N(1-\alpha)}{\bar\tau(Z)\, M_{\mathrm{open}}(Z)}\, L_{\mathrm{s},q}(\kappa_q; Z), \\ W(Z \to Z^{+u}) &= \frac{N\alpha}{\bar\tau(Z)}\, p_u^{\mathrm{pose}}(Z), \end{aligned} \quad (9)$$

where a posing action outputs u with probability p_u^{pose}. All other rates are zero. The cancellation rate $\kappa_q \simeq (N-1)W(Z \to Z^{+q})/N$ is the success rate of the other researchers of q. The probability

$P_t(Z)$ that the knowledge hypergraph is in state Z at time t obeys the master equation

$$\frac{\mathrm{d}}{\mathrm{d}t}P_t(Z) = \sum_{Z'}\big[W(Z' \to Z)\,P_t(Z') - W(Z \to Z')\,P_t(Z)\big]. \tag{10}$$

The first term is the probability flowing into Z, and the second the probability flowing out. The state is the knowledge hypergraph alone: agents enter only through N, α and the means and transforms of their PH laws.

First order: weak mechanisms. The baseline leaves out mechanisms such as target policies, memory and failure lists. A weak mechanism that adds no slow variable changes only three inputs of a question: the probability π_q that an entering agent chooses it, the success per entry $\ell_q = L_{\mathrm{s},q}(\kappa_q)$, and the mean time τ_q spent on it. By the renewal reward theorem, the growth rate is

$$\lambda = \frac{N(1-\alpha)}{\bar{\tau}}\sum_q \pi_q \ell_q, \qquad \bar{\tau} = \alpha\tau_{\mathrm{pose}} + (1-\alpha)\sum_q \pi_q \tau_q, \tag{11}$$

which is $\sum_q W(Z \to Z^{+q})$ when $\pi_q = 1/M_{\mathrm{open}}$. At fixed N and α, a small change of the inputs gives the relative change

$$\frac{\delta\lambda}{\lambda} = \frac{\delta\big(\sum_q \pi_q \ell_q\big)}{\sum_q \pi_q \ell_q} - \frac{\delta\bar{\tau}}{\bar{\tau}}. \tag{12}$$

The first term is the change in results per cycle. The second is the cost of time: a mechanism that raises ℓ_q but lengthens actions, such as failure lists, can lower λ. Table 1 and Appendix C.2 give the effect of relaxing each assumption. Relaxing the public setting, for example by keeping refutations private, adds a slow variable, who knows which failure, and so extends the state. First order fails in two cases. First, if many agents start on one new question together, they run a race (Appendix D.5). Second, at the start, or when one result changes many questions at once, the phases must be kept.

Keeping the phases. In that second case we keep in the state the mean number $x_c(t)$ of researchers in each phase, for each class c of questions with one action law (β_c, Q_c, s_c). Between results,

$$\dot{x}_c = x_c\,Q_c + a_c(t)\,\beta_c, \qquad \lambda(t) = \sum_c x_c(t)\,s_c, \tag{13}$$

where $a_c(t)$ is the rate at which agents reach the first phase of class c: all N agents start at $t = 0$, and an agent that ends an action or is cancelled starts again at once on a uniformly chosen open question, reaching the first phase after a fixed delay $d_c \geq 0$ unless cancelled. Results are random events at the rate λ; each cancels the researchers of every question that has become true or false (Appendix C.1). With constant entry rates, cancellation as a clock of rate κ and $\dot{x}_c = 0$, Eq. (13) gives the rates of Eq. (9), with L_{s} and τ_{R} taken over the whole action, delay included. The correction is what this stationary solution leaves out: the transient of x_c, the delay in a_c and the drop of x_c at each result. At the start $\lambda = 0$ until the first delay ends, whereas Eq. (9) gives a positive rate at once.

Physical interpretation. Write U_q, Q_q and H_q for a statement that is unrevealed, open and established. Posing is $U_q \to Q_q$. Establishing q through the route P is

$$Q_q + \sum_{v\in P} H_v \;\longrightarrow\; H_q + \sum_{v\in P} H_v. \tag{14}$$

Every result is a catalyst for later results, and no result is removed. So the coarse-grained dynamics is a catalytic reaction network that only grows (Hordijk & Steel, 2004). Different routes to the same statement are alternatives. Further physical discussion, including the bound on what can ever be established and a comparison with other growth models, is in Appendices D.1 and C.3.

5 Scaling laws

More agents can finish sooner while doing more work in total. We compare the time $T_L^{(N)}$ needed to establish L new statements with the total agent time $N\mathbb{E}[T_L^{(N)}]$, when the mean is finite (Fan et al., 2026).

More agents can shorten the wait. Twice as many agents can finish in half the time without saving total work. In the model, this holds when cancellation is switched off, each agent's way of working stays unchanged, and the questions and starting knowledge are the same:

$$M_N(t) \overset{d}{=} M_1(Nt), \qquad T_L^{(N)} \overset{d}{=} T_L^{(1)}/N. \tag{15}$$

The mean growth curves coincide when time is measured as Nt. This prediction uses average action rates; it need not hold when agents start together or when one result suddenly changes many tasks.

More knowledge can take longer to search. As the record grows, random search makes useful premises harder to find. Approximate the established count by $x(t)$. If each agent's result rate falls as x^{-d}, with fixed $d \geq 0$, then

$$x(t) \propto (Nt)^{1/(d+1)}. \tag{16}$$

once new results far outnumber the initial knowledge, provided competition is weak and the share of questions with available premises stays steady. In the search model, d counts the premises still to be found by random search. A fixed chance of finding a needed premise eventually avoids this slowdown: $d = 0$ gives linear growth.

More agents also duplicate work. Competition matters when another agent can solve the same question during a typical attempt. The single-agent PH fit gives the success probability q_{s}. Here M_{ready} counts unanswered questions with at least one sufficient premise set already established. This count comes from the task dependencies and current knowledge; when the dependencies are unknown, it must be estimated separately from the PH fit. If all agents research equally difficult ready questions and choose among them uniformly, the competition scale is estimated by

$$q_{\mathrm{s}} = \beta(-Q)^{-1}s, \qquad N^* \approx \frac{M_{\mathrm{ready}}}{q_{\mathrm{s}}}.$$

The mean attempt time cancels: faster attempts also mean faster competitors. This estimate counts one settled question per result. When a result also settles other questions through the closure, as in the ETP world of Section 6.2, we measure the rate on one agent instead. If a campaign of one agent lowers the number of open questions from $M_{\mathrm{open}}(0)$ to $M_{\mathrm{open}}(0) - L$ in the time $T_L^{(1)}$, each open question is settled at the mean rate

$$r_L = \frac{1}{T_L^{(1)}} \ln \frac{M_{\mathrm{open}}(0)}{M_{\mathrm{open}}(0) - L}.$$

Before competition sets in, Eq. (15) holds, so the other $N - 1$ agents settle the question of an agent at the rate $(N - 1)r_L$, and one cancellation is expected during an attempt of mean duration $\mathbb{E}[T]$ when

$$N_L^* = 1 + \frac{1}{r_L\, \mathbb{E}[T]}. \tag{17}$$

When each result settles one question, $r_L \approx q_{\mathrm{s}}/(\mathbb{E}[T]M_{\mathrm{ready}})$ and N_L^* returns the estimate above. Stopping unfinished work can discard useful progress, but can also free agents from long attempts unlikely to succeed. The balance depends on how long successful and unsuccessful attempts take.

Crowding can raise the work per result. When many agents race on the same questions, the benefit depends on the chance of a very fast success. Suppose the probability of success by a short time t grows as t^n. Keeping the phases tracks agents' existing progress and when they join each question. With the same fixed pool, all agents researching and no extra waiting between attempts, this still gives

$$\lambda \propto N^{1/n}, \qquad N/\lambda \propto N^{(n-1)/n}.$$

For a fixed number of such rounds, $\mathbb{E}[T_L^{(N)}] \propto N^{-1/n}$. Keeping the phases matters for the numerical factor through the timing of arrivals; the powers above stay the same in this setting. For $n = 2$, four times as many agents produce only twice as many results per unit time, doubling the work per result. Changing questions, dependencies or adding a minimum action time can change this scaling.

Some steps must wait for earlier results. If the target requires at least D_L actions in sequence, each starting after the previous result and taking at least $t_{\min} > 0$, then $T_L^{(N)} \geq D_L t_{\min}$. Extra agents cannot remove this waiting time. An average action time alone does not provide such a bound.

New questions and shared results support further work. Increasing α shifts time from answering questions to posing them, supplying future work at the cost of current research. Increasing σ makes more findings public, so other agents can use them and stop repeating completed work. When findings stay private, progress also depends on who knows which results.

6 Experiments

The two coarse-graining steps are tested on its own: Section 6.1 tests the phase-type distribution of an action, Section 3.2, on existing agents, and Section 6.2 tests the scaling laws of Section 5.

6.1 First coarse-graining: the phase-type distribution of an action

To test the first coarse-graining, we ran unmodified agents, Haiku 4.5 and Sonnet 5 in Claude Code or the NewtonBench agent and gpt-5.5 in Codex, on three kinds of task: lemmas of the Natural Number Game in Lean (Buzzard & Eugster, 2023), with or without their premises among the lemmas the agent may cite; hidden physical laws in NewtonBench (Zheng et al., 2025); and implications between equational laws from the Equational Theories Project (ETP) (Bolan et al., 2025). For each attempt we recorded its duration and whether its result passed the check. We fit a hyper-Erlang distribution, a mixture of Erlang branches each ending in a result with its own probability, to duration and outcome jointly, by maximum likelihood with the EM algorithm (Asmussen et al., 1996; Thümmler et al., 2006), and choose the number of phases in each branch by five-fold held-out likelihood. Two Erlang branches, five continuous parameters in all, fit every task, model and input within sampling error (Figure 3), and a single exponential phase fits worse. Settings and further tests are in Appendix E.1.

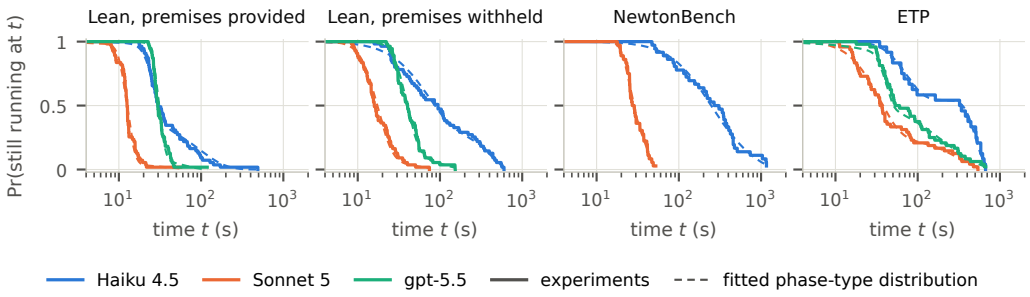

Figure 3: Phase-type distributions fit existing agents' performance. Solid: fraction of attempts still running at time t in the experiments; dashed: the fitted hyper-Erlang distribution with two Erlang branches. Every fit lies within the 95% sampling band, for three kinds of task, three language models and, on Lean, with the premises provided or withheld.

6.2 Second coarse-graining: scaling laws

To test the scaling laws of Section 5 with real agents, Claude Code with Sonnet 5 researches a world of the ETP (Bolan et al., 2025). The world has 40 equational laws; of the 1560 implications between them, 439 are open at the start, 135 true and 304 false. The agents share one record: an agent proves an implication, or refutes it with a finite magma given by its operation table; a checked result enters the record, which then settles every question that follows from it by transitivity or from a known table. Let T_{50} and T_{80} be the first times at which M_{open} has fallen to one half and to one fifth of its initial value. We ran four campaigns at each of $N = 1$, 4 and 8 and one at each of $N = 16$, 64, 128 and 256; those with $N \geq 16$ ran with 16 processes in virtual time, which reproduces the order of events of a simultaneous run (Appendix E.3). We compare them with the coarse-grained dynamics with the phases kept (Section 4.2). Its only inputs are the implication hypergraph of the world and the action law of one agent, a fixed delay followed by phases, fitted on the attempts of the campaigns with $N = 1$ separately for true and false questions. Nothing is fitted to campaigns with more agents.

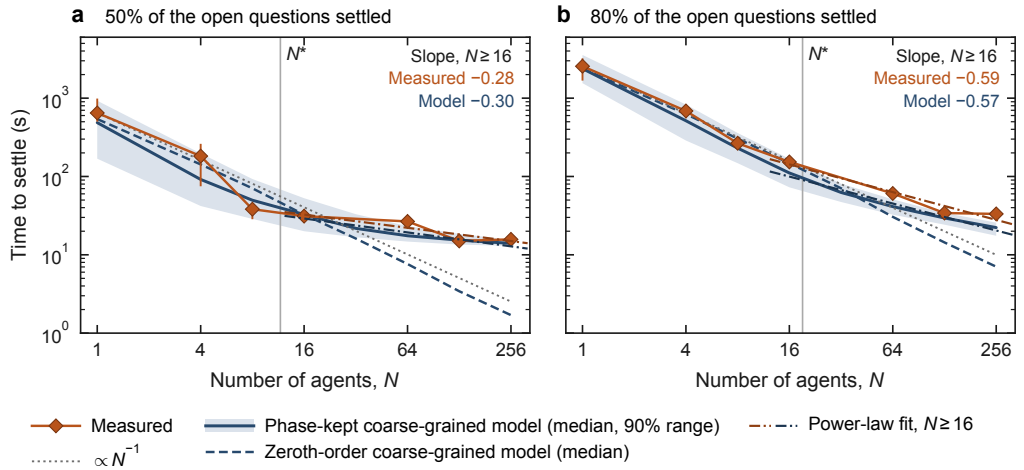

Figure 4: Calibrated on one agent, the phase-kept coarse-grained model reproduces how Claude Code agents scale on a world of the Equational Theories Project (ETP), while the zeroth-order model keeps falling as N^{-1}. Time to settle (a) half and (b) four fifths of the 439 open questions against the number of agents N. Measured points are the mean of four campaigns, with their range, for $N \leq 8$ and one campaign for larger N. The only inputs of both models are the action law fitted on the attempts of one agent and the implication hypergraph of the world; nothing is fitted to campaigns with more agents. The vertical line is the onset of competition N^* predicted from the campaigns of one agent, Eq. (17); the power-law fits are least-squares fits for $N \geq 16$, with their slopes given in each panel.

Up to about 16 agents, T_{50} and T_{80} fall roughly as N^{-1}, the collapse of Eq. (15); beyond, they level off (Figure 4). The coarse-grained dynamics with the phases kept reproduces both parts, with log–log slopes close to the measured ones below and above 16 agents, while the zeroth order, with the same inputs, keeps falling as N^{-1} (dashed lines; Appendix E.3). The difference lies in how actions begin. The zeroth order uses stationary rates, which are positive at once. With the phases kept, every action of the fitted law begins with the delay and a first phase that cannot end in a result, so no result appears in the first seconds, whatever N; and a result that settles many questions at once cancels all their researchers, who begin again. Such cancellations become common beyond the onset of competition N^* of Section 5, where one attempt expects one cancellation; counting the questions that the closure settles with each result, the campaigns of one agent place N^* at the bend (vertical lines; Appendix E.3). With 128 and 256 agents, T_{50} is close to the floor $D_L t_{\min}$ of Section 5, with $D_L = 1$, since every question here can be settled directly, and $t_{\min}$ the delay of a refutation, since half of the open questions cannot be settled without refutations.

7 Discussion

We modelled collective AutoResearch as agents interacting with a growing knowledge hypergraph, described each agent by the phase-type distribution of the duration and outcome of its actions, and connected this local model to coarse-grained collective dynamics. In the experiments, phase-type distributions fit existing agents on three kinds of task, and the coarse-grained dynamics, given the distribution of a single agent, predicts how the completion time of many agents sharing a record changes with their number, in close agreement with the measurements.

For the further work, the model treats questions of one class as equally hard, but the same agent acts with a different distribution on easy and on hard questions (Appendix E.1). Easy questions are then settled early while attempts on hard ones keep failing, and once the easy questions are gone, many agents gather on the hard ones and stay stuck there. A distribution that depends on the difficulty of each question is the next step for the model.

AI USE STATEMENT

Generative AI assistants supported literature review and manuscript revision. Authors remain responsible for verification and the final claims.

REFERENCES

P. W. Anderson. More is different. *Science*, 177(4047):393–396, 1972. doi: 10.1126/science.177.4047.393.

Søren Asmussen, Olle Nerman, and Marita Olsson. Fitting phase-type distributions via the EM algorithm. *Scandinavian Journal of Statistics*, 23(4):419–441, 1996.

Nicolas Behr, Vincent Danos, and Ilias Garnier. Stochastic mechanics of graph rewriting. In *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science*, LICS '16, pp. 46–55. ACM, 2016. doi: 10.1145/2933575.2934537. URL https://www.pure.ed.ac.uk/ws/files/25448314/final_paper_submission159_2.pdf.

Nicolas Behr, Vincent Danos, and Ilias Garnier. Combinatorial conversion and moment bisimulation for stochastic rewriting systems. *Logical Methods in Computer Science*, 16(3), 2020. doi: 10.23638/LMCS-16(3:3)2020. URL https://arxiv.org/abs/1904.07313.

Blaž Bertalanič and Carolina Fortuna. The Ringelmann effect in multi-agent LLM systems: A scaling law for effective team size. *arXiv preprint arXiv:2606.02646*, 2026.

Federico Bianchi, Yongchan Kwon, Aneesh Pappu, and James Zou. Harnessing the collective intelligence of AI agents in the wild for new discoveries. *arXiv preprint arXiv:2606.10402*, 2026.

Mogens Bladt and Bo Friis Nielsen. *Matrix-Exponential Distributions in Applied Probability*, volume 81 of *Probability Theory and Stochastic Modelling*. Springer, New York, 2017.

Daniil A. Boiko, Robert MacKnight, Ben Kline, and Gabe Gomes. Autonomous chemical research with large language models. *Nature*, 624:570–578, 2023. doi: 10.1038/s41586-023-06792-0.

Matthew Bolan, Joachim Breitner, Jose Brox, Nicholas Carlini, Mario Carneiro, Floris van Doorn, et al. The equational theories project: Advancing collaborative mathematical research at scale. *arXiv preprint arXiv:2512.07087*, 2025.

Kevin Buzzard and Jon Eugster. Natural number game (Lean 4 version). https://github.com/leanprover-community/NNG4, 2023. Original Lean 3 version by K. Buzzard and M. Pedramfar (2019); game engine lean4game by A. Bentkamp, J. Eugster and P. Massot.

Claudio Castellano, Santo Fortunato, and Vittorio Loreto. Statistical physics of social dynamics. *Reviews of Modern Physics*, 81(2):591–646, 2009. doi: 10.1103/RevModPhys.81.591.

Anjan Chakravartty. Scientific realism. The Stanford Encyclopedia of Philosophy, 2017. URL https://plato.stanford.edu/entries/scientific-realism/. Substantive revision June 12, 2017; accessed September 16, 2026.

Lingjiao Chen, Jared Quincy Davis, Boris Hanin, Peter Bailis, Ion Stoica, Matei Zaharia, and James Zou. Are more LLM calls all you need? Towards scaling laws of compound inference systems, 2024. URL https://arxiv.org/abs/2403.02419.

Johan S. G. Chu and James A. Evans. Slowed canonical progress in large fields of science. *Proceedings of the National Academy of Sciences*, 2021. doi: 10.1073/pnas.2021636118.

Giordano De Marzo, Claudio Castellano, and David Garcia. AI agents can coordinate via majority-following beyond human scale. *Science Advances*, 12(33):eaea6091, 2026. doi: 10.1126/sciadv.aea6091.

John Dewey. *Logic: The Theory of Inquiry*. Henry Holt and Company, New York, 1938. URL https://books.google.com/books?id=iV0aAAAAIAAJ.

Kefan Dong and Tengyu Ma. STP: Self-play LLM theorem provers with iterative conjecturing and proving. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 14114–14136. PMLR, 2025. URL https://proceedings.mlr.press/v267/dong25h.html.

Batu El, Jinhee Paeng, Fatih Dinc, Shiye Su, Mete Erdogan, Aneesh Pappu, Haotian Ye, Wanjia Zhao, Surya Ganguli, and James Zou. Physics of agents: Statistical mechanics predicts collective behavior of AI agents. *arXiv preprint arXiv:2608.16578*, 2026.

Jingru Fan, Dewen Liu, Yufan Dang, Huatao Li, Yuheng Wang, Wei Liu, Feiyu Duan, Xuanwen Ding, Shu Yao, Lin Wu, Ruijie Shi, Wai-Shing Leung, Yuan Cheng, Zhongyu Wei, Cheng Yang, Chen Qian, Zhiyuan Liu, and Maosong Sun. Towards a science of collective AI: LLM-based Multi-Agent systems need a transition from blind Trial-and-Error to rigorous science, 2026. URL https://arxiv.org/abs/2602.05289.

Ariel Flint, Luca Maria Aiello, Romualdo Pastor-Satorras, and Andrea Baronchelli. Group size effects and collective misalignment in LLM multi-agent systems. *Proceedings of the National Academy of Sciences*, 123(34):e2531697123, 2026. doi: 10.1073/pnas.2531697123.

Giorgio Gallo, Giustino Longo, Stefano Pallottino, and Sang Nguyen. Directed hypergraphs and applications. *Discrete Applied Mathematics*, 42(2–3):177–201, 1993.

Ali E. Ghareeb, Benjamin Chang, Ludovico Mitchener, et al. A multi-agent system for automating scientific discovery. *Nature*, 655:497–505, 2026. doi: 10.1038/s41586-026-10652-y.

Dror Givon, Raz Kupferman, and Andrew Stuart. Extracting macroscopic dynamics: model problems and algorithms. *Nonlinearity*, 17(6):R55–R127, 2004. doi: 10.1088/0951-7715/17/6/R01.

Juraj Gottweis, Wei-Hung Weng, Alexander Daryin, et al. Accelerating scientific discovery with Co-Scientist. *Nature*, 655:487–496, 2026. doi: 10.1038/s41586-026-10644-y.

Jaakko Hintikka. *Inquiry as Inquiry: A Logic of Scientific Discovery*. Kluwer, Dordrecht, 1999.

Wim Hordijk and Mike Steel. Detecting autocatalytic, self-sustaining sets in chemical reaction systems. *Journal of Theoretical Biology*, 227(4):451–461, 2004.

Yubin Kim, Ken Gu, Chanwoo Park, Chunjong Park, Samuel Schmidgall, A. Ali Heydari, Yao Yan, Zhihan Zhang, Yuchen Zhuang, Yun Liu, Mark Malhotra, Paul Pu Liang, Hae Won Park, Yuzhe Yang, Xuhai Xu, Yilun Du, Shwetak Patel, Tim Althoff, Daniel McDuff, and Xin Liu. Towards a science of scaling agent systems. *arXiv preprint arXiv:2512.08296*, 2026. URL https://arxiv.org/abs/2512.08296v3. Version 3, revised April 8, 2026; first posted in 2025.

Philip Kitcher. The division of cognitive labor. *The Journal of Philosophy*, 87(1):5–22, 1990. doi: 10.2307/2026796.

Thomas G. Kurtz. Solutions of ordinary differential equations as limits of pure jump Markov processes. *Journal of Applied Probability*, 7(1):49–58, 1970. doi: 10.2307/3212147.

Larry Laudan. *Progress and Its Problems: Towards a Theory of Scientific Growth*. University of California Press, Berkeley, 1977.

Sean D. Lawley. Extreme first-passage times for random walks on networks. *Physical Review E*, 102:062118, 2020. doi: 10.1103/PhysRevE.102.062118.

Junyou Li, Qin Zhang, Yangbin Yu, Qiang Fu, and Deheng Ye. More agents is all you need. *Transactions on Machine Learning Research*, 2024. URL https://arxiv.org/abs/2402.05120.

Xinze Li, Nanyun Peng, Simone Severini, and Patrick Shafto. The network structure of Mathlib. *arXiv preprint arXiv:2604.24797*, 2026.

Chris Lu, Cong Lu, Robert Tjarko Lange, Yutaro Yamada, Shengran Hu, Jakob Foerster, David Ha, and Jeff Clune. Towards end-to-end automation of AI research. *Nature*, 651:914–919, 2026a. doi: 10.1038/s41586-026-10265-5.

Shuo Lu, Weicheng Meng, Aijing Yu, Kun Shao, Jian Luan, Ran He, and Jian Liang. Topological collapse of higher-order interactions bottlenecks collective intelligence in AI agent societies. *arXiv preprint arXiv:2608.15519*, 2026b.

Pascal Massart. The tight constant in the Dvoretzky–Kiefer–Wolfowitz inequality. *The Annals of Probability*, 18(3):1269–1283, 1990.

Patrick Massot. leanblueprint: a plasTeX plugin to build formalization blueprints. `https://github.com/PatrickMassot/leanblueprint`, 2023.

Hazime Mori. Transport, collective motion, and Brownian motion. *Progress of Theoretical Physics*, 33(3):423–455, 1965. doi: 10.1143/PTP.33.423.

Buddhika Nettasinghe and Kang Zhao. Dynamics of human-AI collective knowledge on the web: A scalable model and insights for sustainable growth. *arXiv preprint arXiv:2601.20099*, 2026.

Arnab Pal and Shlomi Reuveni. First passage under restart. *Physical Review Letters*, 118:030603, 2017. doi: 10.1103/PhysRevLett.118.030603.

Chen Qian, Zihao Xie, YiFei Wang, Wei Liu, Kunlun Zhu, Hanchen Xia, Yufan Dang, Zhuoyun Du, Weize Chen, Cheng Yang, Zhiyuan Liu, and Maosong Sun. Scaling large language model-based multi-agent collaboration. In *International Conference on Learning Representations*, 2025. URL `https://arxiv.org/abs/2406.07155`.

Shlomi Reuveni. Optimal stochastic restart renders fluctuations in first passage times universal. *Physical Review Letters*, 116:170601, 2016. doi: 10.1103/PhysRevLett.116.170601.

Bernardino Romera-Paredes, Mohammadamin Barekatain, Alexander Novikov, et al. Mathematical discoveries from program search with large language models. *Nature*, 625:468–475, 2024. doi: 10.1038/s41586-023-06924-6.

Andrey Rzhetsky, Jacob G. Foster, Ian T. Foster, and James A. Evans. Choosing experiments to accelerate collective discovery. *Proceedings of the National Academy of Sciences*, 112(47):14569–14574, 2015.

Hidenori Tanaka. When is collective intelligence a lottery? multi-agent scaling laws for memetic drift in LLMs. *arXiv preprint arXiv:2603.24676*, 2026.

Michael te Vrugt, Jens Bickmann, and Raphael Wittkowski. How to derive a predictive active field theory: a step-by-step tutorial, 2022. URL `https://arxiv.org/abs/2212.00078`.

Axel Thümmler, Peter Buchholz, and Miklós Telek. A novel approach for phase-type fitting with the EM algorithm. *IEEE Transactions on Dependable and Secure Computing*, 3(3):245–258, 2006. doi: 10.1109/TDSC.2006.27.

Mark Towers, Ariel Kwiatkowski, Jordan Terry, John U. Balis, Gianluca De Cola, Tristan Deleu, Manuel Goulão, Andreas Kallinteris, Markus Krimmel, Arjun KG, Rodrigo Perez-Vicente, Andrea Pierré, Sander Schulhoff, Jun Jet Tai, Hannah Tan, and Omar G. Younis. Gymnasium: A standard interface for reinforcement learning environments, 2024. URL `https://arxiv.org/abs/2407.17032`. NeurIPS 2025 Datasets and Benchmarks Track.

Hanchun Wang. Negative knowledge as failure-aware shared memory for AutoResearch. In *ICML 2026 AI4Research Workshop*, 2026. URL `https://github.com/hch-wang/Negative_Knowledge`. arXiv:2606.21024.

Pengsong Zhang, Heng Zhang, Huazhe Xu, et al. Scaling laws in scientific discovery with AI and robot scientists. *arXiv preprint arXiv:2503.22444*, 2025.

Yifan Zhang, Yunheng Zou, Shaokun Zhang, Jian Hu, Hao Zhang, Binfeng Xu, Jan Kautz, and Yi Dong. Agora: Git as shared memory for collective AutoResearch, 2026. URL `https://arxiv.org/abs/2609.18094`.

Tianshi Zheng, Kelvin Kiu-Wai Tam, Newt Hue-Nam K. Nguyen, Baixuan Xu, Zhaowei Wang, Jiayang Cheng, Hong Ting Tsang, Weiqi Wang, Jiaxin Bai, Tianqing Fang, Yangqiu Song, Ginny Y. Wong, and Simon See. NewtonBench: Benchmarking generalizable scientific law discovery in LLM agents, 2025. URL https://arxiv.org/abs/2510.07172. ICLR 2026.

Kevin J. S. Zollman. The communication structure of epistemic communities. *Philosophy of Science*, 74(5):574–587, 2007.

Robert Zwanzig. Memory effects in irreversible thermodynamics. *Physical Review*, 124:983–992, 1961. doi: 10.1103/PhysRev.124.983.

Organization of the appendices

Each appendix supports one section of the main text and is cited from it; material that supports no statement of the main text has been left out. Appendix A states the philosophical setting of the model and the modelling assumptions. Appendix B gives the microscopic model of Section 3 in full: the record and the agents' views, the read and write rules, publication and boundary rules, the phase-type representation of the action law, the reference agent used in the simulations. Appendix C derives the results of Section 4: the exact balances, the elimination of the phases, the baseline closure and its corrections. Appendix D proves the predictions of Section 5: time scaling and retrieval-dependent growth, with conditional supply and same-question race calculations. Appendix E gives the experiments of Section 6: the settings, the fitted action laws, the registered predictions and their outcomes, and the numerical checks of the algebra.

A Philosophical setting and modelling assumptions

A.1 Setting

We assume that the validity of a result is constrained independently of researchers' beliefs (Chakravartty, 2017), treat research as inquiry ending in settled results that later inquiry uses (Dewey, 1938), and take the community as the unit producing knowledge (Kitcher, 1990; Zollman, 2007). Research as alternating questions and deductive moves follows the interrogative model of inquiry (Hintikka, 1999), and progress measured by solved problems follows Laudan (1977); both are one account among several, cited as precedents and not as consensus. The hidden hypergraph of Section 3.1 represents this constraint; it is a modelling device, not a claim about the ontology of science, and it need not be enumerated or reconstructed. Formal validity is relative to fixed axioms and libraries; empirical domains require their own evidential standards. A formal library records part of this structure: the dependency network of the Lean library Mathlib gives, for each declaration, the set of declarations its proof uses (Li et al., 2026).

Checking belongs to research: the check accepts, rejects or leaves unresolved a submitted derivation, and adds neither a separate verification activity nor a labelled node for the certificate. Sound checks and no retraction are assumptions, and true and false labels are final. The count M of Eq. (6) counts solved problems without weighting by importance, the measure of progress of Laudan (1977); in the blueprint convention (Massot, 2023) true corresponds to *proved* and undetermined to *stated*.

A.2 Modelling assumptions

We fix a domain and statements of comparable granularity, so that a direct derivation fits within one research action, including retrieval, reasoning and checking. We make five assumptions. **A1 (validation).** An independent check accepts or rejects every submitted derivation, and accepted results are never retracted. **A2 (joint prerequisites).** A fixed hidden structure specifies which sets of premises jointly suffice for each statement. **A3 (conditional description).** How long an agent takes and whether it succeeds depend on the task and on the record only through a few declared inputs, such as whether a premise set of the target is established, how many of its premises the agent holds, the size of the record and the failed attempts recorded on the target. **A4 (finite capacity).** An agent holds at most C statements for immediate use and performs one action at a time. **A5 (visibility).** An agent sees the public record and its own unpublished findings.

Validation determines which outputs count as knowledge; prerequisite relations connect research tasks; and finite time and capacity make the distribution of knowledge across agents consequential. Sharing can thus change both what agents can establish and how long they take.

B The microscopic model in full

This appendix gives the update rules of the model of Section 3 that the main text leaves out, the definition of the phase-type law, and the agent used in the simulator.

B.1 Update rules

The hidden directed hypergraph $\mathcal{H}^\star = (V^\star, E^\star)$ is finite and fixed; a simulator receives it explicitly, a real environment supplies an equivalent rule. Premises need not be known when a question is posed, and several incoming hyperedges are alternative sufficient premise sets. A failure case is a record attached to a question, not a vertex, listing the statements its author loaded for it; it does not change the label. Agent i sees $\mathcal{O}_i(t) = \mathcal{H}_t \cup h_i(t)$ and holds $K_i(t) \subseteq V(\mathcal{O}_i(t))$ with $|K_i(t)| \le C$. Within each view the true nodes are the least closure of the true seeds under the recorded edges, a node becomes false only through a checked refutation, both labels are final, and canonical identities of nodes and edges keep repeated publication from increasing M.

Loading inserts an accessible statement into K_i; a full base first evicts an eligible statement, or the load has no effect if none is eligible. At termination any protection of the current action's loads is released, and a newly established result enters the solver's base under the same capacity rule. Publication inserts nothing into other bases. With memory, K_i persists between actions; without memory it is cleared and refilled at the start of the next action. In the public setting agents are coupled only through publication.

The campaign starts from a nonempty public seed set, true initially, optional known open targets, empty private hypergraphs, bases drawn uniformly from $\min(C, M_0)$ seeds, and $N, C \ge 1$, $\sigma, \sigma_{\mathrm{f}} \in [0, 1]$. Each genuinely new output receives one Bernoulli publication decision when produced, with probability σ for a question, a verified derivation or a refutation and σ_{f} for a failure case; a negative decision leaves it private. A positive decision releases the output with a private derivation chain supporting it, one verified derivation for each premise established only in the agent's view down to public results, chosen by a fixed ordering; the released statements and edges enter the public sets without duplicates, new premises as undetermined, and public true labels follow by closure. Publication takes no time. All affected closures are then recomputed, and the research of every agent whose target has become true or false in its view is cancelled without a failure case; the completing and cancelled agents choose new actions at once, the others keep theirs. Simultaneous completions are a null event, a closure cascade belongs to one update, and with finite phase-type laws on a finite graph finitely many events occur in finite time.

B.2 Phase-type laws

A phase-type law with finitely many outcomes is the joint law of the absorption time and the absorbing state of a continuous-time Markov chain on a finite set of transient phases $\{1, \dots, m\}$ and one absorbing state per outcome o. It is specified by the initial row β on the phases, the $m \times m$ sub-generator Q of the transitions among phases, with $Q_{\phi\psi} \ge 0$ for $\phi \ne \psi$ and negative diagonal, and the exit-rate columns $\mathbf{r}_o \ge 0$ into each outcome, with $\sum_o \mathbf{r}_o = -Q\mathbf{1}$ so that every phase is eventually absorbed. The duration T is the absorption time and the outcome J the absorbing state; for a research action the outcomes are a result, a rejected submission and giving up, and $s = \mathbf{r}_{\mathrm{result}}$ is the success exit vector of Section 4. Since βe^{Qt} is the row of phase probabilities at time t,

$$\begin{aligned} \Pr(T > t) = \beta e^{Qt}\mathbf{1}, \qquad f_o(t) = \beta e^{Qt}\mathbf{r}_o, \\ \mathbb{E}[e^{-sT}\mathbf{1}_{\{J=o\}}] = \beta(s\mathbf{I} - Q)^{-1}\mathbf{r}_o, \qquad \mathbb{E}[T] = \beta(-Q)^{-1}\mathbf{1}, \end{aligned} \tag{18}$$

where f_o is the density of ending at time t with outcome o and $\Pr(J = o) = \beta(-Q)^{-1}\mathbf{r}_o$ (Bladt & Nielsen, 2017). A hyper-Erlang law is the subclass in which the chain is a mixture of branches, each a series of identical exponential phases ending in one outcome. At a fixed record the description is exact when the agent's procedure has finitely many states with exponential holding times; for other agents it is a fitted approximation, and phase-type laws are dense among joint laws of a positive duration and finitely many outcomes.

B.3 The reference agent

The *reference search* is a fixed policy whose action law is known exactly; Appendix D.4 derives its search-chain approximation and Appendix E.4 checks the algebra against its simulation. The model does not assume this policy. At each boundary the agent poses with probability α and researches otherwise; if no target is open it poses. A posing action lasts an exponential time of mean τ_{pose} and then, with probability π, outputs a uniform draw from the unrevealed statements adjacent to a held

statement in $\mathcal{H}^\star$,

$$A_i = \{v \in V^\star \setminus V(\mathcal{O}_i) : \exists (P, q) \in E^\star,\ v \in P \cup \{q\},\ (P \cup \{q\}) \cap K_i \neq \emptyset\},$$

into h_i, with no output if A_i is empty. Without memory, each action starts with a uniform set of up to C established results from its view, taking all of them if fewer are available. A research action takes a uniform undetermined target q in the agent's view and searches: loads occur at rate ν_{s}, each with probability ω a missing established premise of q and otherwise a uniform established statement not held, a full base evicting a uniform statement not loaded in this action. If no missing premise is available, the directed choice falls back to a uniform load; an empty load pool gives no update. The agent attempts a derivation at rate μ, which passes with probability p if the gate (3) is open and otherwise ends the action with a failure case listing its loads; it gives up at rate γ; it does not attempt while its loads lie within a visible failure case on q; and a result adds q to K_i, evicting a uniform held statement if full after releasing protection. In environments that admit refutation a third exit ends the action with a checked refutation of q.

C Derivations for Section 4

This appendix derives the coarse-grained dynamics of Section 4.2 at its two orders. The zeroth order is the baseline: it follows from the exact growth balance by two closures, eliminating the phases and closing the entries. The first order is the linear response of the growth rate to the mechanisms the baseline leaves out. We work in the public setting, $\sigma = \sigma_{\mathrm{f}} = 1$, without refutation or pooling, and assume non-explosion and finite moments, so that the balances hold in differential form.

C.1 Zeroth order: the baseline

Assumptions. The baseline rests on eight assumptions. (i) Agents choose open questions uniformly. (ii) They keep no memory between actions. (iii) They use no failure lists. (iv) At entry, an agent holds a uniform random set of $\min(C, M)$ public true statements, independent of its target. (v) Other researchers cancel an action at a rate κ_q that does not depend on its phase. (vi) The inputs of an action stay fixed while it runs. (vii) After an initial transient, the phase occupancies relax faster than the inputs change. (viii) All agents are alike. Table 1 states what relaxing each one does.

Exact starting point. A researcher of q in phase ϕ ends with a result at the rate $s_{q\phi}$, so results on q appear at the rate $\lambda_q = \sum_\phi X_{q\phi} s_{q\phi}$ and $\mathrm{d}\mathbb{E}[M]/\mathrm{d}t = \mathbb{E}[\lambda]$ with $\lambda = \sum_q \lambda_q$, Eq. (7). This is exact but not closed: it needs the phase occupancies $X_{q\phi}$, whose own balance,

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathbb{E}[X_a] = \mathbb{E}\big[X_a Q_a + U_a - C_a\big], \qquad (C_q)_\phi = X_{q\phi}\big(\lambda_q - s_{q\phi}\big), \tag{19}$$

contains the internal transitions Q_a of each class a of actions, the entries U_a into new actions, and the cancellations C_a: a success from any phase of q cancels every other researcher of q, so a researcher in phase ϕ is cancelled at the rate $\lambda_q - s_{q\phi}$. The product $X_{q\phi}(\lambda_q - s_{q\phi})$ couples the occupancies of different agents, and the entries depend on the record; neither is a function of M, M_{open} and M_{ready} alone.

Step 1: eliminating the phases. Under (v)–(vii), first take a common law and cancellation rate within each of three classes: posing, research on a ready question (R) and research on a question that is not ready (U), each with the sub-generator Q_a, initial row β_a and exit rates $\mathbf{r}_a = -Q_a\mathbf{1}$ of its PH law, and for R the success exits s_{R}. Replacing the pair term by the independent clock, $\mathbb{E}[X_{q\phi}(\lambda_q - s_{q\phi})] \to x_{q\phi}\kappa_q$, the mean occupancies $x_a = \mathbb{E}[X_a]$ obey the linear system

$$\dot{x}_a = x_a(Q_a - \kappa_a \mathbf{I}) + \lambda_a^{\mathrm{in}}\beta_a, \qquad \lambda_a^{\mathrm{in}} = \pi_a \sum_b x_b(\mathbf{r}_b + \kappa_b \mathbf{1}), \qquad \lambda = x_{\mathrm{R}} s_{\mathrm{R}}, \tag{20}$$

with $\kappa_{\mathrm{pose}} = \kappa_{\mathrm{U}} = 0$ and entry probabilities $\pi_{\mathrm{pose}} = \alpha$, $\pi_{\mathrm{R}} = (1-\alpha)\rho$, $\pi_{\mathrm{U}} = (1-\alpha)(1-\rho)$, where $\rho = M_{\mathrm{ready}}/M_{\mathrm{open}}$. Variation of constants solves it, and multiplying by s_{R} gives the success flux as a memory of past entries,

$$\lambda(t) = x_{\mathrm{R}}(0) e^{(Q_{\mathrm{R}} - \kappa \mathbf{I})t} s_{\mathrm{R}} + \int_0^t \lambda_{\mathrm{R}}^{\mathrm{in}}(u)\, k(t-u)\, \mathrm{d}u, \qquad k(u) = \beta_{\mathrm{R}} e^{(Q_{\mathrm{R}} - \kappa \mathbf{I})u} s_{\mathrm{R}}. \tag{21}$$

The integral of the kernel is the success per entry $L_{\mathrm{s}}(\kappa) = \beta_{\mathrm{R}}(\kappa\mathbf{I} - Q_{\mathrm{R}})^{-1} s_{\mathrm{R}}$, and the mean occupation is $\tau_{\mathrm{R}}(\kappa) = \beta_{\mathrm{R}}(\kappa\mathbf{I} - Q_{\mathrm{R}})^{-1}\mathbf{1}$, the two functionals of Eq. (8). When the entries change slowly compared with the kernel, assumption (vii), the occupancies are stationary, $x_a = \lambda_a^{\mathrm{in}} \beta_a(\kappa_a \mathbf{I} - Q_a)^{-1}$, so that $x_a \mathbf{1} = \lambda_a^{\mathrm{in}} \tau_a$ and $\lambda = \lambda_{\mathrm{R}}^{\mathrm{in}} L_{\mathrm{s}}(\kappa)$: the phases are gone, and the action law enters only through L_{s} and the mean times. For weak competition, $L_{\mathrm{s}}(\kappa) = q_{\mathrm{s}} - \kappa\, \mathbb{E}[T \mathbf{1}_{\{J=\mathrm{result}\}}] + O(\kappa^2)$ and $\tau_{\mathrm{R}}(\kappa) = \mathbb{E}[T] - \kappa\, \mathbb{E}[T^2]/2 + O(\kappa^2)$, with q_{s} the success probability of one action. For heterogeneous questions, retain a separate class for each law and cancellation rate and sum their success kernels.

Step 2: closing the entries. Under (i)–(iv) and (viii) the entries form a renewal process of identical agents: an agent that ends an action starts the next at once, posing with probability α and otherwise choosing an open question uniformly. Summing the stationary occupancies over classes, $N = \sum_a x_a \mathbf{1}$, gives the total entry rate $N/\bar{\tau}$ with the mean cycle $\bar{\tau} = \alpha\tau_{\mathrm{pose}} + (1-\alpha) M_{\mathrm{open}}^{-1} \sum_q \tau_q$, and uniform choice gives each open question the entry rate $\lambda_q^{\mathrm{in}} = N(1-\alpha)/(\bar{\tau} M_{\mathrm{open}})$. Each entry succeeds with $L_{\mathrm{s},q}(\kappa_q)$ averaged over the holdings at entry: under (iv) the number j_0 of premises of a ready k-premise route held at entry is hypergeometric. With $c = \min(C, M)$, a whole route is covered with probability

$$\chi_q = \begin{cases} \dfrac{\binom{M-k}{c-k}}{\binom{M}{c}} = \dfrac{(c)_k}{(M)_k}, & k \le c, \\ 0, & k > c, \end{cases} \tag{22}$$

several routes combining as a union. The cancellation rate of a researcher of q is the success rate of the $N-1$ others on q, and the cycle time depends on it, so both are solved together,

$$\bar{\tau} = \alpha\tau_{\mathrm{pose}} + \frac{1-\alpha}{M_{\mathrm{open}}} \sum_{q\ \mathrm{open}} \tau_q(\kappa_q; Z), \qquad \kappa_q \simeq \frac{N-1}{N} W(Z \to Z^{+q}), \tag{23}$$

with $\kappa_q = 0$ on questions that are not ready. When no question is open, use $\alpha = 1$ and $\bar{\tau} = \tau_{\mathrm{pose}}$ at this boundary. The posing kernel p_u^{pose} is the probability that one posing action outputs u, for the reference agent the average of $\pi \mathbf{1}\{u \in A_i\}/|A_i|$ over the entry holdings, with zero contribution from an empty A_i. Multiplying entry rate by success per entry gives the rates of Eq. (9), and for any function f of the record

$$(\mathcal{L}_Z f)(Z) = \sum_{Z'} W(Z \to Z') \big[f(Z') - f(Z) \big], \tag{24}$$

where Z' runs over Z^{+q} and Z^{+u}; this is the generator of the master equation (10), and readiness is recomputed from the hidden hypergraph after every jump. When all ready questions share one law and one cancellation rate κ, the per-agent throughput is

$$\frac{\lambda}{N} = \frac{(1-\alpha)\rho L_{\mathrm{s}}(\kappa)}{\bar{\tau}}, \qquad \bar{\tau} = \alpha\tau_{\mathrm{pose}} + (1-\alpha)\big[\rho\, \tau_{\mathrm{R}}(\kappa) + (1-\rho)\tau_{\mathrm{U}}\big], \tag{25}$$

entries into ready questions times success per entry, and with $\kappa \simeq (N-1)\lambda/(N M_{\mathrm{ready}})$ the counts obey $\mathrm{d}M/\mathrm{d}t = \lambda$ and $\mathrm{d}M_{\mathrm{open}}/\mathrm{d}t = N\alpha b/\bar{\tau} - \lambda$, with b the probability that a posing action yields a question. These deterministic equations need in addition a closure of ρ and b on the counts and a concentration argument (Kurtz, 1970) that the growing graph does not supply.

Keeping the phases. Two situations violate (vii): at the start all agents are at entry, and in ETP one result can settle many questions through the record, so that all their researchers restart together. Take, as in ETP, no posing and every open question ready. Group the open questions into classes c with action laws (β_c, Q_c, s_c) and failure exits $g_c = -Q_c \mathbf{1} - s_c$, and let $M_{\mathrm{open},c}$ count those of class c. All N agents start at $t = 0$; the completion flow $\sum_b x_b(s_b + g_b)$ and the cancelled agents start again at once, each on a uniformly chosen open question. A start waits the delay d_c of its class, and $a_c(t)$ is the flow of those not cancelled meanwhile. A result falls on class c with probability $x_c s_c/\lambda$ and on a uniformly chosen open question of that class; the record adds it and recomputes the labels. If ΔM_c questions of class c become true or false, the share $\Delta M_c / M_{\mathrm{open},c}$ of x_c, and of the class-c agents still in their delay, is cancelled. For ETP the classes are the true and the false targets, and a refutation adds a refuting table of smallest order drawn from a pool.

C.2 First order: the effect of each dropped mechanism

A mechanism that the baseline drops, if weak and without a slow variable, changes only the three inputs of a question: the probability π_q that an entering agent chooses it, its success per entry $\ell_q = L_{\mathrm{s},q}(\kappa_q)$, and its mean occupation τ_q. The renewal argument of Step 2 with these inputs gives Eq. (11), and differentiating $\log\lambda$ at fixed N and α, with the posing law fixed, gives the linear response

$$\delta\lambda = \frac{N(1-\alpha)}{\bar\tau_0}\sum_q \big(\ell_q^{(0)}\,\delta\pi_q + \pi_q^{(0)}\,\delta\ell_q\big) - \frac{\lambda_0}{\bar\tau_0}\,\delta\bar\tau, \qquad \delta\bar\tau = (1-\alpha)\sum_q \big(\tau_q^{(0)}\delta\pi_q + \pi_q^{(0)}\delta\tau_q\big), \tag{26}$$

where the superscript 0 marks the baseline. The first term is the change in results per cycle, the second the cost of time: a mechanism that raises success but lengthens the cycle can lower λ. A mechanism whose effect is not small, or that adds a variable outliving one action, extends the state Z instead, and one that breaks the separation of time scales requires the phases of Eq. (20). Table 1 lists the cases; the coefficients follow.

Competition and the restart index. Cancellation enters ℓ_q and τ_q through κ_q: to first order $\delta\ell_q = -\kappa_q\mathbb{E}[T\mathbf{1}_{\{J=\mathrm{result}\}}]$ and $\delta\tau_q = -\kappa_q\mathbb{E}[T^2]/2$, so for questions of equal difficulty

$$\partial_\kappa \log\lambda\big|_{\kappa=0} = \frac{(1-\alpha)\rho}{\bar\tau_0}\,\frac{\mathbb{E}[T^2]}{2} - \tau_{\mathrm{s}}, \qquad \tau_{\mathrm{s}} = \mathbb{E}[T \mid J = \mathrm{result}], \tag{27}$$

with $\bar\tau_0$ the cycle without cancellation. Cancellation helps only when the mean remaining duration of an action seen at a random moment exceeds the mean duration of a successful action, the restart criterion (Reuveni, 2016; Pal & Reuveni, 2017) with the outcome marked; for $\rho = 1$ and $\alpha = 0$ the sign is that of $\Phi - 1$ with $\Phi = \mathbb{E}[T^2]/(2\,\mathbb{E}[T]\,\tau_{\mathrm{s}})$. The mark matters: with $\beta = (1/2, 1/2)$, $Q = \mathrm{diag}(-1,-4)$ and success from the slow branch, $s = (1,0)^\top$, the slope is $-3/20$, and with success from the fast branch, $s = (0,4)^\top$, it is $3/5$, although both laws have the same duration distribution and success probability. Since $\kappa_q \propto N - 1$, this is the first departure from the time scaling of Appendix D.3.

Target policy, failure lists, memory. A policy σ_q gives $\delta\pi_q = \sigma_q - 1/M_{\mathrm{open}}$ and moves $\bar\tau$ when durations differ across questions. Failure lists move ℓ_q from the lower towards the upper bound of Eq. (33) and can lengthen τ_q, both computed from the search chain with the suppressed attempt clock; their write rates carry the factor N, so they change the prefactor of the growth curve but not the time scaling. Memory across actions makes the holdings at entry lag the record. In a comparison model with a fixed record of $M > C$ results, full bases, a load clock of rate ν per agent and uniform eviction, the mean holding fraction of a result obeys

$$\dot a_v = \nu\frac{1-a_v}{M-C} - \nu\frac{a_v}{C}, \qquad a_v(t) = \frac{C}{M} + \Big(a_v(0) - \frac{C}{M}\Big)e^{-t/T_K}, \qquad T_K = \frac{C(M-C)}{\nu M}, \tag{28}$$

so the distribution of j_0 at entry is shifted by the holding deficit of recent results, and ℓ_q, τ_q change by the corresponding averages of the conditional law.

C.3 Comparison model

If every candidate were revealed, no agent competed for a question, and each route $P \to q$ took an independent positive delay D_{Pq} from the moment its premises are established, the establishment times would satisfy $t_q = \min_{P\to q}\{\max_{v\in P} t_v + D_{Pq}\}$ with $t_{\mathrm{seed}} = 0$, an AND–OR growth recursion, first-passage percolation for single premises. It is a comparison model, not a consequence of the jump model, whose delays are neither independent nor stationary; no universal fluctuation class follows either way.

D Derivations for Section 5

This appendix derives the reach bound, the supply limit, the time scaling, the growth laws of the search chain and the crowding law.

Table 1: Assumptions of the baseline and what relaxing each one does. A small violation moves the inputs π_q, ℓ_q, τ_q of Eq. (11) and acts through Eq. (26); a violation that adds a slow variable extends the state Z. The first row is the public setting of Section 4; the others are assumptions (i)–(viii) of Appendix C.1.

Assumption	Input moved	State extended	Effect when relaxed
Public setting: all outputs public, true premises only, no refutation or pooling	publication changes available premises, cancellation and the laws in each agent's view	yes when private knowledge affects future actions	rates depend on who knows which results or failures; a quantitative closure is not derived
(i) Uniform choice of open questions	π_q, and $\bar{\tau}$ if durations differ	no	first order, Eq. (26); correlated choice also raises cancellation
(ii) No memory across actions	ℓ_q, τ_q, via the holdings at entry	yes if holdings relax slowly: a_v of Eq. (28)	entry holdings change the averages of the conditional action law
(iii) No failure lists	ℓ_q, up to the bound of Eq. (33); τ_q can grow	yes: the failures recorded on each question	first order while q has few lists; a longer cycle can lower λ
(iv) Entry holdings uniform and independent of the target	ℓ_q, τ_q, via the holdings at entry	no	biased by memory; premises shared within a team would add a slow variable (not modelled)
(v) Independent, phase-blind cancellation clock	ℓ_q, τ_q, via cancellation	no, unless many attempts start together	first order: the restart index, Eq. (27); synchronized race: wall time $\propto m^{-1/n}$, work $\propto m^{(n-1)/n}$, Eq. (36)
(vi) Inputs frozen during an action	the law itself, e.g. a question becoming ready mid-search	only for fast changes: the phase occupancies of Eq. (20), with transfers	the action law must respond to changes in its inputs
(vii) Fast relaxation of phase occupancies	none; rates depend on past entries	only for sudden changes: the phases, or the kernel of Eq. (21)	a synchronized start or a sudden change requires the phase dynamics or the kernel
(viii) Identical agents	π_q, ℓ_q, τ_q per agent type	no	not derived

D.1 Capacity closure

The reach bound used in Section 4.2 is the following.

Proposition 1 (Capacity-constrained reachability)**.** *In the public setting without pooling, let $E^\star_C$ be the routes of at most C premises and T_t the public true set. Then $T_t \subseteq \mathrm{cl}_{E^\star_C}(S_0)$ and $M(t) \le |\,\mathrm{cl}_{E^\star_C}(S_0)|$ for every N. Pooling the bases of at most m agents would replace C by mC.*

For $A \subseteq V^\star$, define the inflationary map

$$\Phi_C(A) = A \cup \{q : \exists (P, q) \in E^\star_C,\ P \subseteq A\}.$$

On a finite graph, iterating from S_0 terminates at the least closed set containing S_0. In the public setting every new establishment uses a recorded edge with $P \subseteq K_i \subseteq T_{t-}$ and $|P| \le C$. Induction over events proves Proposition 1.

D.2 Supply-limited and research-limited production

Every result establishes a question that was open initially or posed later, so the long-run rate of results cannot exceed the arrival rate of posed questions within the reach of Appendix D.1. With the posing success probability b, the reachable fraction f_{solv} and a common uncancelled success probability $q_{\mathrm{s}} = L_{\mathrm{s}}(0)$ frozen, the supply and the research capacity are $N\alpha b f_{\mathrm{solv}}/\bar\tau$ and $N(1-\alpha)q_{\mathrm{s}}/\bar\tau$. If research keeps up with the feasible arrivals when supply is limiting, and effort concentrates on ready questions when capacity is limiting, then

$$\frac{\lambda}{N} \simeq \frac{1}{\bar\tau} \min\{\alpha b f_{\mathrm{solv}},\ (1-\alpha)q_{\mathrm{s}}\}. \tag{29}$$

The two rates match at $\alpha^* = q_{\mathrm{s}}/(b f_{\mathrm{solv}} + q_{\mathrm{s}})$. This is a conditional crossover: uniform selection does not by itself ensure the two attainment assumptions, and state-dependent readiness and action times move the actual optimum. In a finite pool that needs no new questions, $\alpha = 0$ is best.

D.3 Time scaling, hitting times and equal budget

Let $\mathcal{L}_N$ be the generator of Eq. (24) with N agents. Under the baseline with $\kappa_q = 0$, the entry holdings, the target choice $1/M_{\mathrm{open}}$, the posing kernel and the cycle time $\bar\tau$ do not depend on N, so every rate is N times its value for one agent and $\mathcal{L}_N = N\mathcal{L}_1$ on the finite state space of records. The transition semigroups then satisfy $e^{t\mathcal{L}_N} = e^{Nt\mathcal{L}_1}$, so from the same initial record the law of $(Z_N(t))_{t\ge0}$ is the law of $(Z_1(Nt))_{t\ge0}$; this is Eq. (15) for M, a function of Z. The hitting time

$$T_L^{(N)} = \inf\{t : M_N(t) - M(0) \ge L\} \tag{30}$$

is a functional of the path, hence $T_L^{(N)} \overset{d}{=} T_L^{(1)}/N$, including the event $T_L = \infty$, whose probability does not depend on N. Define the mean increment $m_N(t) = \mathbb{E}[M_N(t) - M(0)]$. At equal agent time $B = Nt$, $m_N(B/N) = \mathbb{E}[M_1(B) - M(0)] = m_1(B)$, so the equal-budget difference $\Gamma_N(B) := m_N(B/N) - m_1(B)$ is zero. The statement is about laws, not paths: the fluctuations of $M_N(t)$ are those of $M_1(Nt)$. Two mechanisms break $\mathcal{L}_N = N\mathcal{L}_1$: a cancellation rate $\kappa_q > 0$, which depends on $N-1$; and entry holdings that lag the record, since $\lambda \propto N$ moves the record further within the relaxation time T_K of Appendix C.2. Failure lists retained in Z with their write rates keep the identity, because their rates carry the same factor N and their effect on the action law does not depend on N.

Computing the competition scale from a PH law. For a ready question with law (β, Q, s), write $q_{\mathrm{s}} = \beta(-Q)^{-1}s > 0$ and $\tau = \beta(-Q)^{-1}\mathbf{1}$ for its success probability and mean uncancelled duration. At fixed inputs, uniform target choice and weak competition give the success rate of the other agents on this question as

$$\kappa \simeq \frac{(N-1)(1-\alpha)q_{\mathrm{s}}}{\bar\tau_0 M_{\mathrm{open}}}, \qquad \bar\tau_0 = \alpha\tau_{\mathrm{pose}} + (1-\alpha)\big[\rho\tau + (1-\rho)\tau_{\mathrm{U}}\big],$$

where ready questions share the same law and $\rho = M_{\mathrm{ready}}/M_{\mathrm{open}}$. For $(1-\alpha)\rho > 0$, setting the exposure $\kappa\tau$ to one defines the estimate

$$N^* - 1 \simeq \frac{M_{\mathrm{ready}}\bar\tau_0}{(1-\alpha)\rho q_{\mathrm{s}}\tau}. \tag{31}$$

For $\alpha = 0$ and $\rho = 1$, $\bar{\tau}_0 = \tau$, so this reduces to $N^* - 1 \simeq M_{\mathrm{ready}}/q_{\mathrm{s}}$. Neglecting the one-agent correction gives the expression in the main text. Multiplying all PH rates by the same factor changes τ but not this research-only estimate of N^*. The environment must still supply M_{ready}; the PH law alone does not determine the competition scale. This criterion measures exposure to cancellation, not a guaranteed bend in the growth curve. Its effect depends on the duration–outcome law and on changes in the record.

D.4 Sequential search, dilution and directed retrieval

For the example agent on a ready target with one premise set of $k \le C$ statements, j_0 of them held at entry, the research action is a chain whose stage counts the premises held, provided the record is much larger than the loads of one action and loads evict no held premise. The next premise is found at rate $a_j = \nu_{\mathrm{s}}[\omega + (1-\omega)(k-j)/U]$, where ω is the precision of retrieval and U the number of visible results not held; attempts occur at rate μ and the agent gives up at rate γ.

Proposition 2 (Search as a sequential chain)**.** *For this chain, without failure lists and with independent cancellation at rate κ,*

$$L_{\mathrm{s}}(\kappa) = \frac{p\mu}{\mu+\gamma+\kappa} \prod_{j=j_0}^{k-1} \frac{a_j}{a_j+\mu+\gamma+\kappa}. \tag{32}$$

For fixed positive rates and $j_0 = 0$, blind search, $\omega = 0$, gives $L_{\mathrm{s}}(0) = \Theta(M^{-k})$, whereas $\omega > 0$ gives a lower bound that does not depend on the record size, Eq. (34)*.*

To leading order in the number of loads over M, the chain is Markov without failure lists. Its transient matrix has $Q_{jj} = -(a_j+\mu+\gamma)$ and $Q_{j,j+1} = a_j$ for $j < k$, and $Q_{kk} = -(\mu+\gamma)$. The successful exit vector has $s_k = p\mu$ and zero elsewhere; the failed exits have rate $\mu+\gamma$ for $j < k$, a rejected attempt or giving up, and $(1-p)\mu+\gamma$ at k. Starting in phase j, let $b_j(z) = \mathbb{E}_j[e^{-zT}\mathbf{1}_{\mathrm{result}}]$. Competing exponential clocks give

$$b_k(z) = \frac{p\mu}{\mu+\gamma+z}, \qquad b_j(z) = \frac{a_j}{a_j+\mu+\gamma+z}\, b_{j+1}(z),$$

and iteration proves Eq. (32); $z = \kappa$ implements independent cancellation. The intrinsic means obey $\tau_k = 1/(\mu+\gamma)$ and $\tau_j = [1 + a_j\tau_{j+1}]/(a_j+\mu+\gamma)$, and supply the durations in Eq. (25).

Negative knowledge. Within the chain approximation, for a ready target without memory and with $p = 1$, the success per entry without cancellation lies between

$$L_{\mathrm{s}}^0 = \frac{p\mu}{\mu+\gamma} \prod_{j=0}^{k-1} \frac{a_j}{a_j+\mu+\gamma} \quad \text{and} \quad L_{\mathrm{s}}^\infty = \frac{p\mu}{\mu+\gamma} \prod_{j=0}^{k-1} \frac{a_j}{a_j+\gamma}, \tag{33}$$

the values without lists and with every premature attempt suppressed. While a list suppresses attempts, the killing rate of a search stage drops from $\mu+\gamma$ to γ; without memory and with $p = 1$ no list contains a full premise set, so the final attempt is never suppressed, and each search stage is passed with a probability between $a_j/(a_j+\mu+\gamma)$ and $a_j/(a_j+\gamma)$, which proves the bounds. With memory or $p < 1$ a list can delay the final attempt and the lower bound can fail.

Dilution and directed retrieval. For $\omega = 0$ and $j_0 = 0$, $a_j = \nu_{\mathrm{s}}(k-j)/U$ with U of the order of M, so each factor of Eq. (32) is of order $1/M$, $L_{\mathrm{s}}(0) = \Theta(M^{-k})$, while the mean attempt time tends to $1/(\mu+\gamma)$. For fixed $\omega > 0$, $a_j \ge \omega\nu_{\mathrm{s}}$ and each factor increases with a_j, so

$$L_{\mathrm{s}}(0) \ge \frac{p\mu}{\mu+\gamma} \Big(\frac{\omega\nu_{\mathrm{s}}}{\omega\nu_{\mathrm{s}}+\mu+\gamma} \Big)^k, \tag{34}$$

and success and mean duration tend to positive limits as $M \to \infty$.

Growth laws. Suppose ready questions remain a constant positive fraction, cancellations and memory transients are negligible, and a deterministic count closure is justified. If its per-agent

drift is ax^{-d} with fixed $a > 0$ and $d \geq 0$, integration gives $x^{d+1} = M(0)^{d+1} + (d+1)aNt$, giving Eq. (16) once growth dominates the initial count. For the blind chain with $j_0 = 0$, $d = k$ and

$$x(t)^{k+1} = M(0)^{k+1} + (k+1)aNt. \tag{35}$$

Directed search gives instead $d = 0$ and linear growth before structural depletion. Both cases need sustained supply and a separation between attempt relaxation and record growth. With j_0 premises held at entry, $L_{\mathrm{s}}(0) = \Theta(M^{-(k-j_0)})$: well-mixed holdings leave $j_0 = 0$ dominant for large records and change only the prefactor, whereas holdings correlated with the target change the exponent to $1/(k - j_0 + 1)$. Failure lists move the rates between the bounds of Eq. (33).

D.5 Races and crowding on one question

An exact same-question race. For m researchers who start together on one question, each needing n independent exponential stages of rates $r_1, \ldots, r_n$, with no failures, restarts or exchange of knowledge, the first completion $T_{(1)}$ and the total agent time B_m satisfy, as $m \to \infty$,

$$\mathbb{E}[T_{(1)}] \sim \Gamma(1 + 1/n) \left(\frac{n!}{m \prod_{\ell=1}^{n} r_\ell} \right)^{1/n}, \qquad \mathbb{E}[B_m] = m\, \mathbb{E}[T_{(1)}]. \tag{36}$$

With $T = \sum_\ell E_\ell$ and independent $E_\ell \sim \mathrm{Exp}(r_\ell)$, $F_T(t) = ct^n + O(t^{n+1})$ with $c = \prod_\ell r_\ell / n!$, so $\Pr(T_{(1)} > b_m x) \to e^{-x^n}$ for $b_m = (cm)^{-1/n}$; a tail bound and dominated convergence give the mean, and $B_m = mT_{(1)}$ pathwise (Lawley, 2020). For $n > 1$ the wall time falls as $m^{-1/n}$ while the work grows as $m^{(n-1)/n}$; for $n = 1$ the work is constant.

Crowding: the sublinear regime. Let n be the smallest number of phases a research action passes through, from an initial phase to the exit towards a result. For a phase-type law with sub-generator Q, initial row β and success exit column s, the numbers $\beta Q^m s$ vanish for $m < n-1$ and $\beta Q^{n-1} s > 0$. Two consequences follow. For small t, $\Pr(T \leq t, J = \text{result}) = \int_0^t \beta e^{Qu} s \, \mathrm{d}u = \beta Q^{n-1} s \, t^n / n! + O(t^{n+1})$, so n is the slope of this curve on a log-log plot at short times. For large κ, $L_{\mathrm{s}}(\kappa) = \beta(\kappa \mathbf{I} - Q)^{-1} s = \sum_{m \geq 0} \beta Q^m s \, \kappa^{-m-1} = A\kappa^{-n} + O(\kappa^{-n-1})$ with $A = \beta Q^{n-1} s$, and $\tau_{\mathrm{R}}(\kappa) = \kappa^{-1} + O(\kappa^{-2})$. For the chain of Appendix D.4 with $j_0 = 0$, $n = k+1$ and $A = p\mu \prod_j a_j$.

Now let every agent research ready questions, $\alpha = 0$ and $\rho = 1$, with $N \gg M_{\mathrm{ready}}$, so that κ exceeds every rate of the law. Then the cycle time is $\bar{\tau} \simeq \tau_{\mathrm{R}}(\kappa) \simeq 1/\kappa$, each question receives entries at the rate $N\kappa / M_{\mathrm{ready}}$, and each entry yields a result with probability $A\kappa^{-n}$, so $W(Z \to Z^{+q}) \simeq NA\kappa^{1-n} / M_{\mathrm{ready}}$. The cancellation rate of a researcher is the success rate of the others on its question, $\kappa \simeq W(Z \to Z^{+q})$. Solving, $\kappa = (NA/M_{\mathrm{ready}})^{1/n}$ and

$$\lambda = M_{\mathrm{ready}} \kappa = (NA)^{1/n} M_{\mathrm{ready}}^{(n-1)/n}, \qquad \frac{N}{\lambda} \propto N^{(n-1)/n}, \tag{37}$$

within this homogeneous, frozen-input closure. The exponent agrees with the exact race above, where the first of m fresh researchers finishes in a time of order $m^{-1/n}$. If the cycle contains time that cancellation cannot shorten, posing or research on questions that are not ready, $\bar{\tau}$ tends to a positive constant as κ grows and the same argument gives $\lambda \propto N^{1/(n+1)}$ beyond a second crossover.

Keeping phases and allowing agents to arrive at different times. Consider a controlled arrival model with fixed action laws, no positive minimum duration and no time spent away from ready research. Each agent follows the full PH chain after joining a question; its successful completion curve is $G(t) = \int_0^t \beta e^{Qu} s \, \mathrm{d}u \sim ct^n$, with $c = A/n!$. Approximate the initial batch by a Poisson count of mean b, and later arrivals by an independent Poisson process of rate a. The probability that the question is still open at time t is then

$$S(t) = \exp\left[-bG(t) - a \int_0^t G(u) \, \mathrm{d}u \right]. \tag{38}$$

This retains the ages of the attempts without replacing them by a stationary success rate. For exactly b initial agents, the first factor would instead be $[1 - G(t)]^b$. The Poisson description is an arrival approximation for a fixed population, not its exact projection.

With $b = 0$, the short-time form is $S(t) \simeq \exp[-act^{n+1}/(n+1)]$. The lifetime U of the question therefore has a scale $\ell = [(n+1)/(ac)]^{1/(n+1)}$, with $\mathbb{E}[U] \propto \ell$ and $\mathbb{E}[U^2] \propto \ell^2$. At a high arrival rate almost all entrants stay until the first result, so the mean number present over a question's life is, to leading order,

$$m \simeq \frac{a\mathbb{E}[U^2]}{2\mathbb{E}[U]} \propto a\ell.$$

In a fixed population that keeps all agents busy on a fixed pool of equally difficult ready questions, $m = N/M_{\text{ready}}$. Eliminating a gives

$$\mathbb{E}[U] \propto (cm)^{-1/n}, \qquad \lambda = \frac{M_{\text{ready}}}{\mathbb{E}[U]} \propto N^{1/n} M_{\text{ready}}^{(n-1)/n}. \tag{39}$$

An initial batch with b proportional to m adds a term bct^n to the exponent of $S(t)$ and preserves this scaling, while changing its coefficient. Thus keeping the phases and arrival times can change the waiting-time coefficient without changing the population exponent. For a fixed number of successive rounds with the same arrival rule, mean completion time scales as $N^{-1/n}$ and total agent time as $N^{(n-1)/n}$. This does not establish the same law for evolving pools or dependencies, and a positive minimum duration cuts off the short-time regime.

If arrivals are instead imposed externally with $a \propto N$ and no initial batch, then $\mathbb{E}[U] \propto N^{-1/(n+1)}$. That is a different population protocol: the exponent cannot be inferred from keeping the phases alone. Appendix E.5 tests the fixed-population arrival approximation in a designed environment.

E Experiments

E.1 Single-agent experiments

This subsection gives the settings and results summarised in Section 6.1. The agents were Claude Code with Haiku 4.5 and with Sonnet 5, Codex with gpt-5.5 at low reasoning effort, and the vanilla agent of NewtonBench (Zheng et al., 2025) with both Claude models. On the 55 named lemmas of the Natural Number Game (Buzzard & Eugster, 2023) in Lean, the knowledge base of $C = 8$ lemmas either contains the premises of the reference proof or only distractors; on the Equational Theories Project (ETP) (Bolan et al., 2025) a target is an implication between two laws, true or false; NewtonBench asks for a hidden physical law at three levels of difficulty. An action ends at the first accepted check on Lean, at the first verified proof or refutation on ETP, and with the agent's last turn on NewtonBench, or at a cap of the agent's own program, which is an exit without a result; outcomes are decided by running the check again. On Lean, Claude Code with Haiku and Codex ran twice, three days apart, and Claude Code with Sonnet once; on NewtonBench, Haiku ran twice and Sonnet once; on ETP, Claude Code attempted 24 implications with Haiku and 48 with Sonnet, and Codex the same 48. Figures 3 and A1 show the first run of each agent.

We fit within the hyper-Erlang subclass, mixtures of Erlang branches each ending in a result with its own probability; parameters are fitted by EM (Asmussen et al., 1996) and branch orders chosen by five-fold held-out likelihood (Thümmler et al., 2006). Duration and outcome are fitted jointly, since the success before cancellation depends on how the outcome varies with the duration. The fitted distribution must reproduce the joint law of duration and outcome, move with the declared inputs and not with undeclared ones, and agree across repeated runs.

Every fitted distribution lies within the 95% Dvoretzky–Kiefer–Wolfowitz band of the empirical distributions (Massart, 1990); each has two Erlang branches, and a single exponential phase is worse by 0.3 to 1.1 in held-out log-likelihood per action. The shape belongs to the agent and the task: on the same lemmas Codex and Claude Code with Sonnet are regular ($c_v^2 = 0.16$ to 0.33) and Claude Code with Haiku heavy-tailed ($c_v^2 = 1.0$ to 1.6); Sonnet is nearly deterministic on NewtonBench ($c_v^2 = 0.09$) and heavy-tailed on ETP ($c_v^2 = 1.9$). The declared input moves the distribution as expected: without the premises, Claude Code with Haiku succeeds with probability 0.85 instead of 0.98 and takes 161 instead of 53 s on average, and Sonnet and Codex slow from 14 to 20 s and from 33 to 46 s while still succeeding almost always. A phase-type distribution has no strict earliest time, so $t_{\min}$ is read from the data, and the short-time order n cannot be resolved with tens of attempts.

Figure A2 compares the two functionals of Eq. (8) for the two runs of each agent. A paired permutation test over $\kappa\,\mathbb{E}[T] \in [0.01, 10]$ finds no difference between the runs on Lean ($p > 0.06$), and

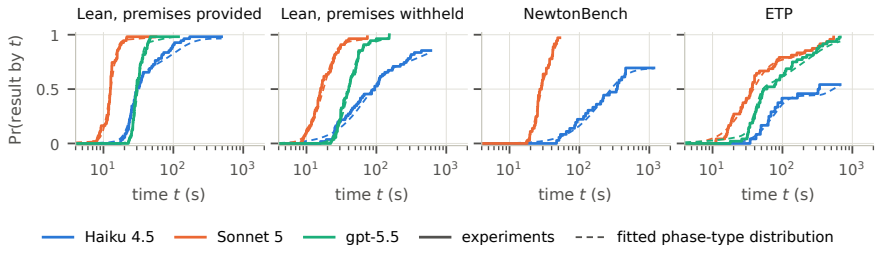

Figure A1: The attempts of Figure 3: fraction ended with a result by time t, on a logarithmic time axis. The curves level off at the success probability q_s; the gap to one is the fraction that ended without a result. Solid lines are the experiments, dashed lines the fitted hyper-Erlang distributions.

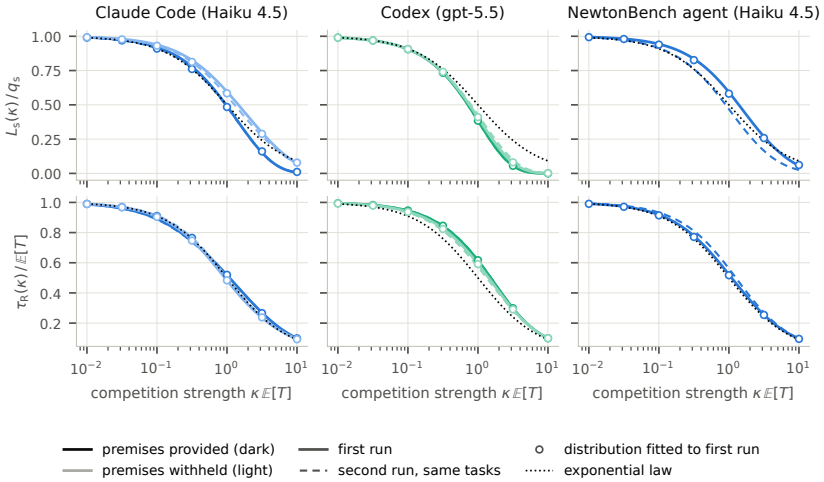

Figure A2: The two functionals of Eq. (8) through which the collective equations use the distribution of an action. Top: the probability of a result before cancellation, relative to the success probability, $L_s(\kappa)/q_s$; bottom: the mean time on the question until the action ends or is cancelled, relative to the mean duration, $\tau_R(\kappa)/\mathbb{E}[T]$. Lines are the two runs, circles the hyper-Erlang distribution fitted to the first run; for an exponential duration independent of the outcome both equal $1/(1+\kappa\,\mathbb{E}[T])$ (dotted).

the distribution fitted to the first run is within 0.03 of either run in L_s/q_s and $\tau_R/\mathbb{E}[T]$. The NewtonBench agent is the exception: its second run succeeded on 31 of 36 tasks against 25 ($p = 0.008$), which a change of the model service between the runs or the small number of tasks could explain.

E.2 Scaling experiments

This subsection gives the setting of the registered ETP campaigns summarised in Section 6.2; the campaigns with up to 256 agents are in Appendix E.3. The test world has 40 equational laws and 1560 implications, of which 439 are open at the start, 135 true and 304 false; it was generated only after the prediction program and its calibration were frozen. Claude Code with Sonnet 5 ran four campaigns at each of $N = 1$, 4 and 8. The coefficients come from the campaigns with $N = 1$, and the predictions for larger N were registered with file hashes before those runs. The times to settle 90% of the open questions at $N = 4$ and 8 fell inside the registered prediction intervals, 31% and 4% above the collapse in agent time, but the registered primary test failed, 5 of 8 checks inside their intervals where 6 were required, because fewer attempts were cancelled than predicted; with four campaigns per N the design resolves the collapse only to about $\pm 30\%$.

E.3 ETP WITH UP TO 256 AGENTS

The ETP campaigns of Figure 4 add one campaign at each of $N = 16$, 64, 128 and 256 on the same test world to the four at each of $N = 1$, 4 and 8. Campaigns with $N \geq 16$ were run in virtual time with 16 Claude Code processes: an attempt reads the record only when it starts, so it can be run whenever the record at its start is fixed, and results are applied in the order of their virtual times, which reproduces a simultaneous run exactly; a replay of the records checks every attempt. Each campaign ends when 90% of the open questions are settled. In the campaigns with $N \leq 8$ the agents started 2 s apart; in the others, and in the model, they start together.

The law of one agent is fitted by maximum likelihood on the attempts of the four campaigns with $N = 1$, 164 on true and 78 on false questions: a fixed delay of 7 s followed by a chain of four phases for true questions, and of 12 s followed by three phases for false ones, chosen by the Akaike criterion among chains of up to five phases and mixtures of two chains. The model integrates the mean phase occupancies of the coarse-grained dynamics with the phases kept (Section 4.2) and draws the results as events on the exact closure of the world; 200 runs per N give the median and the 90% range.

Least-squares fits of $\ln T_L$ against $\ln N$ give slopes, between $N = 1$ and 16, of -1.17 and -1.04 for the measured T_{50} and T_{80} and -0.99 and -1.11 for the model; between 16 and 256, of -0.28 and -0.59 measured and -0.30 and -0.57 for the model (dash-dotted in Figure 4), and for T_{90} of -0.52 measured and -0.62 for the model (Figure A3). The zeroth order, the jump process of Eq. (9) with the same laws, the delay included in their functionals, and self-consistent κ_q, keeps slopes near -1 and gives $T_{50} = 1.7$ s at $N = 256$, less than the delay of either law (dashed in Figure 4). Only a refutation settles a false question, and 304 of the 439 open questions are false, so settling half of them needs refutations, and the floor of Section 5 for T_{50} is the 12 s delay of the law of a refutation. The collapse line $T_{50}^{(1)}/N$, with the measured $T_{50}^{(1)} = 647$ s, reaches this floor near $N = 50$; the measured T_{50} is 15.0 s at $N = 128$ and 15.7 s at $N = 256$.

The onset of competition follows from Eq. (17). In this world one result settles many questions through the closure (the first refuting table of the campaign with $N = 128$ settled 104), so the rate r_L is taken from the four campaigns with $N = 1$: for half, four fifths and nine tenths of the open questions, $T_L^{(1)} = 647$, 2555 and 3390 s and the logarithm in r_L is $\ln 2$, $\ln 5$ and $\ln 10$. With $\mathbb{E}[T] = 88$ s, the mean duration of an attempt of the fitted law, the laws of true and false questions weighted by the attempts of one agent on each, this gives $N_{50}^* = 12$, $N_{80}^* = 19$ and $N_{90}^* = 18$, where the measured and the modelled curves bend (Figures 4 and A3). The estimate $M_{\mathrm{ready}}/q_{\mathrm{s}}$, which counts one question per result, would be 225, 90 and 45 with M_{ready} the open questions left at these levels. The share of the attempts cancelled by the results of other agents is 18%, 36%, 56%, 70%, 85% and 92% at $N = 4$, 8, 16, 64, 128 and 256.

The time to settle 90% of the open questions exceeds the model from about 16 agents on, by a factor of about two for $N \geq 64$ (Figure A3). Figure A4 shows where the difference lies. At 90% the measured campaign with $N = 128$ leaves 41 true and 2 false questions open, a run of the model 20 and 15. The open true questions follow from 6 of the 33 direct implications of the order of the laws; once one of them is proved, the closure settles many others. A single agent seldom proves these implications directly: in the four campaigns with $N = 1$ they were established by a direct proof in 55 of 164 cases and otherwise through the closure of other results, and a direct attempt on them took a median of about 100 s, against 38 s for all true questions of these campaigns. The fitted law of one agent therefore averages them with easier questions. With many agents the easy questions go first, most agents then work on these implications at once, and few remain for the other open questions; the last tenth is set by questions whose own law is slower than the average.

E.4 CHECKS OF THE ALGEBRA

Two scripts check the algebra: `paper-scripts/check_phase_reduction.py` compares the marked resolvents, the renewal throughput and the phase dynamics with direct finite-chain solutions and checks the restart counterexample and the race integral; `paper-scripts/check_ch4_search.py` simulates the reference search and reproduces Eq. (32), the bounds of Eq. (33) and the throughput of Eq. (25) within two per cent. The chain approximation errs by 4 to 40% when one attempt loads a sizable part of the record and overestimates success by a factor of 8 at $C = 4$ when held premises are evicted, falling to 2.5% at $C = 200$;

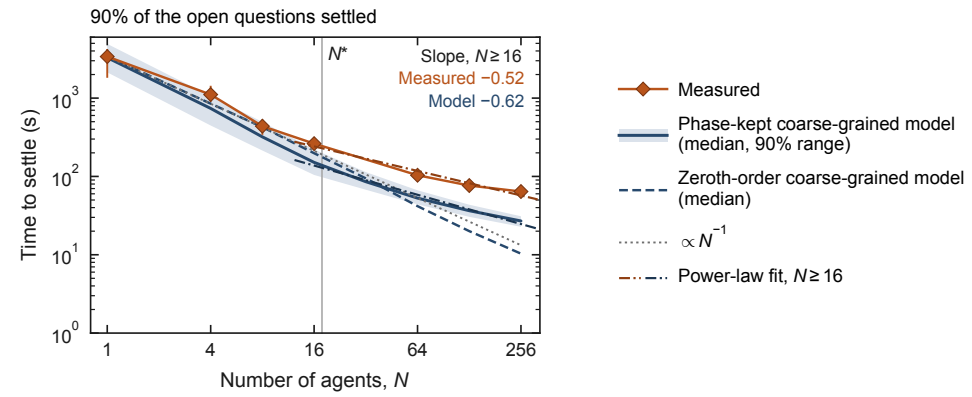

Figure A3: Time to settle 90% of the 439 open questions of the ETP world against N, measured and with the two orders of the coarse-grained model, as in Figure 4; the vertical line is N^*_{90} of Eq. (17).

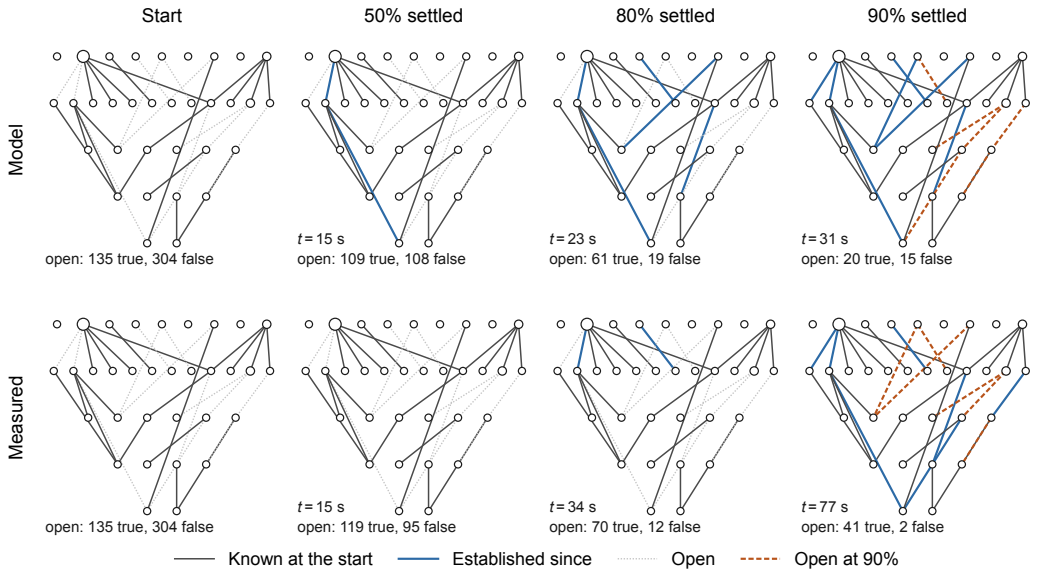

Figure A4: The order of the laws of the ETP world filling in, for a run of the model (top) and the measured campaign (bottom), both with $N = 128$. Nodes are laws, laws that imply each other merged; edges are the direct implications, the stronger law above. Snapshots at the start and when 50%, 80% and 90% of the open questions are settled; each panel gives the time and the numbers of open true and false questions.

with two researchers per ready question the cancellation rate exceeds $(N-1)\lambda/(NM_{\text{ready}})$ by up to 14%.

E.5 Crowding in a designed corridor world and on FrozenLake

The crowded exponent $1/n$ needs an action law whose success curve starts as t^n, which random explorers in Gymnasium lack: they can succeed at their first step. We therefore built a corridor world with the step law of a one-row slippery FrozenLake-v1 map (Towers et al., 2024): pushing right advances one tile with a small probability and otherwise leaves the agent in place, so an attempt on a question n tiles away is a sum of n geometric waits. Independent corridors fix the number of ready questions, attempts never fail, and all questions are alike. From single-agent runs and the

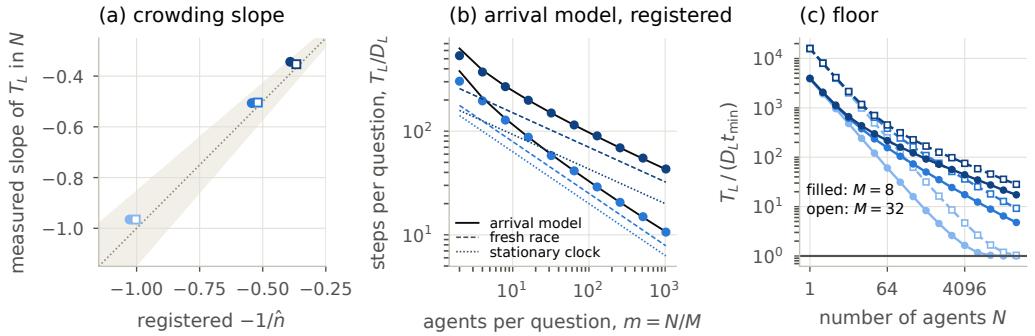

Figure A5: Designed corridor world. (a) Slope of T_L in N over the registered windows against the registered $-1/\hat{n}$, with 90% intervals; shaded: the registered tolerance; filled and open markers: two numbers of ready questions. (b) Registered test of Eq. (40) on new worlds: time per question against agents per question, runs (dots), arrival model (black), fresh race (dashed) and stationary clock (dotted), for $n = 2$ and 3. (c) T_L relative to the minimum $D_L t_{\min}$ for $n = 1, 2, 3$.

transition table we registered, before the multi-agent runs, the collapse below N^*, the slope $-1/\hat{n}$ of T_L in N, and the prefactors of the stationary clock (Eq. (37)) and of the fresh race (Eq. (36)). The slopes and the position of the onset agree with the registration, only $n = 1$ reaches the minimum time per question, and both prefactors are too fast (Figure A5).

Arrival on a question. In the crowded regime a question lives about as long as one attempt, so its phase occupancies need not relax before it is solved, and agents do not start together. Approximate the initial batch on a new question by a Poisson count of mean b and subsequent entries by independent Poisson arrivals at rate a. With $x(t)$ the mean occupancies of Eq. (13) for this one question, $x(0) = b\beta$ and entry rate a, independent agents leave it open after time t with probability

$$S(t) = \exp\Big(-\int_0^t x(u)\, s\, \mathrm{d}u\Big) = \exp\Big(-b\, G(t) - a\int_0^t G(u)\, \mathrm{d}u\Big), \qquad G(t) = \Pr(T \le t,\ J = \text{result}). \tag{40}$$

For exactly b initial agents, the first factor is $[1-G(t)]^b$; the exponential form above uses the Poisson approximation. In the corridor b is the mean share $1/M$ of the agents released by the previous question, and a follows from Little's law: the mean crowd over the life of a question equals N/M. No parameter is fitted. Registered for new corridor worlds before they were run, the model predicts the time per question except for crowds of a few agents, while the fresh race and the stationary clock are too fast; all three give the exponent $-1/n$ (Figure A5b).

The corridor world is built to satisfy the assumptions of Appendix D.5, so it checks the reduction numerically; it does not test q_s, heterogeneous questions, time that cancellation cannot shorten, or whether n can be measured for language-model agents.

On standard FrozenLake maps with dense holes, with and without slip, an attempt can succeed at its first step, so n is undefined and no crowded power law was predicted or observed. Beyond a few N^* the time falls to the depth floor $D_L t_{\min}$ computed from the map (Figure A6). On the slippery lake the floor already binds near N^*, so the registered test that the onset moves with $1/q_s$ failed there.

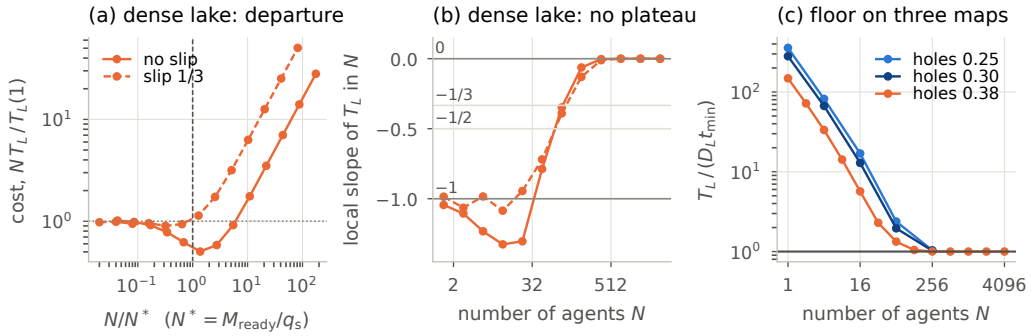

Figure A6: FrozenLake with dense holes. (a) Cost of the target relative to one agent against N/N^*, without and with slip. (b) Local slope of T_L in N: from -1 to 0, with no plateau at $-1/2$ or $-1/3$. (c) T_L relative to the depth floor on three maps without slip; the two sparser maps are the earlier runs.