REALFIN: Benchmarking Agent Skills for Multi-Tool Orchestration under Tool Explosion

Anonymous ACL submission

Abstract

Large language models are increasingly moving from answering questions to acting as agents that plan actions, call tools and operate in external environments. Solving tasks in these settings requires constructing and executing multi step workflows rather than producing a single response. The key difficulty lies not only in computation but in choosing appropriate tools and organizing them into effective workflows, since realistic environments contain many tools with related yet non identical functions and several plausible solution paths may coexist for the same task. However, many existing benchmarks primarily emphasize static question answering and final outputs, often using simplified tool environments and short workflows, which provides limited visibility into how tools are chosen and workflows are executed. A natural direction is therefore to evaluate agents through the processes they carry out in realistic environments where overlapping tools and genuine workflows are present. We introduce REALFIN to support this direction. REALFIN contains tasks derived from authentic financial analysis workflows, pairs each task with an executable reference program and operates in realistic multi tool environments while recording execution traces. It enables joint evaluation of outcomes and execution behavior and supports process oriented assessment of tool choice, workflow organization and temporal grounding in addition to answer accuracy.

1 Introduction

Large language models are increasingly moving from answering questions to acting as agents. Modern systems do not only produce textual outputs but plan actions, call tools and operate within external environments (Yao et al., 2023; Shinn et al., 2023; Wang et al., 2024). In these settings, obtaining an answer requires constructing and executing multi step workflows rather than performing a single step generation. This shift brings evaluation closer to the way real tasks are completed in practice.

The key difficulty in this agent setting lies in choosing and organizing tools rather than in computation alone. Realistic environments contain many tools whose functions are related yet not identical, and several distinct workflows may appear reasonable for the same task. Multiple solution paths therefore coexist and the agent must decide which tools to invoke and how to arrange them. This phenomenon, in which overlapping tools create abundant and plausible workflows, can be viewed as tool explosion and it places central importance on agent skills in tool selection and workflow orchestration (Qin et al., 2024a,b; Patil et al., 2024).

However, existing benchmarks do not primarily target this form of agent behavior. Many widely used evaluations are organized as static question answering (Chen et al., 2021, 2022; Zhu et al., 2021; Deng et al., 2022; Guo et al., 2025) tasks that emphasize final answers, providing limited visibility into how tools are chosen and workflows are executed. In a number of settings the tool environments are simplified (Qin et al., 2024b; Liu et al., 2024) and the workflows required to complete tasks are short. These benchmarks serve important purposes but focus on different aspects of capability from those that arise when rich tool ecosystems and extended workflows are naturally present.

A natural direction is therefore to evaluate agents through the processes they execute in realistic environments. When tools overlap in function and workflows are required to reach answers, the behavior of interest is the way agents plan, choose tools, ground temporal references and coordinate intermediate steps. Evaluations that observe such processes complement answer accuracy and offer a more complete view of the abilities that underlie reliable tool use.

We propose REALFIN from this perspective. REALFIN consists of tasks constructed from authentic

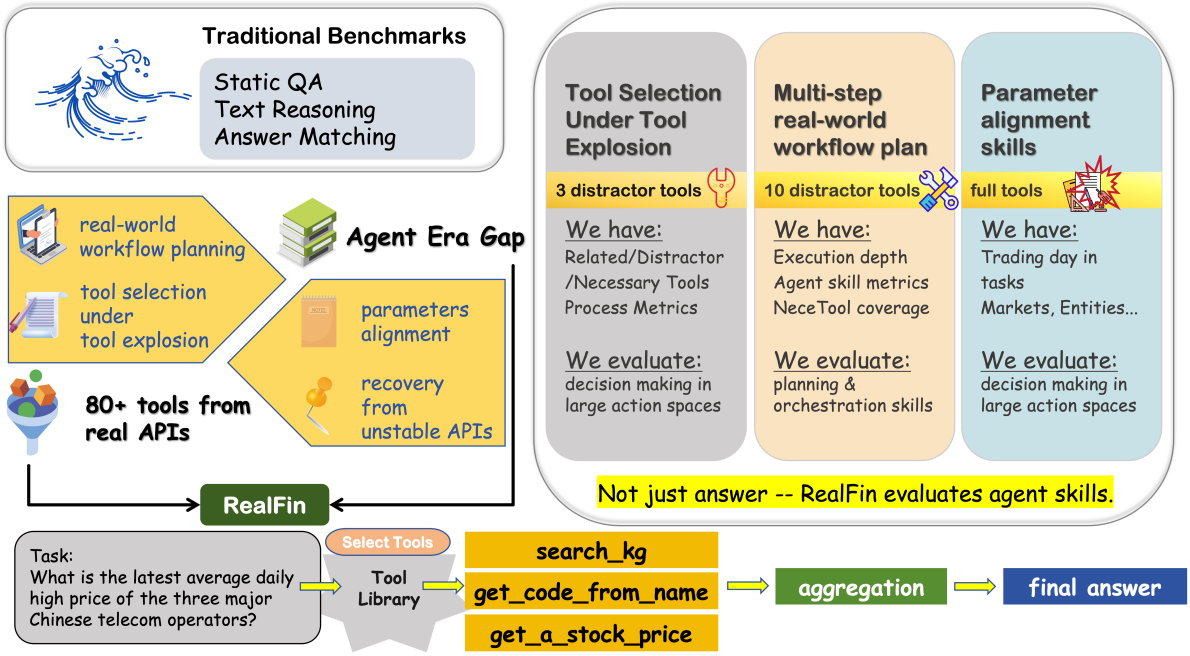

Figure 1: Overall concept of REALFIN: workflow-grounded financial tasks, executable golden code, and a realistic multi-tool environment with tool explosion.

financial analysis workflows that require multi step tool use. Each task is paired with an executable reference program that specifies one valid solution path while still permitting alternative workflows in practice. The benchmark is built on realistic financial application programming interfaces and records execution traces, enabling the study of both outcomes and processes in environments where tool abundance arises naturally. As shown in Table 1, REALFIN is positioned against prior benchmarks in terms of tool scale, workflow realism, executable evaluation, and process-level metrics.

Our contributions are threefold.

- We advance a workflow-grounded benchmark for financial agents, framing tasks as complete analysis workflows rather than static QA, instantiated in REALFIN.
- We show that realistic financial APIs naturally give rise to tool explosion, and we instantiate this phenomenon in a live multi-tool environment with redundancy and overlapping functionality.
- We promote a process-based view of agent evaluation, where models are assessed through their tool-use behavior rather than only final answers, and REALFIN makes this view measurable under tool explosion.

2 Related Work

2.1 Financial reasoning benchmarks

Existing financial benchmarks mainly evaluate static question answering and numerical reasoning rather than executable workflows. Representative datasets such as FinQA, ConvFinQA, TAT-QA and FinEval (Chen et al., 2021, 2022; Zhu et al., 2021; Guo et al., 2025) are built on fixed corpora with a single golden answer and do not require real API interaction. More recent financial-agent benchmarks such as InvestorBench and FinToolBench move closer to realistic financial decision-making and tool use, but emphasize broad decision tasks or large-scale tool execution rather than process-oriented diagnosis of workflow construction under overlapping tools (Li et al., 2025). REALFIN instead focuses on executable multi-step workflows, tool redundancy, and fine-grained process evaluation in live financial environments.

2.2 General tool-use and agent benchmarks

General agent benchmarks study tool use and environment interaction through multi-tool or web-based tasks (Qin et al., 2024a; Li et al., 2023; Zhou et al., 2024; Liu et al., 2024; Mialon et al., 2024). Related agent frameworks such as ReAct and Reflexion highlight the importance of interleaving reasoning with actions and improving be-

Benchmark	Tool Selection	Multi-Tool Workflow	Real-API Robust-ness	Temporal Reasoning	Alt. Tool Paths	Tool Propen-sity
FinQA	✗	✗	✗	✗	✗	✗
ConvFinQA	✗	✗	✗	✗	✗	✗
FinanceBench	✗	✗	✗	△	✗	✗
FinBen	△	△	✗	△	✗	✗
FinSearchComp	✓	△	△	✓	✗	△
ToolBench	✓	✓	△	✗	✗	✓
GAIA	✓	△	△	△	✗	△
BFCL	✓	△	△	✗	✗	✓
RealFin (Ours)	✓	✓	✓	✓	✓	✓

Table 1: Benchmark coverage across capability dimensions. Symbols: ✓ = supported, ✗ = not supported, △ = partially or indirectly supported. Dimensions: Tool Selection = choosing correct tools from large catalogs; Multi-Tool Workflow = composing multi-step domain workflows; Real-API Robustness = robustness to unstable real-world APIs and execution noise; Temporal Reasoning = grounding to trading calendars and temporal anchors; Alt. Tool Paths = support for multiple valid executable tool workflows; Tool Propensity = tendency to invoke tools instead of zero-tool answering.

havior through feedback, while Gorilla studies accurate API invocation at scale (Yao et al., 2023; Shinn et al., 2023; Patil et al., 2024). Many existing benchmarks employ simulated or mirrored APIs for reproducibility. Tool explosion usually appears only as a byproduct rather than a primary focus, whereas our work explicitly targets agent behavior in redundant multi-tool environments with executable workflows and process-level analysis.

3 Benchmark Design

3.1 Workflow-Driven Task Construction

Workflow grounded task design We construct REALFIN directly from full financial analysis workflows used by human analysts. Tasks are decomposed into core analytical intents intermediate variables and dependency structures. Each instance is aligned with an executable solution path so the benchmark evaluates planning and execution under tool explosion instead of single step answering.

Template based large scale task generation Workflow skeletons are transformed into parameterized templates and instantiated under financial validity constraints. Tasks are executed in a real API environment that preserves overlapping and substitutable tools so multiple valid tool chains may solve the same task.

Reflecting tool explosion through a redundant multi tool environment We construct a multi tool execution environment from real financial APIs (Akshare (King, 2022)) that intentionally preserves functional overlap and substitutable tools. The tools are almost freely accessible, and Only one endpoint requires an API key with a sufficient free quota. Tasks are not bound to a single correct tool. The golden workflow specifies one valid solution path while alternative paths may also exist. This embeds tasks in a naturally redundant tool space so that tool explosion arises from the environment itself rather than from artificial task design.

Executable alignment of tasks and solutions Every task in REALFIN is paired with an executable golden program. The program encodes the full solution process including entity resolution, tool selection, parameter planning, temporal alignment, and aggregation. Gold trajectories, derived from real analyst workflows, are encoded as executable reference programs to ensure correctness and reproducibility. These trajectories reflect realistic procedures rather than imposing robustness assumptions, and multiple valid workflows may exist. Agents are not required to follow a single reference path, as many tasks admit alternative executable paths with comparable outcomes. Programs are executed against real APIs, and samples

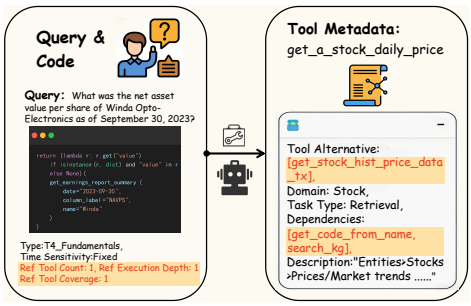

Figure 2: Data and tool structure in REALFIN. The left side illustrates the task-side representation including query text and executable code, while the right side shows the structured tool metadata such as domain, task type, dependencies, and alternatives.

that are non-executable, unstable, or ambiguous are removed. This strict alignment enables process-level evaluation instead of answer-only scoring, offering a more comprehensive measure of agent performance.

Use of AI Assistants in Methodology During the preparation of this paper, large language models were used as assistants for various tasks. However, all technical ideas, dataset design, experiments, and conclusions were conceived and verified by the authors.

3.2 Dataset Structure and Properties

This section overviews the structural properties of REALFIN, including its data fields, tool specifications, task distributions, and comparison with representative benchmarks. Figure 2 illustrates the joint design of the data and tool structures.

Temporal structure Tasks are categorized into Fixed, Relative, and Realtime classes, introducing trading-calendar alignment, rolling windows, and data drift. Temporal grounding thus becomes an evaluated agent skill rather than simple post-processing.

Tools and multi-path structure REALFIN contains about 85 overlapping financial tools, forming a shallow hub-dominated dependency graph with multiple valid tool chains per task. Tool explosion therefore arises naturally, requiring agents to select and orchestrate tools rather than merely produce final answers.

Table 2 shows that REALFIN uniquely combines real APIs, temporal sensitivity, multi-tool redundancy, and executable workflows, making it suitable for evaluating agent skills under tool explosion. We further verify that golden workflows align with real analyst practice by comparing them with standard operating procedures (Appendix D).

Benchmark scope REALFIN is designed as a diagnostic benchmark for studying agent behavior under tool explosion rather than proposing new architectures. It isolates the effect of tool-space complexity on decision-making and provides a testbed for future methods. The hierarchical packaging strategy serves as a lightweight structural aid rather than a full solution, and its limited improvements highlight the intrinsic difficulty of dense tool ecosystems.

Generality beyond finance Although instantiated in finance, the core challenge—decision-making under redundant tools and multi-step workflows—is domain-independent. Similar settings appear in software engineering toolchains, web environments, and scientific workflows. Finance provides a realistic, temporally grounded environment for observing these behaviors, and the resulting insights extend beyond the financial domain.

3.3 Quality Control and Evaluation Protocol

We implement a quality-control and evaluation protocol that secures reliable dataset construction, executable workflows, and reproducible and interpretable evaluation outcomes.

Data quality control We apply validity checks to all template-instantiated tasks, filtering nonexistent or ambiguous entities, invalid market codes, non-trading days, and insufficient historical windows. Each task is paired with a golden reference program executed against real financial APIs; instances that fail or return empty results are repaired or removed. Golden programs are periodically rerun to verify output consistency and monitor temporal drift using rule-based checks and selective manual inspection. As a result, every sample is semantically valid and end-to-end executable. Execution success exceeds 97%, with high output stability; only a small fraction of temporally labeled tasks change across dates due to genuine market non-stationarity.

Evaluation protocol We adopt a process-aware evaluation protocol rather than answer-only scoring. Besides outcome accuracy, we compute process metrics from execution traces, including entity binding, temporal grounding, tool selection, routing order, parameter planning, and multi-step orchestration, enabling diagnosis of agent skills

Table 2: Benchmark comparison. RealFin denotes the full benchmark, while RealFin-Hard denotes the subset of tasks whose reference workflows require at least five tool calls.

Benchmark	Domain	#Tasks	#Tools/APIs	Avg Tool Calls	Avg Exec Depth	Exec	Time-Sens.	Answer type
RealFin	Finance (tool)	1524	85	3.78	2.12	✓	✓	Fast Live exec
RealFin-Deep	Finance (tool)	482	85	6.10	2.10	✓	✓	Live exec
API-Bank	Tool-aug LLM	274	53	2.07	≈1.50	✓	×	Simu APIs
BFCL	General	5551	/	≈2.55	≈1.50	✓	×	AST
ToolBench	General	126,486	16,464	3.71	≈2.00	✓	×	Static
PACIFIC	Fin conv-QA	2757	Document	0.00	≈1.00	×	×	Static
WebArena	Web agent	812	Web	≈2.8	≈2.80	✓	✓	Rule-based
ConvFinQA	Fin num-reason	14115	computation	≈1.00	≈1.50	✓	×	Static
FinQA	Fin num-reason	8281	computation	≈2.00	≈1.50	✓	×	Static
FinSearchComp	Fin QA	635	search engine	1–2	1–2	×	✓	LLM-Judge

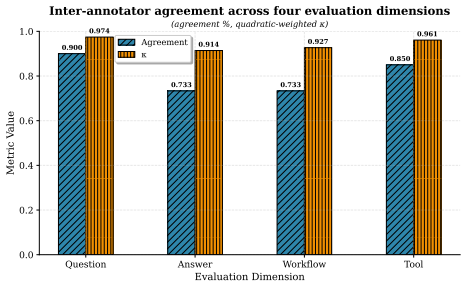

Figure 3: Inter annotator agreement across four evaluation dimensions

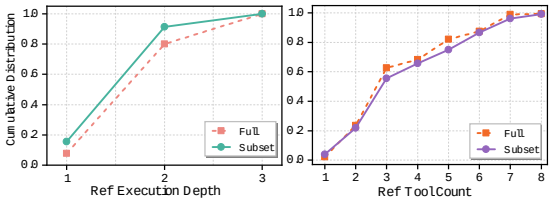

Figure 4: **Structural similarity between the full benchmark and the 128-core subset.**

and failure modes. Because real APIs may fail, we report metrics on both all traces and the subset where every API call succeeds. REALFIN supports snapshot evaluation for reproducibility and live evaluation to preserve non-stationarity; both produce consistent qualitative conclusions.

3.4 Human validation of data quality

We conduct a human validation study to assess the reliability of questions answers workflows and tools in the benchmark. The questionnaire is in Appendix F.

We sample 60 tasks covering diverse workflow depths and task types, each independently reviewed by two annotators (120 double-annotated instances). Twelve annotators with backgrounds in quantitative finance and computer science evaluate four dimensions—question clarity, answer correctness, workflow correctness, and tool usability—using a 0–1–2 ordinal scale. As shown in Figure 3, raw agreement ranges from 0.73 to 0.90 and quadratic weighted κ reaches 0.91–0.97, indicating near-perfect agreement. Most disagreements occur between scores 1 and 2 rather than 0 and 2, suggesting differences in strictness rather than validity judgments. These results confirm that the tasks, answers, workflows, and tools are consistently interpreted by annotators.

4 Experiments

This section evaluates whether large language models possess the agent skills needed to complete financial workflows under tool explosion. Agent skills include tool selection in redundant tool sets, multi-step workflow orchestration, semantic grounding of entities and time, parameter planning, and error recovery. We vary tool exposure (Full, BM25, Oracle, OraA-3, OraA-10), add structured assistance (package routing and workflow hints), and analyze process-level metrics and tool traces. Results show systematic deficiencies in these skills within large and functionally overlapping tool ecosystems.

4.1 Experimental Setup

Evaluation subsets. We evaluate on a 128-task core subset. Its execution-depth and tool-count distributions match the full benchmark (Figure 4), with no significant differences by Kolmogorov–Smirnov tests ($p = 0.088, 0.548$). The subset therefore serves as a compact but faithful representation of the full dataset.

(1) Baselines. *Full* exposes all tools to the model, simulating real systems and the strongest level of tool explosion. *BM25* exposes only the Top-20 tools retrieved by BM25, simulating retrieval-based tool selection and testing whether shallow tool discovery mitigates tool explosion.

Table 3: Evaluation metrics.

Metric (Abbrev.)	Definition
Outcome-level metrics	
Outcome Accuracy (Acc)	Answer-only.
Execution Success Rate (ESR)	Calls succeed.
Process-level metrics	
Tool Selection Recall (TSR)	Correct tools succeed.
Ordering Accuracy (OA)	Call order right.
Entity Grounding Accuracy (EGA)	Entities correct.
Temporal Grounding Accuracy (TGA)	Time correct.
Parameter Planning Success (PPS)	Parameters correct.
Orchestration Success Rate (OSR)	Workflow correct.
Recovery Success Rate (RSS)	Errors fixed.

(2) Controlled exposure. *(i) Oracle (Ora).* Expose only tools used in the reference workflow, largely removing tool explosion and isolating execution, parameter planning, and temporal grounding. *(ii) OraA-k.* Extend Oracle with functionally similar alternatives and k distractors, increasing tool redundancy while keeping task semantics fixed. We report OraA-3 and OraA-10.

(3) Package routing. *Pkg* groups tools by functionality; the model first selects a package and then a tool within it. This provides limited structural guidance and examines whether hierarchical routing alleviates difficulty in large tool spaces.

(4) Unfair oracle-guided setting. In this unfair oracle-guided setting, the model answers with oracle trajectory hints including the correct tool set and solution path but not the final answer.

4.2 Metrics

We report both outcome-level and process-level metrics. Outcome metrics, such as ACC and ESR, assess task completion and execution reliability. Process metrics are derived from structured tool-use traces, which capture the sequence of tool invocations, intermediate steps, and decisions made by the agent during task execution. These traces allow us to decompose the agent's behavior into key components, such as tool selection (TSP), workflow orchestration (OA, OSR), parameter planning (PPS), temporal and entity grounding (TGA, EGA), and recovery from failures (RSS). Each of these metrics evaluates specific aspects of the agent's decision-making process, providing a detailed understanding of its performance in handling tool explosion. For detailed mathematical definitions of the metrics, please refer to Appendix C.

4.3 Main results

Table 4 reports overall performance on the RealFin-128 core subset, evaluating agent skills including tool discovery and selection, multi-step workflow orchestration, parameter planning, entity and temporal grounding, and error recovery.

Tool explosion and systematic skill degradation. Tool explosion refers to environments where multiple functionally related tools coexist for the same sub-task and must be composed into multi-step workflows, rapidly expanding the decision space. As shown in Table 4, performance declines consistently as tool-space size increases: from Oracle to OraA-10, ACC drops from roughly 20–35% to 5–15%, while OSR falls from about 5–15% to near zero for most models. Similar degradation appears across PPS and RSS, indicating that multiple agent skills—including tool selection, orchestration, parameter planning, and recovery—deteriorate systematically under tool explosion.

Controlled tool-space scaling. To isolate the effect of tool-space size, we keep tasks and reference workflows fixed while progressively increasing both functionally related and distractor tools. Figure 6 shows that ACC, OSR, and PPS decrease steadily from Oracle to OraA-3, OraA-10, and Full, with a clear threshold effect: changes from Oracle to OraA-3 are moderate, while the drop from OraA-3 to OraA-10 and Full is much sharper. This indicates that behavioral degradation is driven primarily by the structure of the tool space rather than by task difficulty or data noise.

Mitigation via tool-space structure. Reducing the candidate pool with BM25 retrieval partially restores performance: compared with Full, ACC and OSR increase by roughly 2–4×, with PPS improving similarly. Package routing, which keeps the number of tools unchanged but organizes them hierarchically, also improves results: OSR increases from near zero to roughly 10–20%. These findings suggest that structured interfaces mitigate decision complexity without reducing tool availability.

Answer correctness vs. workflow completion. ACC and OSR are often decoupled: ACC may reach 20–40% while OSR remains in single digits or near zero. Many correct answers therefore arise from guessing or partial tool use rather than correct end-to-end workflows, showing that answer-level metrics alone overestimate agent capability.

Task solvability and robustness of results. With oracle trajectory hints, models achieve 83% accuracy, confirming that tasks are solvable and that most difficulty lies in autonomous tool discovery

Table 4: Model performance across tool-exposure regimes. ACC = outcome accuracy; OSR = orchestration success; PPS = parameter-planning success; RSS = recovery success. Model versions: Claude (claude-4.5-sonnet), GPT (gpt-5.1), Gemini (gemini-3-pro), Qwen (qwen-max), DeepSeek (deepseek-v3.2), Doubao (doubao-seed-1.6), Kimi (kimi-k2).

	Baseline								Constrained tool spaces												Package-route			
Model	Full*				BM25*				Oracle				OraA-3				OraA-10				pkg			
	ACC	OSR	PPS	RSS	ACC	OSR	PPS	RSS	ACC	OSR	PPS	RSS	ACC	OSR	PPS	RSS	ACC	OSR	PPS	RSS	ACC	OSR	PPS	RSS
Cluster A																								
Claude	9.4	0.0	0.8	8.5	25.0	14.1	29.7	36.6	27.3	3.1	13.3	9.9	36.7	14.1	48.4	17.9	26.6	1.6	13.3	9.8	23.4	17.2	18.0	8.3
GPT	1.6	0.0	0.8	9.4	6.4	4.7	11.7	26.0	10.9	2.3	7.8	5.8	8.6	11.7	35.2	15.4	4.3	0.0	2.9	11.1	30.5	15.6	18.0	6.6
Gemini	8.6	0.0	0.8	8.5	10.0	5.5	7.8	13.8	14.8	2.3	12.5	9.1	9.4	0.8	9.4	11.4	7.0	1.6	1.6	4.9	38.3	12.5	17.2	10.7
Cluster B																								
Qwen	3.1	0.0	2.3	6.7	9.6	2.3	6.2	15.4	16.4	1.6	14.1	9.1	14.1	10.9	35.2	13.8	4.7	0.0	3.1	5.7	23.4	13.3	14.8	10.7
DeepSeek	30.5	0.0	8.6	11.6	9.4	18.0	26.6	35.8	32.0	5.5	13.3	8.3	3.9	0.0	0.0	36.4	0.8	0.0	0.0	30.2	19.5	13.3	14.1	9.9
Doubao	33.6	0.8	16.4	6.6	11.7	7.8	23.4	37.4	36.7	4.7	19.5	9.1	0.0	0.0	0.0	31.4	0.0	0.0	0.0	26.4	42.2	11.7	14.8	11.6
Kimi	10.2	0.8	7.8	7.4	4.7	13.3	21.1	30.9	30.5	0.8	14.1	8.3	1.6	0.0	0.0	17.4	0.0	0.0	0.0	15.7	14.1	2.3	4.7	9.1
Unfair oracle-guided setting's ACC $\approx 83.59\%$																								

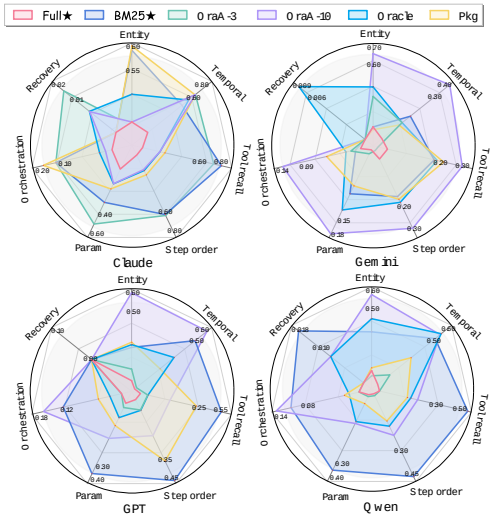

Figure 5: Agent-skill radar visualization across seven process metrics. Outer envelopes correspond mainly to OraA-3 and coarse-pkg settings, while Full and OraA-10 settings contract sharply toward the center.

and workflow planning. Bootstrap analysis (Table 17, Appendix E) further shows that observed differences are statistically robust: we identify 19 large effects ($|\Delta| \geq 0.05$, $|z| \geq 2$), while execution metrics drop by 30–76%, indicating systematic degradation rather than sampling noise.

4.4 Agent-skill-oriented analysis

To move beyond answer-level metrics, we analyze agent behavior in REALFIN using process-based measures. The radar plot in Figure 5 summarizes seven dimensions of agent capability, including tool recall, step consistency, parameter planning, entity and temporal grounding, workflow execution success, and error recovery.

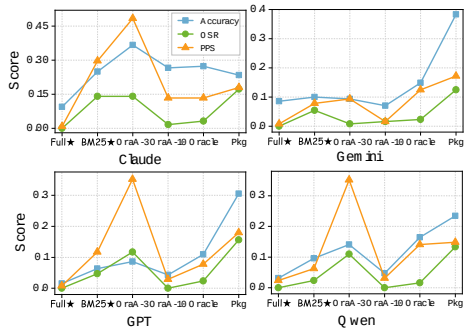

Figure 6: **Effect of distractor intensity.** ACC, OSR, and PPS under Oracle, OraA-3, OraA-10, Full, and BM25 on the 128-core subset. Performance peaks at OraA-3 and drops sharply as distractors increase.

Tool selection and limited replanning dominate failures. Tool recall and step consistency decline together across exposure settings, indicating that many ordering errors originate from incorrect tool selection rather than step planning itself. Tool selection therefore acts as a root capability for workflow construction, and failures propagate to later orchestration stages as the tool space expands. Error recovery also remains low and deteriorates under tool explosion: once models commit to an incorrect tool or path, they rarely abandon it and replan, behaving more like one-shot planners with limited reflection.

Process-level analysis. Process-level metrics reveal key challenges in multi-step financial workflows. As shown in Table 12, ESR is generally high while ACC remains much lower, indicating that errors mainly arise after tool execution, such as in parameter construction or result interpretation. TSR remains low even under Oracle conditions, revealing systematic underuse of available tools. EGA consistently exceeds PPS, suggesting

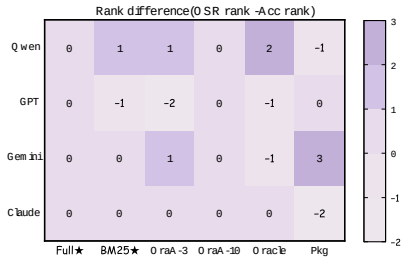

Figure 7: **Rank–flip heatmap.** Difference between OSR-based and ACC-based model rankings across exposure regimes. Large flips occur in moderate tool spaces; very large spaces flatten OSR.

that entity grounding is easier than parameter planning, while low OSR highlights the difficulty of full workflow orchestration. Additional analysis in Table 18 (Appendix E) further shows that many failures originate from tool retrieval in large action spaces, and even with correct tools, agents often struggle with ordering, parameter selection, and temporal grounding.

Execution reshapes model rankings. Model rankings induced by ACC and OSR diverge substantially (Figure 7). The gap is largest under medium tool spaces (Oracle, OraA-3), where models may answer correctly without completing workflows, or execute workflows yet miss the final value. Under very large tool spaces (OraA-10, Full), OSR collapses toward zero and execution-based rankings flatten. Table 18 further shows that most failures occur at tool selection, with $\mathrm{ACC} \mid \mathrm{selection} \gg \mathrm{ACC}$, indicating that the main bottleneck lies in tool selection and workflow orchestration rather than single-step tool invocation.

Interface structure shapes agent behavior. Structured interfaces (e.g., package routing or retrieval-based exposure) stabilize tool selection and improve orchestration success by imposing external structure on decisions. However, these gains mainly reduce decision burden rather than strengthen intrinsic skills, as recovery and replanning improve only marginally (Appendix C).

4.5 Analysis under controlled interventions

Tool explosion enlarges the agent's action space. Rather than making causal claims, we conduct two **quasi-experimental** interventions to examine how (i) tool-space size and (ii) context length **relate to** agent behaviour.

Intervention 1 tool-space size We fix the prompt context and vary only the number of available tools. We evaluate $k \in \{0, 3, 6, 10\}$ distractors under the G+R+kN configuration while keeping tasks and prompts identical. Qwen is reported as a representative model because it avoids ceiling and floor effects and shows clear variation in ACC, TSR and OSR across k (Table 20).

As k increases, TSR and OSR monotonically decrease and OSR is essentially zero for $k \geq 6$. This pattern is consistent with enlarged tool spaces making tool selection and orchestration more difficult, although changing k may also influence other prompt properties.

Intervention 2 context length We then fix the tool set (OraA-3) and vary only prompt length using short and long versions of the same tasks on a 30-task subset. Both Claude and Qwen are evaluated. Despite different absolute scores, both models show the same qualitative trend (Table 19).

Longer contexts are consistently associated with lower ACC, TSR and PPS, and OSR drops to zero for both models, indicating degraded tool selection and orchestration even when the action space is fixed.

Further quantitative and conditional error analyses are provided in Appendix E.1.

5 Conclusion

We presented REALFIN, a benchmark for evaluating LLM agents in realistic financial workflows. It emphasizes agent skills rather than answer matching, including tool selection, multi-step orchestration, parameter planning, temporal grounding, and error recovery. Each task is paired with an executable reference program and interacts with real financial APIs, enabling process-level evaluation.

Experiments under controlled tool exposure demonstrate tool explosion effects: adding related tools while fixing task semantics consistently degrades accuracy and orchestration success. The main performance drop concentrates on tool selection and cascades to execution, and answer-based and process-based metrics yield different model rankings. With workflow guidance, accuracy increases markedly, showing that tasks are solvable and the difficulty mainly lies in agent skills.

REALFIN serves as a diagnostic benchmark, offering an executable and process-aware testbed for studying agent behavior in dense tool ecosystems and time-sensitive environments.

Limitations

Although REALFIN is designed to characterize tool use and workflow execution in real financial environments, the benchmark has several limitations. First, its construction is constrained by the financial tool ecosystem and market coverage currently supported, and it does not exhaust all financial subdomains or data-source types. Model behavior under different market institutions or tool systems may therefore exhibit additional challenges. Second, despite incorporating temporal sensitivity and real APIs, our tasks and workflows are still abstracted and parameterized, and REALFIN covers only a representative subspace of the fully open-world setting. Third, the quality-control and evaluation protocols aim to separate system noise from model behavior, but inevitably depend on implementation details and runtime environments, motivating future studies on cross-platform and cross–data-source consistency.

The majority of tasks in the REALFIN dataset are constructed in Chinese, with a smaller portion consisting of English tasks (a few dozen). This language distribution will be clarified in the revised manuscript. We acknowledge this language imbalance and plan to expand the dataset to include more tasks in multiple languages as part of future work.

Overall, these limitations do not undermine the role of REALFIN in studying agent skills in realistic tool environments. Instead, they point to future extensions, including broader market coverage, richer tool types, increased task openness, and transferable evaluation across ecosystems.

Ethics Statement

REALFIN is an evaluation benchmark rather than a deployed decision-making system. It is designed for scientific study of agent behaviour and does not provide investment advice or trading signals. Nevertheless, evaluating agents in financial settings raises several ethical considerations.

First, we clarify the provenance of the data used in REALFIN. All tasks are generated from publicly accessible financial data sources and public API interfaces. The benchmark does not contain personal data, non-public market information, or proprietary client records. No human subjects were involved, and no manual annotations of personal content were collected.

Second, we consider the risk of potential misuse. Although REALFIN focuses on evaluation rather than training or deployment, stronger agentic capabilities in finance could in principle be used for harmful purposes, including market manipulation, automated speculation, or amplified misinformation about financial assets. To mitigate these risks, the benchmark provides no trading APIs and does not include tasks that directly optimise or recommend financial actions. The benchmark is intended for controlled research environments rather than end-user decision automation.

Third, we reflect on possible systemic effects. Real financial APIs and time-varying data reflect real-world conditions, but agents evaluated on REALFIN should not be assumed to be safe or reliable for high-stakes use. We discourage interpreting benchmark scores as certification for deployment in financial decision-making pipelines; rigorous domain-specific oversight remains necessary.

Fourth, we discuss fairness and accessibility. The benchmark reflects financial markets and APIs that are available in our environment and may not represent global financial practices equally. We release detailed documentation and open-source code to support transparency and reproducibility, while acknowledging that API access may depend on regional availability.

Overall, we emphasize that REALFIN aims to support responsible research into agent evaluation by making processes observable and failure modes explicit, rather than by enabling autonomous financial decision systems. Researchers using this benchmark should consider the broader societal impact of agentic AI in finance and avoid applications that could lead to individual or systemic harm.

References

Zhiyu Chen, Wenhu Chen, Charese Smiley, Sameena Shah, Iana Borova, Dylan Langdon, Reema Moussa, Matt Beane, Ting-Hao Kenneth Huang, Bryan R. Routledge, and William Yang Wang. 2021. Finqa: A dataset of numerical reasoning over financial data. In Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021, pages 3697–3711. Association for Computational Linguistics.

Zhiyu Chen, Shiyang Li, Charese Smiley, Zhiqiang Ma, Sameena Shah, and William Yang Wang. 2022. Convfinqa: Exploring the chain of numerical reasoning in conversational finance question answering. In Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022, pages 6279–6292. Association for Computational Linguistics.

Yang Deng, Wenqiang Lei, Wenxuan Zhang, Wai Lam, and Tat-Seng Chua. 2022. PACIFIC: towards proactive conversational question answering over tabular and textual data in finance. In Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022, pages 6970–6984. Association for Computational Linguistics.

Xin Guo, Haotian Xia, Zhaowei Liu, Hanyang Cao, Zhi Yang, Zhiqiang Liu, Sizhe Wang, Jinyi Niu, Chuqi Wang, Yanhui Wang, Xiaolong Liang, Xiaoming Huang, Bing Zhu, Zhongyu Wei, Yun Chen, Weining Shen, and Liwen Zhang. 2025. Fineval: A chinese financial domain knowledge evaluation benchmark for large language models. In Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL 2025 - Volume 1: Long Papers, Albuquerque, New Mexico, USA, April 29 - May 4, 2025, pages 6258–6292. Association for Computational Linguistics.

Albert King. 2022. Akshare. https://github.com/akfamily/akshare.

Haohang Li, Yupeng Cao, Yangyang Yu, Shashidhar Reddy Javaji, Zhiyang Deng, Yueru He, Yuechen Jiang, Zining Zhu, K. P. Subbalakshmi, Jimin Huang, Lingfei Qian, Xueqing Peng, Jordan W. Suchow, and Qianqian Xie. 2025. INVESTORBENCH: A benchmark for financial decision-making tasks with llm-based agent. In Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2025, Vienna, Austria, July 27 - August 1, 2025, pages 2509–2525. Association for Computational Linguistics.

Minghao Li, Yingxiu Zhao, Bowen Yu, Feifan Song, Hangyu Li, Haiyang Yu, Zhoujun Li, Fei Huang, and Yongbin Li. 2023. Api-bank: A comprehensive benchmark for tool-augmented llms. In Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023, pages 3102–3116. Association for Computational Linguistics.

Xiao Liu, Hao Yu, Hanchen Zhang, Yifan Xu, Xuanyu Lei, Hanyu Lai, Yu Gu, Hangliang Ding, Kaiwen Men, Kejuan Yang, Shudan Zhang, Xiang Deng, Aohan Zeng, Zhengxiao Du, Chenhui Zhang, Sheng Shen, Tianjun Zhang, Yu Su, Huan Sun, and 3 others. 2024. Agentbench: Evaluating llms as agents. In The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024. OpenReview.net.

Grégoire Mialon, Clémentine Fourrier, Thomas Wolf, Yann LeCun, and Thomas Scialom. 2024. GAIA: a benchmark for general AI assistants. In The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024. OpenReview.net.

Shishir G. Patil, Tianjun Zhang, Xin Wang, and Joseph E. Gonzalez. 2024. Gorilla: Large language model connected with massive apis. In Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024.

Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong, Xiangru Tang, Bill Qian, Sihan Zhao, Lauren Hong, Runchu Tian, Ruobing Xie, Jie Zhou, Mark Gerstein, Dahai Li, Zhiyuan Liu, and Maosong Sun. 2024a. Toolllm: Facilitating large language models to master 16000+ real-world apis. In The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024. OpenReview.net.

Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong, Xiangru Tang, Bill Qian, Sihan Zhao, Lauren Hong, Runchu Tian, Ruobing Xie, Jie Zhou, Mark Gerstein, Dahai Li, Zhiyuan Liu, and Maosong Sun. 2024b. Toolllm: Facilitating large language models to master 16000+ real-world apis. In The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024. OpenReview.net.

Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. Reflexion: language agents with verbal reinforcement learning. In Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023.

Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. 2024. Voyager: An open-ended embodied agent with large language models. Trans. Mach. Learn. Res., 2024.

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R. Narasimhan, and Yuan Cao. 2023. React: Synergizing reasoning and acting in language models. In The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023. OpenReview.net.

Shuyan Zhou, Frank F. Xu, Hao Zhu, Xuhui Zhou, Robert Lo, Abishek Sridhar, Xianyi Cheng, Tianyue Ou, Yonatan Bisk, Daniel Fried, Uri Alon, and Graham Neubig. 2024. Webarena: A realistic web environment for building autonomous agents. In The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024. OpenReview.net.

Fengbin Zhu, Wenqiang Lei, Youcheng Huang, Chao Wang, Shuo Zhang, Jiancheng Lv, Fuli Feng, and Tat-Seng Chua. 2021. TAT-QA: A question answering benchmark on a hybrid of tabular and textual content in finance. In Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021, pages 3277–3287. Association for Computational Linguistics.

Appendix

A Dataset Statistics

Type distribution. Table 5 summarizes the distribution of task types in REALFIN.

Takeaways.

- 1,524 valid instances, no broken records.
- Aggregate questions (T3) constitute nearly half of the corpus.

Type	Count	Proportion
T3_Aggregate	677	44.4%
T1_Snapshot	412	27.0%
T4_Fundamentals	328	21.5%
T2_Compare	107	7.0%

Table 5: Distribution of task types in the 1,524-instance dataset.

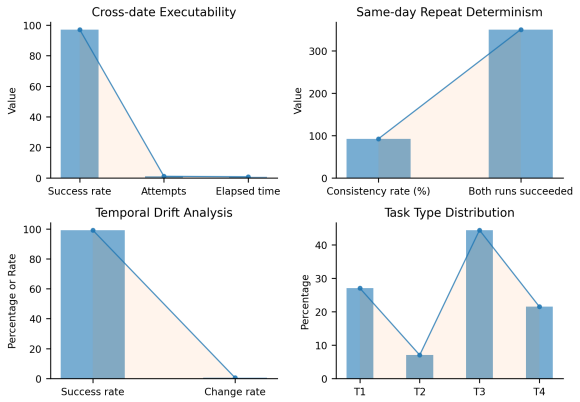

Figure 8: **Code analysis experiments.** We evaluate executability across dates, determinism across repeated runs, and temporal drift of time-sensitive tasks. High executability with non-trivial drift confirms that REALFIN captures realistic, time-varying financial environments.

Execution depth.

Tool usage statistics. **Takeaways.**

- 567 single-tool vs. 957 multi-tool instances (37.2% vs. 62.8%).
- 433 instances (28.4%) require actual chaining (depth ≥ 2).

Time Sensitivity	Count	Proportion
Fixed	794	52.1%
Relative	629	41.3%
Realtime	101	6.6%

Table 6: Time-sensitivity labels for dataset instances.

Time sensitivity. **Type $\times$ Time Sensitivity.** **Depth $\times$ Time Sensitivity.** **Takeaways.**

- Fixed-date queries dominate; relative-time queries are substantial.
- Realtime queries are mostly shallow and largely in T1.
- Deeper chains skew toward relative-time computation.

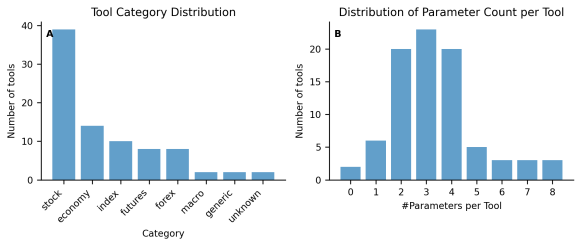

Figure 9: **Tool ecosystem statistics.** Left: distribution of tool categories. Right: distribution of parameter counts per tool. The benchmark exhibits heterogeneous and partially redundant tools, operationalizing the challenge of tool explosion.

Type	Fixed	Realtime	Relative
T1_Snapshot	315	97	0
T2_Compare	102	4	1
T3_Aggregate	49	0	628
T4_Fundamentals	328	0	0

Table 7: Joint distribution of task type and time sensitivity.

Workflow complexity signals. **Takeaways.**

- Control flow is pervasive despite shallow average depth.
- Multi-tool usage is common; long chains are less frequent.

Extended Dataset Statistics.

B Tool Library Statistics

Domain distribution. **Takeaways.**

- Stock-related tools dominate the library.
- Long tail of mixed-domain tools reflects cross-asset workflows.

Task types of tools. **Takeaways.**

- The library is primarily retrieval-oriented; computational tools are secondary.

Dependency graph properties. **Graph statistics.**

- Leaves: 46
- Internal nodes: 39
- Longest chain length: 4

Hub tools (fan-in).

Signal	Count	Proportion
Control flow in code	1436	94.2%
Multiple tools (>1)	957	62.8%
Chained tools (depth ≥ 2)	433	28.4%

Table 8: Indicators of workflow complexity.

Statistic	**Value**
Number of questions	1,524
Executable golden codes	1,524
Tool functions	85
Markets covered	A, HK, US, FX, Macro
Task types (T1/T2/T3/T4)	27% / 18% / 44% / 11%
Avg. execution depth	2.12
Avg. tool calls per instance	3.78
Multi-tool questions	73%
Temporal regime: Fixed	28.3%
Temporal regime: Relative	52.1%
Temporal regime: Realtime	19.6%
Golden program success rate	97.15%
Avg. attempts per program	1.06
Repeat-run consistency	87.64%
Temporal drift proportion	0.55%

Table 9: **Statistics of REALFIN.** All questions are paired with executable golden programs. Executable analysis is based on automatically re-running all golden programs across dates.

- `get_code_from_name`: 18
- `search_kg`: 4
- `get_index_history`: 3

Fan-out.

- `calculate_gold_forex_change`: 2
- `get_a_stock_daily_price`: 2

Takeaways.

- The dependency graph is shallow with a few hub tools.

Error models. Takeaways.

- 97.6% of tools expose structured error models.
- Parameter and internal errors dominate.

Real APIs vs. virtual tools. Takeaways.

- Rough balance between wrapper tools and virtual/algorithmic tools.
- Akshare is the dominant external data provider.

Domain	Count	Proportion
stock	37	43.5%
index	10	11.8%
economy	10	11.8%
futures	7	8.2%
forex	7	8.2%
generic	4	4.8%
mixed domains	8	–

Table 10: Domain coverage of tools. Mixed-domain rows aggregate tools annotated with multiple domains.

Metric	Value
Total tools	133
External-API tools	65
Virtual / algorithmic tools	68
Unique Akshare endpoints	66

Table 11: Real vs. virtual tool implementation statistics.

C Full metric definitions

This appendix provides full mathematical definitions of the metrics introduced in Section §4. We distinguish between outcome accuracy and process-based agent-skill metrics.

C.1 Notation

Let the dataset be $\mathcal{D} = \{1, \ldots, N\}$. For each instance $i \in \mathcal{D}$:

- y_i : reference answer
- $\hat{y}_i$: model's final answer
- $T_i^* = \langle t_{i,1}^*, \ldots, t_{i,m_i}^* \rangle$: gold tool sequence
- $\hat{T}_i = \langle \hat{t}_{i,1}, \ldots, \hat{t}_{i,k_i} \rangle$: model tool sequence
- $\mathcal{T}_i^*$: set of gold tools (ignoring order)
- $\hat{\mathcal{T}}_i$: set of tools used by the model
- d_i^* : gold workflow depth $= |T_i^*|$
- $\hat{d}_i$: model workflow depth $= |\hat{T}_i|$
- $r_{i,j}$: execution result of tool call $\hat{t}_{i,j}$

We write $\mathbf{1}[\cdot]$ for the indicator function.

C.2 Outcome-Level Metrics

Outcome Accuracy (Acc).

$$\text{Acc} = \frac{1}{N} \sum_{i=1}^{N} \mathbf{1}[\hat{y}_i = y_i]. \quad (1)$$

C.3 Execution and Orchestration

Execution Success Rate (ESR). Let $\mathrm{OK}(r)$ denote a non-error, non-null result.

$$\mathrm{ESR} = \frac{1}{N}\sum_{i=1}^{N}\prod_{j=1}^{k_i}\mathbf{1}[\mathrm{OK}(r_{i,j})]. \tag{2}$$

Full Orchestration Success Rate (OSR). A sample is successful if all skills succeed:

$$\mathrm{OSR} = \frac{1}{N}\sum_{i=1}^{N}\mathbf{1}[\mathrm{TSR}_i = \mathrm{OA}_i = \mathrm{PPS}_i = \mathrm{EGA}_i = 1].$$

C.4 Tool Selection

Let

$$\mathcal{T}_i^* = \{t_{i,1}^*, \ldots, t_{i,m_i}^*\}, \quad \hat{\mathcal{T}}_i = \{\hat{t}_{i,1}, \ldots, \hat{t}_{i,k_i}\}.$$

Tool Selection Recall (TSR).

$$\mathrm{TSR} = \frac{1}{N}\sum_{i=1}^{N}\frac{|\mathcal{T}_i^* \cap \hat{\mathcal{T}}_i|}{|\mathcal{T}_i^*|}. \tag{3}$$

C.5 Ordering and Workflow Planning

Ordering Accuracy (OA). Let $\mathrm{Subseq}(A, B) = 1$ iff A is a subsequence of B.

$$\mathrm{OA} = \frac{1}{N}\sum_{i=1}^{N}\mathbf{1}\left[\mathrm{Subseq}(T_i^*, \hat{T}_i)\right]. \tag{4}$$

This measures whether the model preserves the logical order of tools.

C.6 Parameter Planning

Define validity of a tool call result:

$$\mathrm{Good}(r) = \mathbf{1}[r \neq \mathrm{NULL},\ r \neq \mathrm{ERROR}].$$

Parameter Planning Success (per instance). For each gold-required tool class c, consider the model's first call to a tool in c and its result r.

$$\mathrm{PPS}_i = \prod_{c \in \mathcal{C}_i^*}\mathrm{Good}(r_{i,c}), \tag{5}$$

where $\mathcal{C}_i^*$ is the set of required tool classes.

Dataset-level PPS.

$$\mathrm{PPS} = \frac{1}{N}\sum_{i=1}^{N}\mathrm{PPS}_i.$$

C.7 Entity Grounding

Let E_i^* be entity set from question + gold program, and $\hat{E}_i$ the entities implied by tool calls.

$$\mathrm{EGA} = \frac{1}{N}\sum_{i=1}^{N}\mathbf{1}\left[\exists e \in \hat{E}_i \text{ matching some } e^* \in E_i^*\right].$$

C.8 Recovery Success Rate (RSS)

Let $\hat{S}_i$ be the set of steps taken by the agent for task i, and S_i^* the reference steps in the gold program.

$$\mathrm{RSS} = \frac{1}{N}\sum_{i=1}^{N}\mathbf{1}\left[\exists s \in \hat{S}_i \text{ such that } s \in S_i^*\right].$$

RSS evaluates the success of the agent in recovering from errors by comparing the steps taken by the agent with those in the reference workflow.

C.9 Temporal Grounding Accuracy (TGA)

Let T_i^* be the set of temporal references from the question + gold program, and $\hat{T}_i$ the temporal references inferred by the agent's tool calls.

$$\mathrm{TGA} = \frac{1}{N}\sum_{i=1}^{N}\mathbf{1}\left[\exists t \in \hat{T}_i \text{ matching some } t^* \in T_i^*\right].$$

TGA measures the agent's ability to correctly interpret and apply temporal information, such as dates, time periods, or relative times, within the context of the task.

C.10 Detailed Performance Metrics

Table 12 presents the comprehensive performance metrics across different retrieval strategies and language models on our benchmark dataset.

C.11 Case Study

C.11.1 Case Study I: Do functionally redundant tools unlock alternative execution paths?

Claude-Sonnet-4.5 achieves higher accuracy under OraA-3 tool exposure than under the oracle Oracle setting. Since OraA-3 differs from Oracle primarily by adding functionally redundant tools, we ask:

> **Do redundant tools unlock alternative executable workflows that are unavailable under the oracle minimal tool set?**

Experimental probe. We evaluate GPT-5.1 in a controlled setting where the agent is given the oracle tool set $\mathcal{G}$ augmented with redundant tools. We then search for instances satisfying:

1. the oracle tool set alone is sufficient to solve the task; and
2. the model nevertheless solves the task through a different workflow that explicitly uses redundant tools.

Across 1,524 questions we identify six such cases. Although rare, they are conceptually important: current agents seldom explore alternative workflows, yet multi-path execution does occur and can change measured accuracy under Top-k exposure.

Representative examples. Figure 10 illustrates two representative cases.

(a) Alternative executable paths. The agent sometimes constructs a workflow different from the golden program but semantically equivalent, showing that the benchmark implicitly admits multiple valid execution paths even when only one is annotated.

(b) Redundant tools as rescue routes. In other cases, the canonical workflow fails due to missing data or transient API errors, while the agent succeeds by switching to a redundant tool. Redundancy thus acts as a fallback mechanism and improves robustness, rather than serving only as distraction.

Takeaway. This case study suggests that:

- redundant tools occasionally unlock alternative valid execution paths that are not used under Oracle settings;
- current agents do not systematically search workflow space, explaining the small number of observed cases;
- when such paths are discovered, they can improve accuracy under Top-K tool exposure.

Taken together with Case Study II, this reveals a nuanced picture: moderate Top-K exposure may help by increasing the number of valid paths, whereas excessive tool exposure may hurt by increasing decision uncertainty and abstention. The key challenge is therefore not only to provide tools, but to guide agents in exploring the space of executable workflows effectively.

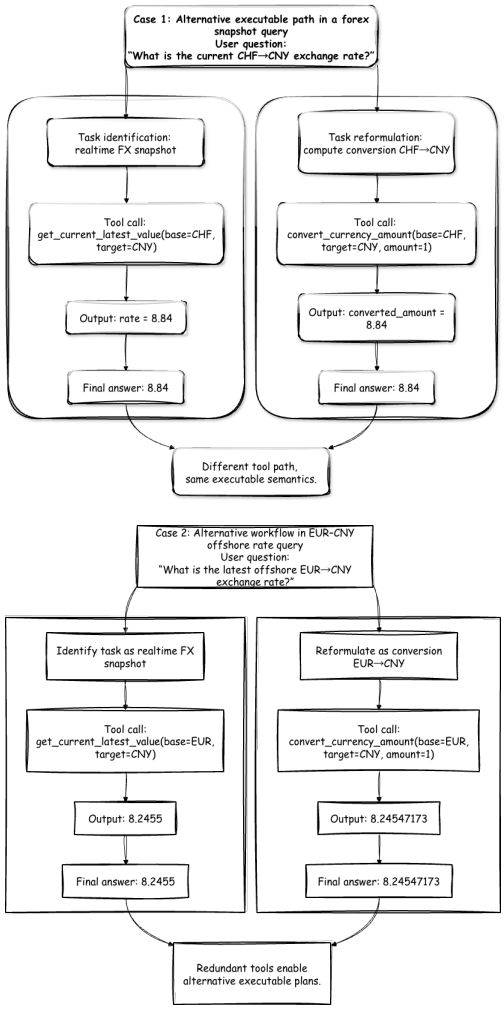

Figure 10: **Two examples of multi-path execution unlocked by redundant tools.** Top: an FX snapshot query solved through an alternative API chain; Bottom: a dividend query where redundant tools provide a fallback path after canonical-tool failure.

C.11.2 Case Study II: Does a Larger Top-K Make Agents Stronger or More Hesitant?

In the previous case study, we showed that OraA-3 exposure can outperform the oracle Oracle setting because it provides additional valid execution paths: agents sometimes reach the correct answer through workflows that differ from the reference program. We now ask a complementary question: **if increasing K creates even more potential paths, why does OraA-10 often perform worse than OraA-3?**

Experimental setup. We focus on the 394-instance paired subset. Each instance is evaluated twice with identical prompts, evaluation protocol, and models; the only difference is the size of the tool list presented to the agent ($k = 3$ vs. $k = 10$).

This paired design isolates the effect of Top-K exposure from all other confounders.

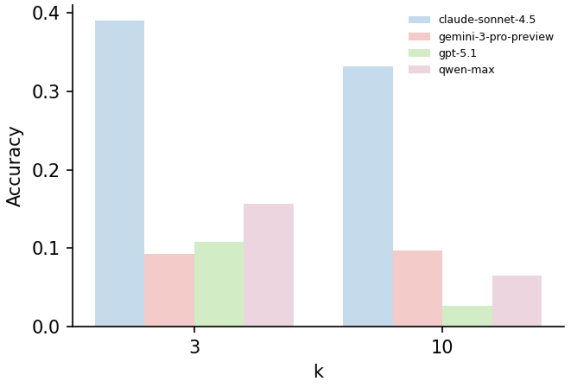

Figure 11: Accuracy comparison between $k = 3$ and $k = 10$ on the 394-item paired subset. Each bar shows mean accuracy across samples for each model.

Observation 1: Increasing K reduces accuracy. Figure 11 shows that OraA-3 consistently outperforms OraA-10 across models on the same tasks. This result is striking because OraA-10 strictly contains OraA-3: if tool availability were the primary bottleneck, accuracy should be non-decreasing in K. Instead, expanding the visible tool set degrades performance, suggesting that the dominant factor is not the absence of the right tool, but the agent's behavior under increased choice.

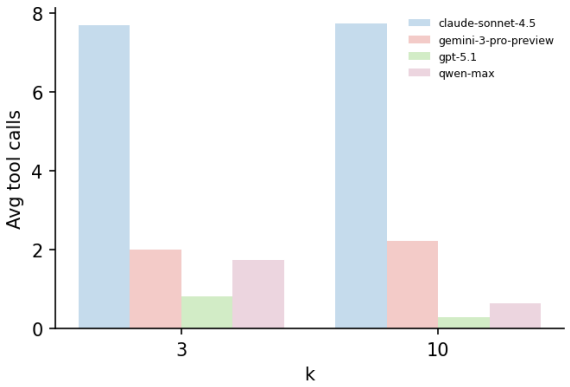

Figure 12: Change in average number of tool calls when increasing from $k = 3$ to $k = 10$. Positive values indicate more tool usage under larger k.

Observation 2: Lower accuracy is not caused by insufficient tool use. Figure 12 shows that the mean number of tool calls does not decrease when moving from $k = 3$ to $k = 10$; in many models it even increases. Therefore, OraA-10 does not fail simply because agents "give up early" or "call too few tools". Instead, errors arise during execution—mis-parameterization, tool misselection among near-duplicates, or inconsistent multi-step composition—rather than from a lack of action.

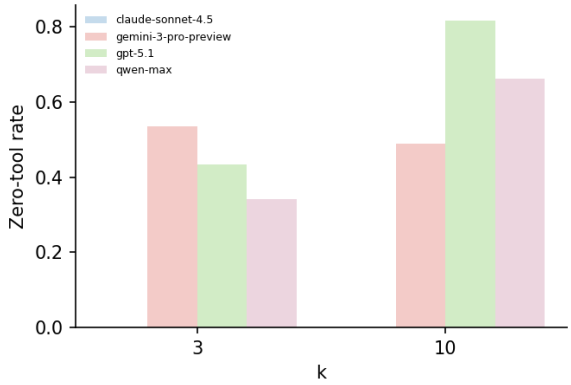

Figure 13: Zero-tool answer rate for $k = 3$ and $k = 10$. Higher values indicate a stronger tendency to avoid tool usage entirely.

Observation 3: Larger K increases hesitation and abstention. Figure 13 shows that OraA-10 settings exhibit a markedly higher zero-tool answer rate than OraA-3. Thus, enlarging the tool set does not only change which tools are used; it changes whether the agent decides to use tools at all. This indicates a shift in strategy: when faced with a larger and more redundant tool space, agents increasingly default to direct answers instead of committing to multi-step tool execution.

Interpretation. Putting these observations together, we obtain the following picture:

- Oracle → OraA-3 improves performance by introducing additional valid execution paths.
- OraA-3 → OraA-10 degrades performance by inflating decision uncertainty and abstention pressure.

In other words, tool exposure has a non-monotonic effect: a moderate number of tools increases the chance of discovering a successful workflow, but excessive tool exposure induces hesitation, increases zero-tool responses, and amplifies composition errors in deep tool chains.

Takeaway. The limiting factor for current agents is therefore not simply tool availability, but their robustness under choice overload. More tools can make agents not only more error-prone, but also more hesitant to act—highlighting tool explosion as a behavioral phenomenon rather than a purely algorithmic scaling issue.

C.11.3 Case Study III: Does coarse tool packaging eliminate the tool-explosion problem?

Top-K exposure limits the visible tool set but does not restructure it. A complementary strategy is to reorganize tools hierarchically into coarse functional packages and require the agent to first select an intent, then a package, and only then a specific tool. This raises the guiding question of this case study:

> **If tools are hierarchically packaged, does tool explosion essentially disappear?**

Hierarchical routing and packaging in practice. Figure 14 shows a representative execution trace under our three-stage routing framework: (i) coarse-grained intent classification, (ii) fine-grained intent routing, and (iii) tool selection within a constrained package. As illustrated in the figure, even after packaging, the agent must still choose among multiple highly similar tools inside the same package and reason over heterogeneous argument schemas and return formats. Thus, packaging restructures the search space, but does not collapse it into a single deterministic route.

Aggregate empirical effect. We evaluate GPT-5.1 under this packaged regime and obtain 180 correct answers out of 520 questions (approximately 30% accuracy). This confirms that hierarchical packaging does help: compared with flat tool exposure, the agent faces fewer irrelevant tools and the routing decisions visualized in Figure 14 become more organized and interpretable.

But tool explosion persists. However, the same execution trace in Figure 14 also reveals why accuracy remains far from satisfactory. Most errors arise from:

- confusion among multiple tools within the same functional family,
- unstable parameter grounding across stages of the workflow,
- failure to normalize heterogeneous output schemas across tools,
- early abandonment of tool use when latency or API noise is encountered.

The branching structure shown in Figure 14 indicates that even after packaging, the agent still faces a large workflow search space rather than a single correct pipeline.

Takeaway. Coarse tool packaging mitigates surface-level search overload but does not eliminate tool explosion. The challenge is structural: the agent must coordinate multi-step workflows across noisy, partially redundant tools. Therefore, tool explosion is not merely a catalog-design artifact; it persists even under hierarchical routing, as visually evidenced in Figure 14.

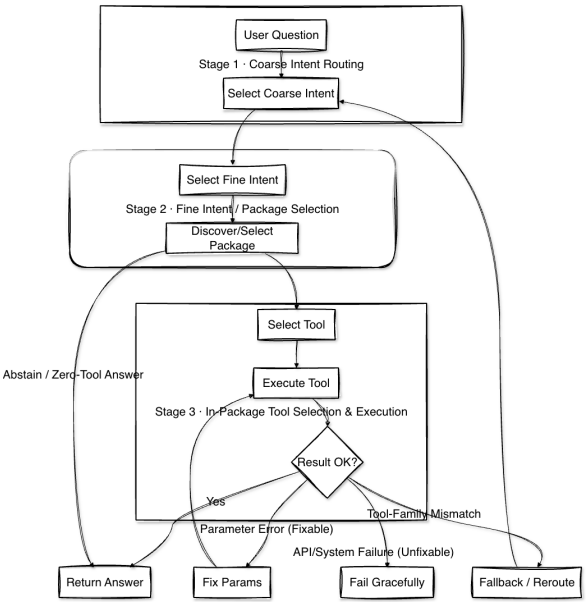

Figure 14: **Hierarchical intent routing and packaged tool selection.** The execution trace demonstrates that even when tools are grouped into packages, the agent must still disambiguate between multiple similar tools and coordinate multi-step workflows, indicating that tool explosion persists in packaged settings.

C.11.4 Case Study IV: Why does accuracy collapse completely beyond execution depth 1?

Across all three subsets (128 / 249 / 394) and all exposure conditions (Oracle, Top-3, Top-10, full), we consistently observe the same phenomenon: tasks with execution depth 1 maintain non-zero accuracy, whereas accuracy rapidly declines as depth increases, and collapses to zero at depth 5. This is not an artifact of specific models or tool exposure settings; rather, it suggests a structural failure mode of current LLM-based agents.

Figure 15 illustrates this monotonic depth-induced degradation: depth 1 tasks are solvable at modest rates, but depth 5 tasks are solved by none of the evaluated agents under any tool-exposure configuration.

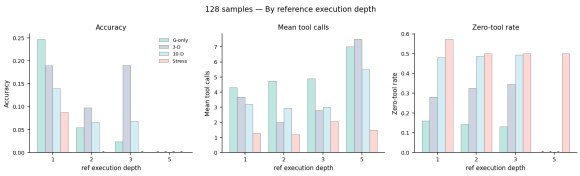

Figure 15: Zero-tool answer rate for $k = 3$ and $k = 10$. Higher values indicate a stronger tendency to avoid tool usage entirely.

Observation 1: Failure is depth-driven, not exposure-driven. Notably, the collapse at depth 5 occurs even in the Oracle (oracle tool set) condition where tool selection burden is eliminated. The same pattern appears under Top-3, Top-10, and full exposure. This indicates that the source of difficulty is not merely tool retrieval or selection; instead, execution depth itself is the dominant factor.

Observation 2: Depth increases tool-use behavior but not accuracy. Across settings, the mean number of tool calls increases with execution depth, yet accuracy simultaneously decreases. This dissociation shows that additional tool use does not translate into improved correctness: agents attempt longer workflows, but cannot maintain coherent execution over multiple steps.

Observation 3: Divergent failure behaviors at high depth. At depth 5, we observe two distinct behavioral responses depending on tool exposure: (i) under full (full-tool exposure), the zero-tool rate rises to around 50%, suggesting deliberate abstention in the face of overwhelming search space; (ii) under Top-K exposure, zero-tool rates drop to nearly zero while tool calls surge, indicating thrashing behavior where the agent repeatedly invokes tools without converging on the correct workflow. These contrastive behaviors point to a common bottleneck—long-horizon planning—modulated by exposure-induced uncertainty.

Why does depth hurt so strongly? The empirical evidence supports a multi-step consistency explanation:

- **state drift**: intermediate tool outputs must be remembered and reused, but partial results are overwritten or inconsistently referenced across steps;
- **temporal misalignment**: trading-day offsets and date normalization must remain consistent across calls, and small inconsistencies compound over depth;
- **parameter propagation errors**: arguments derived from previous tool outputs are misparsed, reformulated, or silently dropped in later calls;
- **error accumulation**: small inaccuracies at early steps propagate forward and dominate the final answer;
- **goal myopia**: agents locally repair tool errors without maintaining a global workflow plan, leading to the high-call "thrashing" observed at depth 5.

Takeaway. Execution depth is the single strongest predictor of failure in REALFIN. Even when tool selection is simplified (Oracle) or tool space is constrained (Top-K), agents fail beyond one or two steps. This suggests that the key challenge in financial agent evaluation is not isolated tool use, but maintaining globally consistent, temporally aligned, multi-step workflows over long horizons. Long-chain orchestration—rather than single-step competence—is therefore the dominant bottleneck for current LLM agents.

C.11.5 Case Study V: Representative failure types

We illustrate how tool explosion changes the error landscape. Failures are not random: they cluster into a small number of systematic types that align with our process metrics.

D Alignment with Real-World Financial SOPs

We present three representative cases to illustrate that RealFin golden workflows closely mirror real-world financial standard operating procedures (SOPs). The cases cover (1) equity quarterly return computation, (2) cross-sectional liquidity aggregation, and (3) precious-metals benchmark price comparison.

E Supplement Experiment Results

Table 17: Model-averaged bootstrap results

Setting	ACC	Exec	Oracle	Prec	Rec
Oracle	0.21 ± 0.03	0.67 ± 0.03	0.76 ± 0.03	1.00 ± 0.00	0.68 ± 0.04
OraA-3	0.09 ± 0.02	0.34 ± 0.02	0.39 ± 0.02	0.98 ± 0.01	0.55 ± 0.04
OraA-10	0.05 ± 0.01	0.24 ± 0.02	0.25 ± 0.02	0.91 ± 0.02	0.35 ± 0.04
Pkg	0.28 ± 0.04	0.34 ± 0.01	0.00 ± 0.00	0.12 ± 0.03	0.13 ± 0.03
BM25*	0.13 ± 0.03	0.60 ± 0.04	0.64 ± 0.04	0.99 ± 0.01	0.43 ± 0.04
Full	0.12 ± 0.02	0.49 ± 0.03	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00

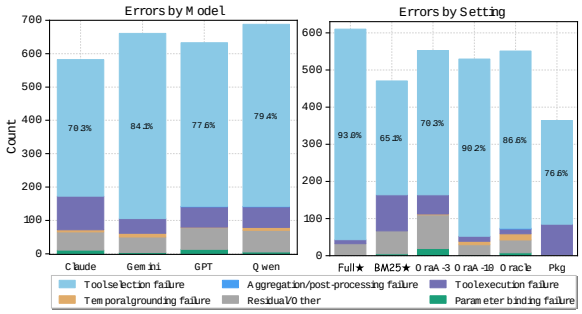

Figure 16: Error distribution.

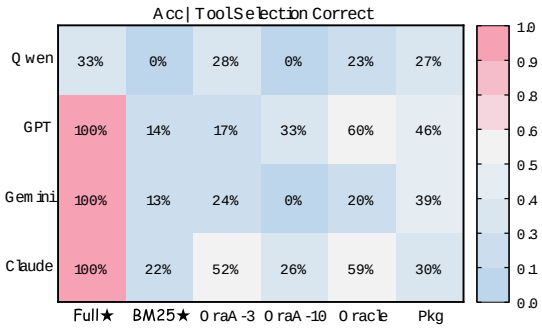

Figure 17: **Conditional accuracy.** Accuracy improves sharply once correct tools are selected.

E.1 Error analysis

E.1.1 Quantitative error breakdown

Figure 16 shows that the majority of errors are caused by *tool selection* (78.1%), with execution errors (10.5%) and residual errors (9.0%) forming most of the remainder. Parameter, temporal, and aggregation errors jointly account for less than 3%. Thus the primary bottleneck in REALFIN is deciding which tools to use, not arithmetic or aggregation.

This pattern is consistent across models (70–84% selection error) and exposure regimes: Full and OraA-10 exhibit the highest selection error rates, whereas OraA-3 and BM25 regimes contain more execution mistakes.

E.1.2 Conditional performance given correct skills

We compute accuracy (i) conditioned on correct tool selection and (ii) conditioned on both correct selection and reference-consistent orchestration (Figure 17).

Main observation. Across all models, ACC | Selection is 3–6× higher than raw accuracy, whereas ACC | Orchestration ≈ ACC | Selection. Hence, once the right tools are chosen, ordering and arithmetic rarely fail; the dominant source of failure is tool selection under tool explosion.

F Annotation Guidelines

F.1 Human evaluation procedure

This appendix describes the human evaluation used for RealFin. The goal of the questionnaire is to assess the quality of tasks and reference workflows in the dataset. Annotators are instructed to read the guidelines carefully and base their judgments only on the information provided for each item.

For every item, annotators are given:

- the question description
- the golden answer
- the reference program implementing the reference workflow
- tool schema documentation, including tool names, functionality, and input and output fields

Annotators are asked not to rely on any external resources and not to execute the code. All judgments must be made purely from the provided information.

F.2 Rating dimensions

Each item is scored along four dimensions:

1. **Question reasonableness:** clarity, lack of ambiguity, and plausibility within real financial analysis scenarios
2. **Answer correctness:** consistency of the golden answer with the question and whether the value or conclusion is correct or essentially correct
3. **Reference workflow correctness:** without running the code, whether the tools are appropriate, parameters are reasonable, and the sequence of steps is conceptually able to produce the correct answer
4. **Tool schema usability:** clarity and sufficiency of the tool documentation, including input and output meanings, to reasonably support solving the task

F.3 Scoring scheme

All four dimensions use a three-level scale:

2 fully correct or reasonable 1 minor issues but usable 0 clearly incorrect or unusable

Annotators may optionally add short comments explaining their scores. If a judgment truly cannot be made due to insufficient information, annotators may select unable to judge and briefly explain the reason. Annotators complete the questionnaire independently.

F.4 Example annotation item

Table 12: Detailed performance comparison across retrieval strategies and models. **Oracle**: Ground truth tool selection. **BM25**: BM25-based retrieval (top-20). **OraA-k**: Oracle with top-k augmentation. **Full**: Full tool library. Metrics include Accuracy (ACC), Exact Success Rate (ESR), Tool Selection Rate (TSR), Overall Accuracy (OA), Exact Goal Achievement (EGA), Task Goal Achievement (TGA), Partial Pass Score (PPS), Overall Success Rate (OSR), and Retrieval Success Score (RSS).

Strategy	Model	ACC	ESR	TSR	OA	EGA	TGA	PPS	OSR	RSS
Oracle: Ground Truth Tool Selection										
Oracle	Doubao-1.6	0.367	0.703	0.281	0.266	0.344	0.545	0.195	0.047	0.091
Oracle	Kimi-K2	0.305	0.750	0.227	0.203	0.398	0.636	0.141	0.008	0.083
Oracle	Claude-4.5	0.273	0.742	0.227	0.219	0.500	0.568	0.133	0.031	0.099
Oracle	DeepSeek-3.2	0.320	0.734	0.211	0.203	0.539	0.318	0.133	0.055	0.083
Oracle	Gemini-3-Pro	0.148	0.750	0.211	0.203	0.445	0.114	0.125	0.023	0.091
Oracle	GPT-5.1	0.109	0.883	0.133	0.133	0.273	0.318	0.078	0.023	0.058
Oracle	Qwen-Max	0.164	0.797	0.227	0.211	0.414	0.523	0.141	0.016	0.091
BM25: BM25-based Retrieval (Top-20)										
BM25	Doubao-1.6	0.117	0.141	0.648	0.563	0.336	–	0.234	0.078	0.374
BM25	Claude-4.5	0.250	0.203	0.719	0.602	0.586	0.000	0.297	0.141	0.366
BM25	DeepSeek-3.2	0.094	0.172	0.641	0.578	0.586	–	0.266	0.180	0.358
BM25	Gemini-3-Pro	0.100	0.719	0.227	0.195	0.195	0.000	0.078	0.055	0.138
BM25	GPT-5.1	0.064	0.633	0.375	0.352	0.289	0.000	0.117	0.047	0.260
BM25	Kimi-K2	0.047	0.250	0.516	0.484	0.453	–	0.211	0.133	0.309
BM25	Qwen-Max	0.016	0.734	0.211	0.188	0.125	0.000	0.063	0.023	0.154
OraA-k: Oracle with Top-k Augmentation										
OraA-10	Doubao-1.6	0.000	0.000	0.250	0.250	0.690	0.390	0.000	0.000	0.264
OraA-10	Kimi-K2	0.000	0.141	0.164	0.148	0.586	0.082	0.000	0.000	0.157
OraA-10	Claude-4.5	0.266	0.594	0.227	0.211	0.445	0.636	0.133	0.016	0.098
OraA-10	DeepSeek-3.2	0.008	0.000	0.309	0.285	0.699	0.286	0.000	0.000	0.302
OraA-10	Gemini-3-Pro	0.070	0.758	0.063	0.063	0.391	0.114	0.016	0.016	0.049
OraA-10	GPT-5.1	0.043	0.884	0.130	0.130	0.130	0.036	0.029	0.000	0.111
OraA-10	Qwen-Max	0.047	0.711	0.086	0.078	0.078	0.136	0.031	0.000	0.057
OraA-3	Doubao-1.6	0.000	0.000	0.305	0.297	0.711	0.659	0.000	0.000	0.314
OraA-3	Kimi-K2	0.016	0.055	0.172	0.164	0.656	0.500	0.000	0.000	0.174
OraA-3	Claude-4.5	0.367	0.820	0.656	0.609	0.438	0.636	0.484	0.141	0.179
OraA-3	DeepSeek-3.2	0.039	0.000	0.367	0.344	0.719	0.341	0.000	0.000	0.364
OraA-3	Gemini-3-Pro	0.094	0.797	0.203	0.188	0.211	0.159	0.094	0.008	0.114
OraA-3	GPT-5.1	0.086	0.836	0.500	0.438	0.258	0.477	0.352	0.117	0.154
OraA-3	Qwen-Max	0.141	0.742	0.484	0.430	0.336	0.477	0.352	0.109	0.138
OraA-0	Claude-4.5	0.000	0.933	0.062	0.062	0.554	0.000	0.000	0.000	0.000
OraA-0	Qwen-Max	0.102	0.836	0.188	0.164	0.281	0.318	0.109	0.008	0.000
OraA-0	Doubao-1.6	0.422	1.000	0.258	0.258	0.563	0.000	0.148	0.117	0.116
OraA-0	Claude-4.5	0.234	1.000	0.266	0.258	0.594	0.000	0.180	0.172	0.083
OraA-0	DeepSeek-3.2	0.195	1.000	0.234	0.234	0.578	0.000	0.141	0.133	0.099
OraA-0	Gemini-3-Pro	0.383	1.000	0.273	0.273	0.641	0.000	0.172	0.125	0.107
OraA-0	Kimi-K2	0.141	1.000	0.133	0.133	0.398	0.000	0.047	0.023	0.091
OraA-0	GPT-5.1	0.305	1.000	0.242	0.242	0.578	0.000	0.180	0.156	0.066
OraA-0	Qwen-Max	0.234	1.000	0.258	0.250	0.555	0.000	0.148	0.133	0.107
OraA-6	Claude-4.5	0.078	0.750	0.117	0.117	0.641	0.364	0.031	0.008	0.008
OraA-6	Qwen-Max	0.078	0.805	0.109	0.102	0.273	0.273	0.047	0.000	0.000
Full: Complete Tool Library										
Full	Doubao-1.6	0.336	0.758	0.227	0.211	0.164	0.614	0.164	0.008	0.066
Full	Claude-4.5	0.094	0.664	0.086	0.086	0.445	0.159	0.008	0.000	0.085
Full	DeepSeek-3.2	0.305	0.711	0.195	0.188	0.391	0.273	0.086	0.000	0.116
Full	Gemini-3-Pro	0.086	0.688	0.086	0.086	0.211	0.000	0.008	0.000	0.085
Full	GPT-5.1	0.016	0.781	0.094	0.086	0.063	0.023	0.008	0.000	0.094
Full	Kimi-K2	0.102	0.766	0.148	0.141	0.305	0.455	0.078	0.008	0.074
Full	Qwen-Max	0.031	0.500	0.086	0.070	0.109	0.045	0.023	0.000	0.067

Type	Representative example
Tool selection failure	*Question:* latest Shanghai silver benchmark price at noon. *Gold:* 17140.0. Oracle: calls `get_sge_benchmark_price` and answers correctly. 10-D / full: calls no tool and outputs free-form numbers.
Temporal grounding failure	*Question:* highest price of Suolian Co. 7 trading days ago. *Gold:* 33.56. Correct tool called, but indexes wrong trading day (returns 33.00).
Execution / answer grounding failure	*Question:* latest closing price of Yilite. *Gold:* 14.06. Correct tools called, but model ignores the "close" field and outputs unrelated values.

Table 13: **Representative failure types.** Tool explosion shifts error structure from computation to (i) tool selection, (ii) temporal grounding, (iii) answer grounding.

Table 14: Equity quarterly return SOP vs. RealFin workflow (Tencent Holdings, Q4 2023).

Analyst SOP	RealFin golden workflow
Identify target security	`get_code("Tencent")`.
Determine quarter start and end dates	Use 2023-09-30 and 2023-12-31.
Retrieve adjusted closing prices	`get_price(close)`.
Ensure price comparability	`adjust = qfq`.
Compute period return	$\frac{P_{\text{end}}}{P_{\text{start}}} - 1$.
Format output	Percentage form, two decimals.

Table 15: Cross-sectional liquidity SOP vs. RealFin workflow (U.S. semiconductor "big three").

Analyst SOP	RealFin golden workflow
Define peer group	`search_kg("U.S. semiconductor big three")`.
Obtain constituent tickers	`get_code_from_name` for each member.
Select reference trading day	`query_date = latest`.
Retrieve trading volume	`get_price(volume)`.
Construct cross-sectional sample	Collect volumes across constituents.
Compute group-level statistic	Arithmetic mean of volumes.

Table 16: Precious-metals fixing SOP vs. RealFin workflow (Shanghai gold benchmark, 2024-05-15).

Analyst SOP	RealFin golden workflow
Specify asset and benchmark	`get_price(gold)`.
Choose pricing sessions	`session = morning / evening`.
Fix calendar date	`query_date = 2024-05-15`.
Retrieve benchmark prices	Two calls to `get_sge_benchmark_price`.
Compute intraday difference	$P_{\text{morning}} - P_{\text{evening}}$.
Reporting format	Return price difference.

Metric	Rate	Meaning
Retrieval fail	0.72	gold tools not retrieved
Orchestration fail	0.03	wrong order / composition
Param/temporal fail	0.17	wrong date or arguments
Full* success	0.08	correct final answer
TSR	0.15	all gold tools selected
OSR	0.03	Full* workflow correct

Table 18: All values are proportions over evaluated samples. OSR $\ll$ TSR indicates difficulty lies in multi-tool orchestration rather than single-tool use.

Model	Context	ACC	TSR	OSR	PPS
Claude	Short	0.300	0.500	0.067	0.467
Claude	Long	0.033	0.167	0.000	0.033
Qwen	Short	0.033	0.333	0.100	0.300
Qwen	Long	0.000	0.133	0.000	0.067

Table 19: Short vs. long context (fixed tool set; 30 tasks each).

Model	k	ACC	TSR	OSR
Qwen	0	0.1016	0.1040	0.0000
Qwen	3	0.1406	0.4844	0.1094
Qwen	6	0.0781	0.0938	0.0000
Qwen	10	0.0469	0.0859	0.0000

Table 20: Intervention on tool-space size (k distractors), fixed prompt context.

Table 21: Example human evaluation item

ID	152
Question	In 2023 the year-on-year growth rate of the M2 money supply was 12.6% in January and changed to another value in December. Compute the change between January and December to two decimal places and report it as a percentage.
Golden answer	-1.8%
Reference program	abbreviated for space
Task type Time sensitivity	T2_Compare Fixed
Example ratings	Question 2 Answer 2 Workflow 2 Tool schema 2